

Consumption, Income Dynamics and Precautionary Savings

Luigi Pistaferri

Thesis submitted for the degree of Ph.D. in Economics

University College, London

1999

ProQuest Number: U643486

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest U643486

Published by ProQuest LLC(2016). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

ABSTRACT

I study the dynamics of household consumption and income. My thesis consists of four papers, detailed as follows.

In the first paper, I test for precautionary savings and excess sensitivity of consumption using a panel of Italian households that has measures of income expectations. The latter provide a powerful instrument for predicting income growth. The empirical analysis allows for a fairly general specification for the stochastic structure of the forecast error. I find that consumption growth is positively correlated with the expected variance of income and uncorrelated with predicted income growth.

In the second paper, I test for the saving for a rainy day hypothesis using the same set of subjective income expectations described above. According to the permanent income hypothesis, household savings should only react to transitory income shocks, as permanent shocks are entirely consumed. I show how subjective expectations can help to identify separately the transitory and the permanent shocks to income, thus providing a powerful test of the theory.

In the third paper, I notice that the theory of full consumption insurance implies absence of consumption mobility between any time periods. This implication requires knowledge of the evolution of the entire consumption distribution. I test this unexplored prediction of the theory using a panel of Italian households. I design a non-parametric test and find substantial mobility of consumption even controlling for possible preference shifts and measurement error. The findings strongly reject the theory of full consumption insurance.

In the final paper I model the conditional variance of income shocks, that is, the appropriate measure of risk emphasised by the theory. I first discriminate

amongst various models of earnings determination that separate income shocks into idiosyncratic transitory and permanent components. I allow for education specific differences in the stochastic process for earnings. The empirical analysis is conducted using data drawn from the 1967-1991 Panel Study of Income Dynamics.

Contents

1	Introduction and background issues	11
	Motivation	11
	Consumption and income dynamics.....	13
	Income dynamics and precautionary savings	15
	Consumption insurance	18
	Data sources	19
	The Bank of Italy Survey of Household Income and Wealth.....	20
	The Michigan Panel Study of Income Dynamics	21
2	Subjective expectations and the excess sensitivity puzzle	24
	Introduction.....	24
	Review of the literature and motivation.....	26
	Predicting income growth	27
	The conditional variance of consumption growth	29
	Non-separability between consumption and leisure.....	30
	The stochastic structure of the forecast errors	31
	Predicting income growth and consumption risk with subjective expectations.....	32
	Expected income growth	32
	Income risk	36
	Sample and specification issue.....	37
	Euler equation estimates	40
	Conclusions	45
3	Superior information, income shocks and the permanent income hypothesis	51
	Introduction.....	51

The estimation strategy	53
Income shocks decomposition	53
The effect of transitory and permanent income shocks on savings	54
Identification	55
Consistency	57
Testing for quadratic preferences	57
Chamberlain's <i>critique</i>	58
The data and the actual implementation of the test	59
The empirical distribution of the income shocks	61
Empirical results	63
Testing for Chamberlain's <i>critique</i>	68
Conclusions	69
4 Consumption insurance or consumption mobility?	76
Introduction	76
Consumption insurance	79
Consumption mobility	80
Extensions	82
Preference specification and heterogeneity	82
Measurement error	85
The data	86
Empirical results	87
Full sample estimates	87
Preference specification and heterogeneity	89
Measurement error	91
Sub-sample estimates	95
Conclusions	96
5 Income risk dynamics and heterogeneity	103
Introduction	103

Modelling the distribution of earnings	107
A brief survey of the literature	107
Aspects of the distribution of earnings	109
The conditional mean of earnings	110
The conditional variance of earnings	114
The validity of the instruments	119
Asymmetry and the lack of identification	120
The data	121
Empirical results	123
Model A	123
Model B	126
The conditional mean	126
The structure of the error term	129
The identification of the ARCH coefficients	131
The conditional variance of the transitory shock	133
The conditional variance of the permanent shock	134
The implications of the empirical estimates	136
A general discussion	136
Precautionary savings and insurance	139
Conclusions	143
6 Conclusions	155
Bibliography	162

No part of this thesis has been presented before to any University or College for submission as part of a higher degree.

Luigi Pistaferrì

ACKNOWLEDGMENTS

I am greatly indebted to Costas Meghir, my Principal Supervisor, and Orazio Attanasio, my Subsidiary Supervisor at University College London, for assiduous suggestions, intuitions and encouragement. Tullio Jappelli deserves a special mention for his generosity and constant support.

Part of this thesis draws from joint work with Tullio Jappelli and Costas Meghir. I would like to thank James Banks, Richard Blundell, François Bourguignon, Christopher Carroll, Victoria Chick, Angus Deaton, Mika I. Kuusmanen, Hamish Low, Sidney Ludvigson, Mario Padula, Jonathan Parker, Christina Paxson, Andrea Tiseno, Guglielmo Weber, two anonymous referees and seminar participants at various universities in Europe and the US for useful suggestions and comments.

For intellectual stimulus and moral support received at various stages of my doctorate, I would like to thank my parents and siblings, Kazu Araki, Luca Benedini, Eve Caroli, Giacinta Cestone, Teresa Delgado, Enzo Formisano, Charles Grant, Luigi Guiso, Marco Pagano, Monica Paiella and Alfonso Rosolia.

Financial support received from Istituto Universitario Navale of Naples, the Italian Ministry of Universities and Scientific Research, the James Lanner Memorial Fund, and the Ente per gli studi monetari bancari e finanziari “Luigi Einaudi” is gratefully acknowledged.

To Marina, above all.

1

Introduction and background issues

1 Motivation

In this dissertation I analyse consumption and savings decisions in the presence of individual income dynamics. In this introductory chapter I will motivate my line of research and discuss the theoretical models that lie in the background of the empirical analyses to be presented in Chapters 2 through 5. The discussion is not meant to be exhaustive; the modern theory of intertemporal consumption choice has been recently surveyed, among others, by Deaton (1992b), Browning and Lusardi (1996) and Attanasio (1999); I refer the interested reader to these excellent pieces of work for a more detailed theoretical treatment.

Testing for the validity of the permanent income hypothesis is complicated for a number of reasons. First, heterogeneity in consumption choices and individual income is likely to be paramount. This has not only prompted several studies based on microeconomic rather than aggregate data, but also started a new line of research in the macroeconomics literature, as recently remarked by Browning *et al.* (1999). While studies based on aggregate data are plagued by aggregation problems (Attanasio and Weber, 1995), microeconometric studies

face data-related problems, the most relevant one probably being the so-called Chamberlain's *critique* (1984). This is essentially a problem deriving from the inconsistency of empirical estimates in short panels.

I approach these problems with some ingenuity. First, I note that Chamberlain's *critique* can be attenuated, if not entirely controlled for, by exploiting data on individual subjective expectations. Second, heterogeneity in preferences and individual resources can be accounted for in a fairly general way. In particular, I consider the possibility of accounting for heterogeneity in income uncertainty. As I will argue in chapter 5, this may have far reaching implications not only on partial equilibrium individual consumption choices but also on the general equilibrium of the economy.

I consider testing for three different theoretical propositions. While tests for such propositions are not novel in the empirical literature, my approach is original because (i) I exploit data that are rarely available to the econometrician, and (ii) the actual format of my tests is different from that proposed so far in the literature. According to the first proposition, consumption growth does not react to predicted income growth; a test of this proposition, also known as *excess sensitivity test*, is presented in chapter 2. According to the second proposition, consumption reacts very strongly to unanticipated permanent income changes, but very weakly to unanticipated transitory income changes; this is the test I perform in chapter 3. Finally, the third proposition states that consumption does not react to idiosyncratic shocks but only to aggregate unanticipated fluctuations; this is the essence of the full consumption insurance hypothesis. A test for such proposition is presented in chapter 4. The modern theory of consumption choices under incomplete markets provide a rationale for the precautionary motive for saving, i.e. individuals will save in the face of uncertain income prospects. In this context, it is important to model and measure income uncertainty. I analyse statistical models for income risk in chapter 5.

2 Consumption and income dynamics

I will start by making explicit the relationship between consumption choices and income dynamics. Consider the standard problem solved by an infinitely lived agent i :

$$\max_{\{c_{it+\tau}\}_{\tau=0}^{\infty}} E_t \sum_{\tau=0}^{\infty} \frac{u(c_{it+\tau})}{(1+\delta)^\tau} \quad (1)$$

subject to the budget constraint:

$$a_{it+\tau+1} = (1+r)(a_{it+\tau} + y_{it+\tau} - c_{it+\tau}) \quad (2)$$

and the transversality condition (“no-Ponzi-game” condition):

$$\lim_{\tau \rightarrow \infty} \frac{a_{it+\tau}}{(1+r)^\tau} = 0 \quad (3)$$

Initial wealth a_{it} is given. The agent is assumed to hold preferences over non durable consumption, $c_{it+\tau}$. The first order conditions for this problem lead to the Euler equation:

$$u'(c_{it}) = \frac{1+r}{1+\delta} E_t u'(c_{it+1}) \quad (4)$$

To specialize the solution, I assume that preferences are quadratic, so that $u(c) = -\left(\alpha c - \frac{\beta}{2} c^2\right)$, and $r = \delta$ for simplicity. This yields the traditional prediction that consumption is a martingale, i.e. that the change in consumption is an innovation, and therefore orthogonal to all past information:

$$\Delta c_{it+1} = \psi_{it+1} \quad (5)$$

as in Hall (1978). A more complicated version of this orthogonality condition is considered and tested in chapter 2. The Euler equation and the intertemporal budget constraint imply the following consumption function:

$$c_{it} = \frac{r}{1+r} \left[a_{it} + \sum_{\tau=0}^{\infty} \frac{E_t(y_{it+\tau})}{(1+r)^\tau} \right] \equiv y_{it}^P \quad (6)$$

where y_{it}^P is the standard definition of permanent income (see Deaton, 1992b). The consumption innovation will in general depend upon the sources of uncer-

tainty of the model. If income is the only source of uncertainty of the model, one can derive the exact form of the consumption innovation ψ . In particular:

$$\psi_{it} = \frac{r}{1+r} \sum_{\tau=0}^{\infty} \frac{(E_t - E_{t-1}) y_{it+\tau}}{(1+r)^\tau} \quad (7)$$

To obtain a relationship between consumption innovation and income innovations one needs to specify a stochastic process for income. If income is well described by an *ARMA*(1, 1) process of the form:

$$y_{it+\tau} = u_{it+\tau} \quad (8)$$

$$u_{it+\tau} = \rho u_{it+\tau-1} + \varepsilon_{it+\tau} - \theta \varepsilon_{it+\tau-1} \quad (9)$$

with $\rho < 1$, it follows that income is a stochastic process stationary in levels and the consumption innovation can be written as:

$$\psi_{it} = \frac{r(1+r-\theta)}{(1+r)(1+r-\rho)} \varepsilon_{it} \quad (10)$$

Of course the two special cases of *AR*(1) and *MA*(1) are obtained by setting $\theta = 0$ and $\rho = 0$, respectively. Whether a shock is transitory or permanent depends upon the value of ρ . In particular, if $\rho = 0$ all shocks to income are transitory; if, instead, $\rho \rightarrow 1$, all shocks become permanent. The marginal propensity to consume out of the income shock is then implicitly governed by the amount of persistence in the income process, measured by the autoregressive coefficient ρ . This is somehow inconvenient, as it does not allow a full distinction between shocks that are mean reverting and shocks that are not. To this aim, one needs to make explicit the distinction between the two. One option available is to consider income as the sum of two components:

$$y_{it+\tau} = p_{it+\tau} + \varepsilon_{it+\tau} \quad (11)$$

where $p_{it+\tau}$ is a permanent component and $\varepsilon_{it+\tau}$ a serially uncorrelated transitory component. While stochastic, the permanent component changes very slowly, if at all. The standard assumption is that it follows a martingale process of the form:

$$p_{it+\tau} = p_{it+\tau-1} + \zeta_{it+\tau} \quad (12)$$

where $\zeta_{it+\tau}$ is serially uncorrelated. Thus:

$$y_{it+\tau} = y_{it+\tau-1} + \zeta_{it+\tau} + \varepsilon_{it+\tau} - \varepsilon_{it+\tau-1} \quad (13)$$

Transitory shocks to individual productivity include overtime labour supply, piece-rate compensation, bonuses and *premia*, etc.; in general, such shocks are mean reverting, e.g. their effect does not last long. On the other hand, some of the innovations to earnings are highly persistent or non-mean reverting, e.g. their effect cumulates over time. Example of permanent innovations are associated to job mobility, long-term unemployment or promotions. The original decomposition of income shocks into transitory and permanent components dates back to Friedman (1957). This practice has become quite standard. Quah (1990) shows that distinguishing between shocks of different nature may also provide an explanation for the *excess smoothness puzzle*.

With this process, and the additional assumption that the transitory and the permanent shock are uncorrelated at all leads and lags, the consumption innovation is:

$$\Delta c_{it} \equiv \psi_{it} = \frac{r}{1+r} \varepsilon_{it} + \zeta_{it} \quad (14)$$

The optimal rule for the agent is to respond on a one-for-one basis to shocks that alter the permanent component of income, but to a much lower extent to shocks that affect income only transitorily (namely, to an amount given by the annuity value). This is the income process that I consider in Chapter 3. A more complicated version of this income process (in logs) is specified and tested in chapter 5.

3 Income dynamics and precautionary savings

This section draws on Caballero (1990) and specializes to the case where labour income is the sum of a martingale permanent component and a serially

uncorrelated transitory shock (i.e., equations 11-12). As in Caballero, the only source of uncertainty of the model is that related to labour income. Suppose that the problem under study is again given by (1-3), with a_{it} given.

Assume however that preferences are of the CARA type: $u(c) = -\frac{1}{\theta}e^{-\theta c}$. The main advantage of CARA preferences is that an analytic solution for consumption can be easily derived even under precautionary savings (which obtains if $u'''(.) > 0$, a case contemplated by CARA or CRRA preferences, but not by quadratic preferences).¹ On the other hand, this assumption has at least two problems. The first is that a solution of negative consumption can be optimal for the agent (see Weill, 1992); the second is that risk aversion is, by assumption, constant. The reaction to risk is therefore unaffected by the level of wealth. While the first problem is much more difficult to tackle,² one might think of attenuating the second by stratifying the available sample according to some exogenous characteristics that make the assumption of constant risk aversion less untenable within *strata*. For instance, the sample could be stratified according to the level of education of the head or initial family income. In both cases, wealth levels will be presumably less dispersed.

Assuming $r = \delta$, the Euler equation of this problem is:

$$e^{-\theta c_{it}} = E_t \left[e^{-\theta c_{it+1}} \right] \quad (15)$$

The *feed-back* solution that satisfies (15) is:

$$c_{it+1} = \Gamma_{it} + c_{it} + \psi_{it+1} \quad (16)$$

where $\Gamma_{it} = \frac{1}{\theta} \ln E_t \left[e^{-\theta \psi_{it+1}} \right]$, and ψ_{it+1} is again the consumption innovation. To find the distribution of the latter, one needs to specify income dynamics.

The *ex-post* intertemporal budget constraint, $\sum_{\tau=0}^{\infty} (1+r)^{-\tau} (c_{it+\tau} - y_{it+\tau}) = a_{it}$, can be alternatively written as follows (Caballero, 1990):

¹However, if preferences are CARA the consumption function fails to be concave in the absence of interest rate risk (Carroll and Kimball, 1996).

²Inserting a non-negativity constraint on consumption amounts to impose a never declining consumption profile.

$$\sum_{\tau=0}^{\infty} \frac{c_{it+\tau}}{(1+r)^\tau} - \sum_{\tau=1}^{\infty} \frac{[y_{it+\tau} - E_t(y_{it+\tau})]}{(1+r)^\tau} = a_{it} + \sum_{\tau=0}^{\infty} \frac{E_t(y_{it+\tau})}{(1+r)^\tau} \equiv \left(\frac{1+r}{r}\right) y_{it}^P \quad (17)$$

Using the income process (11)-(12), one obtains:

$$\begin{aligned} c_{it} \sum_{\tau=0}^{\infty} (1+r)^{-\tau} + \sum_{\tau=1}^{\infty} (1+r)^{-\tau} \sum_{j=1}^{\tau} \Gamma_{it+j-1} + \sum_{\tau=1}^{\infty} (1+r)^{-\tau} \sum_{j=1}^{\tau} \psi_{it+j} \\ - \sum_{\tau=1}^{\infty} (1+r)^{-\tau} \sum_{j=1}^{\tau} \varepsilon_{it+j} - \sum_{\tau=1}^{\infty} (1+r)^{-\tau} \zeta_{it+\tau} = \\ \left(\frac{1+r}{r}\right) y_{it}^P \end{aligned} \quad (18)$$

Taking expectations conditional on the information available at time t and rearranging yields the consumption function:

$$c_{it} = y_{it}^P - \frac{1}{1+r} \sum_{\tau=0}^{\infty} \frac{E_t(\Gamma_{it+\tau})}{(1+r)^\tau} \quad (19)$$

which differs from (6) because of a term that takes into account the effect of the dispersion of the path of future income on current consumption: more risk will depress current consumption and prompt savings. The relation between the consumption innovation $\psi_{it+\tau}$ and the labour income innovation can be identified by replacing back c_{it} in (18), obtaining the condition:

$$\frac{1+r}{r} \sum_{\tau=1}^{\infty} \frac{(\psi_{it+\tau} - \varepsilon_{it+\tau})}{(1+r)^\tau} - \sum_{\tau=1}^{\infty} \frac{\zeta_{it+\tau}}{(1+r)^\tau} + \frac{1}{r} \sum_{\tau=1}^{\infty} \frac{\Gamma_{it+\tau} - E_t(\Gamma_{it+\tau})}{(1+r)^\tau} = 0 \quad (20)$$

In the absence of heteroschedasticity (i.e., $E_t(\zeta_{it+\tau}^2) = \sigma_\zeta^2$ and $E_t(\varepsilon_{it+\tau}^2) = \sigma_\varepsilon^2$ for all τ), the only relation between consumption innovations and labour income innovations that satisfies the expression (20) is therefore:

$$\psi_{it+\tau} = \zeta_{it+\tau} + \frac{r}{1+r} \varepsilon_{it+\tau} \quad (21)$$

Given that: $\Gamma_{it} = \frac{1}{\theta} \ln E_t[e^{-\theta\psi_{it+1}}]$, one can take a second-order Taylor expansion of $e^{-\theta\psi_{it+1}}$ around the conditional mean of ψ_{it+1} to obtain:

$$e^{-\theta\psi_{it+1}} = 1 - \theta\psi_{it+1} + \frac{\theta^2}{2}\psi_{it+1}^2 + rem. \quad (22)$$

taking expectations on both sides yields:

$$E_t e^{-\theta \psi_{it+1}} \simeq 1 + \frac{\theta^2}{2} E_t(\psi_{it+1}^2) \quad (23)$$

for infinitesimal risks one obtains:

$$\Gamma_{it} \simeq \frac{\theta}{2} \left[\sigma_\zeta^2 + \left(\frac{r}{1+r} \right)^2 \sigma_\varepsilon^2 \right] \quad (24)$$

The Euler equation can then be written as:³

$$\Delta c_{it+1} \simeq \frac{\theta}{2} \left[\sigma_\zeta^2 + \left(\frac{r}{1+r} \right)^2 \sigma_\varepsilon^2 \right] + \left(u_{it+1} + \frac{r}{1+r} \varepsilon_{it+1} \right) \quad (25)$$

This is the Euler equation that prevails in the presence of a precautionary motive for saving (Kimball, 1990). I estimate a version of equation (25) based on *CRRA* preferences and accounting for the variance term in chapter 2. In chapter 5 I consider the possibility that income shocks are conditionally heteroscedastic.

4 Consumption insurance

In this section I will describe the main insight of the full consumption insurance hypothesis, for which I will test in chapter 4. As emphasised by Cochrane (1990), full consumption insurance has little to do with self insurance through, say, precautionary savings. These are two different propositions: while the first involves substitution of consumption across states, the second involves substitution of consumption over time.

Suppose that agents have identical preferences of the *CRRA* type, $u(c) = (1 - \gamma)^{-1} c^{1-\gamma}$. Under complete markets, as is known, the social planner's solution will coincide with that obtained by considering the behavior of the fully decentralized economy. I will focus on the former for a matter of convenience. If the social planner maximizes a weighted sum of individual households' utilities, the Lagrangian of the problem can be written as (Deaton, 1997):

$$L = \sum_i \lambda_i \sum_s \sum_t \pi_{st} u(c_{ist}) + \sum_s \sum_t \mu_{st} \left(C_{st} - \sum_i c_{ist} \right)$$

³If $r \neq \delta$, the Euler equation admits a constant term $\frac{(r-\delta)}{\theta}$.

where i , s and t are subscripts for the agent i in the state of nature s in period t , λ_i is the social weight for agent i , μ_{st} is the Lagrange multiplier associated with the resource constraint, π_{st} the probability of state s in time period t , and C_{st} aggregate consumption in state s and time t .

The first order condition can be written in logarithms as:

$$-\gamma \ln c_{ist} = \ln \mu_{st} - \ln \lambda_i - \ln \pi_{st} \quad (26)$$

To obtain the rate of growth of consumption, one subtracts side-by-side from the expression at time $t + 1$:

$$\Delta \ln c_{it+1} = -\gamma^{-1} \Delta \ln \mu_{t+1} + \gamma^{-1} \Delta \ln \pi_{t+1} \quad (27)$$

where I drop the subscript s because only one state is realized in each period. The two terms on the right-hand-side of equation (27) represent aggregate effects. The first is the growth rate of the Lagrange multiplier, the second is the growth rate of the state probabilities. Note that first-differencing has eliminated all household fixed effects.

Versions of equation (27) have been used to test for the assumption of complete markets. If this assumption holds, only aggregate shocks but not idiosyncratic shocks should matter for consumption growth, as the latter are fully insured through state-contingent contracts or other informal mechanisms. I present a novel way to test for this implication in chapter 5, based on the observation that full consumption insurance implies absence of consumption mobility.

5 Data sources

In this final section I will briefly describe the contents and characteristics of the two surveys used in this dissertation. Due to space constraints, my analysis is not meant to be exhaustive; the interested reader is referred to Brandolini (1998) for more details on the Survey of Household Income and Wealth (SHIW), and

to Hill (1986) for information concerning the Panel Study of Income Dynamics (PSID).

5.1 The Bank of Italy Survey of Household Income and Wealth

The SHIW was conducted yearly beginning in 1966. Since 1987 is conducted every other year. The last available survey refers to 1995. In chapter 2 through 5 I will use SHIW panel data available from 1987 to 1995; I will thus omit the description of the characteristics of the survey before 1987. See Brandolini (1998) for more details concerning the historical development of the SHIW.

The basic SHIW sampling unit is the *de facto* family, i.e. a group of individuals linked by blood, marriage or affection, sharing the same dwelling and pooling totally or partially their resources. Institutional population is not included. Individuals who live together exclusively for economic reasons are not considered members of the same family, while in contrast only one unit is recorded in the case of extended families. For such reasons, SHIW-based estimates of average family size tend to exceed the corresponding Census estimates.

The design of the SHIW is such that the sampling procedure occurs in two stages, with municipalities selected in the first stage and families in the second. Municipalities are divided into 51 strata (i.e., 17 regions*3 classes of population size⁴). Municipalities in the first class are always included; the selection of those in the other two classes is random, with probability that has been proportional to demographic size in 1987 and constant since 1989. Families are selected from the registry office records. The sample size is roughly 8,000 in all the survey years I consider. Few modifications have affected representativeness; in 1987, for instance, there was an over-sampling of high-income units.

Since 1989 the SHIW includes a rotating panel component. The proportion of panel families has increased over time from 15 per cent (1989) to about 45 per cent (1995); some of the families are also re-interviewed in all last four surveys. Panel families are selected with criteria similar to those described above; however, since 1991 participation is on a voluntary basis, i.e. only families that express

⁴Classes of population size are: >40,000, between 20,000 and 40,000, <20,000 inhabitants.

their willingness to being re-interviewed are contacted.

The survey is carried out by a private company on behalf of the Bank of Italy. Data are collected in personal interviews, usually between March and May; data on income, consumption and wealth refer to the previous calendar year (in Italy this also coincides with the fiscal year). The gross response rate varies substantially over time; it was 60 per cent in 1987, but it dropped to less than 40 per cent in both 1989 and 1991; since then, has been slightly above 50 per cent. Questionnaires are assessed for reliability and consistency by the Bank of Italy statisticians.

Each survey contains special sections whose scope is to study in detail specific subjects (e.g. intergenerational transfers of wealth, use of health, educational and transport services, social mobility, subjective evaluation of working conditions and future income). The expenditure for durable and non durable goods are available in all years. Estimates of households real estate are also available for all surveys and are based on definitions kept largely unchanged over the years. The basic definitions of income and its components is similar to that used in compiling national accounts; income is recorded net of taxes and social security contributions.

The results of the SHIW are illustrated and commented in the official Bank of Italys publications; the tapes containing microdata are later released in public-use files available free of charge for institutions and researchers in Italy and abroad. As a result, the SHIW is extensively used to study various aspects concerning the behaviour of Italian households.

5.2 The Michigan Panel Study of Income Dynamics

The Panel Study of Income Dynamics (PSID) is a longitudinal study of a sample of US individuals (men, women, and children) and the family units in which they reside. The study is conducted by the Survey Research Center at the Institute for Social Research (University of Michigan).

The PSID started in 1968 with a national sample of 5,000 households. Of these, about 3,000 were representative of the US population as a whole (the

core sample), and about 2,000 were low-income families (the Census Bureau's SEO sample). Since then, individuals from the original households have been reinterviewed every year, irrespective of their living in the same dwelling or with the same people. Adults have been followed as they have grown older, and children have been observed as they advance through childhood and into adulthood, forming households of their own (so called *split-off* households). Information about the original 1968 sample individuals and their co-residents is collected each year (mainly by phone). In order to keep track of demographic changes in the underlying population due to immigration patterns and to increase the representativeness of the sample, in 1990 a representative national sample of 2,000 Latino households was added.

The PSID collects information at both family and individual level; moreover, it also includes information about the areas where sample units live (unemployment rates, food needs, etc.). The core of the survey is to collect data of economic and demographic content, with attention paid in particular to income sources and amounts, employment, family composition changes, and residential location; however, in some waves of the study a set of sociological or psychological questions are also asked. Information gathered in the survey applies to the circumstances of the family unit as a whole (e.g., type of housing) or to particular persons in the family unit (e.g., age, earnings). While some information is collected about all individuals in the family unit, the greatest level of detail is ascertained for the primary adults heading the family unit. Other important topics covered by the PSID include housing and food expenditures, housework time, and health status.

The data used in chapter 5 are drawn from the 1967-1991 family and individual-merged files of the PSID (waves I through XXV). Questions referring to labour income are retrospective; thus, those asked in 1968, say, refer to the 1967 calendar year. The earnings variable is the labour portion of money income from all sources; the variable name in the PSID tapes is "head's money income from labour" and includes the labour part of farm income and business income, wages, bonuses, overtime, commissions, professional practice, labour part of income

from roomers and boarders or business income.⁵ Education level is computed using the PSID variable with the same name.

⁵As noted by Gottshalk and Moffitt (1993), the measure of labour income available in the PSID has sources that may reflect capital income, such as the labour part of farm income and roomers and boarders.

2

Subjective expectations and the excess sensitivity puzzle

1 Introduction

An important implication of the permanent income hypothesis is that individual consumption growth should not respond to expected income growth. The certainty-equivalence version of the model also suggests that consumers do not react to income risk. But for applied economics the fundamental problem is measuring income risk and predicting future income on the basis of variables that are in individuals' information set and can be observed by the econometrician. In this chapter I test the theory of intertemporal consumers choices using data on subjective income expectations to predict realized income growth. The advantage is that no assumption about the process that generates income is required. In the Euler equation I also control explicitly for the potential effect of income risk, predictable changes in households labor supply, and the nature of the aggregate shocks.

The data are drawn from the 1989-1993 rotating panel of the Bank of Italy Survey of Household Income and Wealth (SHIW). The panel offers unique mea-

asures of subjective income and inflation expectations, an annual measure of non-durable consumption that is not affected by seasonality factors, and a wealth of information on financial and real assets. The availability of a good measure of assets is particularly useful for checking for the possibility of asymmetric response of consumption to predicted income growth.

To date, the panel component of the SHIW has not been extensively exploited for econometric purposes. For the purpose at hand, the main limitation of the panel is that it is relatively short. Even though over long periods of time the forecast error in consumption growth should be zero on average, in my case it may not. In short panels the null hypothesis that the coefficient of predicted income growth in the Euler equation is zero is a joint test of the orthogonality condition implied by the permanent income hypothesis and of the maintained assumptions about the particular stochastic structure of the forecast error. Rejection of the null could be attributed either to a failure of the theory or to the inconsistency of the estimator in short panels. My test must therefore be designed to tackle this important econometric problem.

In Section 2 I review the literature on excess sensitivity tests, motivate the methodology and describe how it differs from alternative approaches. The construction of subjective income expectation is presented in Section 3. Here I also compare income expectations with income realizations and discuss the validity of expectations as an instrument for predicting realizations. Data and specification issues are discussed in Section 4. Euler equation estimates, reported in Section 5, indicate that consumption growth is positively correlated with the expected variance of income growth, but uncorrelated with predicted income growth. To check for possible asymmetries in the response of consumption to predicted income growth, I also split the sample according to the level of assets (as in Zeldes, 1989a) and distinguish between positive and negative expected income growth (as in Shea, 1995). In short, I cannot reject the orthogonality conditions implied by intertemporal optimization, but can reject the certainty equivalence version of the permanent income hypothesis. Section 6 summarizes my main findings and how they can be reconciled with the institutional evidence

showing the pervasiveness of borrowing constraints in the Italian economy.

2 Review of the literature and motivation

Several authors have tested the permanent income hypothesis by estimating versions of the following Euler equation with panel data:

$$\begin{aligned} \Delta \ln c_{it+1} = & \alpha' \Delta F_{it+1} + \rho^{-1}(E_{it}r_{it+1} - \delta) + \\ & var_{it}(\Delta \ln c_{it+1} - \rho^{-1}r_{it+1}) + \beta E_{it} \Delta \ln y_{it+1} + \varepsilon_{it+1} \end{aligned} \quad (1)$$

where i is a household index, c_{it+1} a measure of non-durable consumption, F_{it+1} includes predictable indicators of households' preferences (such as age), r_{it+1} is the real after-tax rate of interest, ρ^{-1} the intertemporal elasticity of substitution, δ the rate of time preferences, E_{it} the expectation operator and ε_{it+1} the forecast error. Equation (1) can be derived exactly assuming that preferences are of the isoelastic form and that the distribution of the real interest rate and of consumption growth is jointly lognormal. Alternatively, it can be regarded as a second-order approximation to the first-order conditions of the consumer optimization problem.

Predicted income growth is often added to the Euler equation in order to test the orthogonality condition implied by intertemporal optimization, i.e. that $E_{it} \Delta \ln y_{it+1}$ should not help in explaining consumption growth ($\beta = 0$). It should be noted that the excess sensitivity test I perform has power against some, but not all, alternative consumption models. For instance, while myopic behavior will lead to excess sensitivity in every period, in a model with prudence and borrowing constraints the orthogonality condition may not be violated most of the time (and even perhaps all the time), as households save in the anticipation of future constraints. Empirically, it is very hard to distinguish between precautionary saving and models with liquidity constraints.

The empirical literature faces several serious problems in testing the restriction $\beta = 0$. First, it is difficult to find viable instruments for income growth that are truly exogenous and yet have good predictive power. Second, the condi-

tional variance of the uncertain components - consumption and the real interest rate - is difficult to observe and is therefore generally omitted from the estimation. Third, excess sensitivity may result from a failure to control properly for non-separability between consumption and leisure. Finally, excess sensitivity may also arise spuriously from the misspecification of the stochastic structure of forecast errors. I address these four problems in turn.

2.1 Predicting income growth

Testing for excess sensitivity requires reliable instruments to predict income growth. However, finding such instruments in panel data has proved to be extremely difficult, particularly in US studies. The Panel Study of Income Dynamics (PSID), which has been extensively used to estimate Euler equations, has relatively good data on income but information on consumption is limited to food expenditures. The Consumer Expenditure Survey (CEX) does give detailed measures of consumption, but the information on income is scanty and suffers from severe measurement error. Three approaches have been proposed to enhance the power of the instruments: out-of-sample information, two-sample instrumental variables techniques, and subjective income expectations.

Shea (1995) isolates a subset of households in the PSID whose heads can be matched to labor union contracts. Information on these contracts is then used to construct a measure of expected nominal wage growth. The latter is found to be strongly correlated with actual wage growth (a coefficient of 0.86 with a t -statistic of 3.8). Inflation expectations, however, are estimated on aggregate data through an autoregressive forecasting model. Shea then estimates an equation similar to (1) omitting the conditional variance term and replacing the income term with the expected real wage growth of the household head. He finds that expected wage declines affect consumption more strongly than expected wage increases, a result that is not consistent with either myopia or with the hypothesis that excess sensitivity is due to liquidity constraints.⁶ There are several problems with this

⁶Garcia, Lusardi and Ng (1997) apply a switching regression model with unknown sample separation to data drawn from the CEX and report a similar finding.

approach. One is that it assumes that the history of past inflation is known to each households in the sample. Another is that Shea ends up with a small sample (647 consumption changes drawn from 285 households), often resulting in poor standard errors, particularly if the sample is split according to the asset-income ratio.⁷ Finally, since only food consumption is available in the PSID, he requires an assumption of separability between food and other non-durable expenditures in the household utility function. Yet as Attanasio and Weber (1995) point out, this assumption is rejected in the CEX.

A second possibility is to enhance the power of the test by using two-sample instrumental variable techniques. Lusardi (1996) uses consumption data from the CEX and income data from the PSID, thus overcoming the problem of using just food consumption to estimate the Euler equation. The data are matched by a two-sample instrumental variable estimator. Nonetheless, the adjusted R^2 of the regressions of actual income growth on the instruments (demographic variables, education and occupation dummies) is only about 1 percent (see Lusardi, 1996, Table 4). Even though Lusardi finds evidence of excess sensitivity to predicted income growth, she does not investigate whether excess sensitivity arises from non-separabilities, myopia, liquidity constraints or other sources.

Hayashi's (1985) and Flavin's (1994) approach is the closest in spirit to the one I use in this study. Hayashi uses a unique data set of Japanese households reporting subjective expectations for income and consumption on a quarterly basis. He derives the theoretical covariance between the forecast errors in consumption growth and the subjective income expectations, estimates the parameters of the Euler equation by applying a minimum distance estimator and finds some evidence in favour of excess sensitivity. The procedure does not require assumptions about the nature of the aggregate shocks, and is therefore consistent even in short panels.

The 1967-69 US Survey of Consumer Finances used by Flavin contains a cate-

⁷The effect of expected real wage growth is never significantly different from zero in the regressions in Table 5, p. 195. When Shea splits expected income according to positive and negative expected wage growth he finds an implausible coefficient of 2.242 for expected wage decreases.

gorical variable about expectations of family income changes. These, in addition to lagged disposable income, are used as an instrument for income growth. Using a robust instrumental variable estimator to control for the presence of influential outliers, Flavin finds evidence of excess sensitivity for both high and low asset households. Evaluating the overall predictive power of Flavin's instrument is not easy, because first-stage results are not fully reported. Data are again problematic in this application. The Survey does not contain a consumption measure, which must therefore be inferred from income and assets. The sample size is small (774 observations), especially when the sample is split by assets.

2.2 The conditional variance of consumption growth

The conditional variance term in equation (1) is generally omitted from the estimation.⁸ This is correct only under the certainty equivalence version of the model, which implies that households do not react to the expected variance of consumption growth. However, if the utility function exhibits decreasing risk aversion, prudent households react to expected consumption risk by reducing consumption in period t relative to period $t + 1$, to an extent that depends on the degree of prudence.⁹ The reason the variance term is omitted in actual estimation is not that applied researchers believe in quadratic utility.¹⁰ Rather, it is that it has turned out to be extremely difficult to find suitable proxies for the conditional variance.

If the conditional variance term is omitted, one cannot of course test for quadratic preferences or estimate the degree of prudence. But the consequences of this omission could be far more serious. Ludvigson and Paxson (1997) point out that estimating a linearised Euler equation can bias the coefficient of the intertemporal rate of substitution. Furthermore, insofar as the conditional vari-

⁸A notable exception is Dynan (1993).

⁹Kimball (1990) defines absolute prudence as the ratio between the third derivative and the second derivative of the within-period utility function. With isoelastic preferences, relative prudence $-U''' / cU''$ equals 1 plus relative risk aversion.

¹⁰Research on precautionary saving is in fact steadily growing, see Browning and Lusardi (1996).

ance of consumption is correlated with $E_{it}\Delta \ln y_{it+1}$, the latter will proxy for the omitted effect of consumption risk, generating spurious evidence of excess sensitivity. Carroll (1992) goes one step further, and points out that even Zeldes' (1989a) sample splitting approach may produce spurious evidence in favour of liquidity constraints if one does not control properly for expected consumption risk. In fact, Zeldes' test consists in splitting the sample according to the asset-income ratio: if liquidity constraints are at the root of excess sensitivity, one should find no violation of the orthogonality conditions in the high-asset, and excess sensitivity in the low-asset group. But omitting the conditional variance term creates a spurious correlation between consumption growth and income that is stronger for low-wealth households. The reason is that rich households have greater capacity than poor ones to buffer income fluctuations by drawing down their assets, so that a finding of excess sensitivity in the group of poor households only - as in Zeldes - could be rationalized once the assumption of certainty equivalence is dropped by the theory of intertemporal choices.

There are two ways to solve the problem. One would be to estimate the non-linear Euler equation by the generalized method of moments. The second, which is used here, is to introduce explicit proxies for the conditional variance of consumption in the linearised Euler equation. This approach is more directly comparable with previous studies; it also allows me to use standard statistical tools to test if preferences are quadratic or if households react to expected income risk.

2.3 Non-separability between consumption and leisure

If leisure is an argument of the utility function, and if consumption and leisure are non-separable, today's consumption decisions will be affected by predictable changes in households' labor supply. This implies that consumption growth is positively correlated with predictable growth in hours of work. Since predicted growth in hours will almost surely correlate with predicted income growth, failure to control for labor supply indicators may lead to spurious evidence of excess sensitivity (that is, it could bias the estimated β coefficient upwards). This

point has been forcefully made by Attanasio and Weber (1995) and Meghir and Weber (1996) with CEX data. But the same authors also indicate a way out to this problem. Following their suggestions, I augment equation (1) with labor supply indicators.

2.4 The stochastic structure of the forecast errors

The disturbance term ε_{it+1} in equation (1) is a forecast error, the difference between realized and expected consumption growth. According to the permanent income hypothesis with rational expectations, the conditional expectation of a forecast error must be zero, i.e. $E_{it}(\varepsilon_{it+1}) = 0$. 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 will converge to zero as the dimension of the cross-section gets large. For instance, if the forecast error is the sum of an aggregate and of an idiosyncratic shock, then in a short panel the orthogonality condition fails even if the permanent income model is true: aggregate shocks induce a cross-sectional correlation between expected consumption growth and predicted income growth. 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 not evenly distributed in the population.

For this reason, excess sensitivity tests performed on short panels are in fact joint tests of the null hypothesis that $\beta = 0$ and 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.¹¹ Distinguishing between the two alternatives is difficult, unless the true structure of the forecast error is known. Yet, as will be seen, subjective expectations provide a guide to

¹¹Deaton (1992, p. 147-48) provides an example with non-additive aggregate shocks leading to spurious evidence in favor of excess sensitivity.

modelling the stochastic structure of the forecast error, thereby diminishing the problems one faces when testing for excess sensitivity with short panels.

3 Predicting income growth and consumption risk with subjective expectations

I estimate the Euler equation using the 1989-1993 panel section of the Bank of Italy Survey of Household Income and Wealth (SHIW). Details on sample design, response rates and timing of the interviews have been provided in chapter 1. Here I describe only the wording of the questions and the subjective expectations used to predict income growth and to proxy for consumption risk.¹² Several surveys contain subjective income expectations, but vary considerably as to the way expectations are elicited.¹³ 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.

3.1 Expected income growth

In 1989 and 1991 the following two questions were asked to each labour income recipient.

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

¹²Guiso, Jappelli and Terlizzese (1992) used the same SHIW questions to study the effect of earnings risk on 1989 saving and households' wealth. They also discuss the pros and cons of using subjective income expectations.

¹³The 1982 Japanese Survey of Family Consumption contains information about consumption and income expectations. The Dutch VSB Panel, the 1967 US Survey of Consumer Finances, and the US Survey of Economic Expectations (SEE) contain information on income prospects, but not on expected or actual consumption. Das and Van Soest (1997) and Dominitz and Manski (1998) using the VSB and the SEE, respectively, compare income expectations with realizations.

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 zero; 0-3; 3-5; 5-6; 6-7; 7-8; 8-10; 10-13; 13-15; 15-20; 20-25; >25 percent. If the response is “less than zero”, 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 labour 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), I set the upper bound of the distribution -the open interval- at 35 percent. 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:

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

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.¹⁴ Since I 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.

Let $E_{it}z_{it+1}$ denote the expected growth rate of nominal earnings or pension income, $E_{it}\pi_{it+1}$ the expected rate of inflation and $g_{it}^e = E_{it}z_{it+1} - E_{it}\pi_{it+1}$ the expected growth rate of real earnings. This is the instrument I use for $\Delta \ln y_{it+1}$, the actual growth rate of earnings of the household head. Although each labor

¹⁴More precisely, $x_0 = 0$ for those assigning zero probability to a negative earnings growth event, otherwise it is a value chosen by the respondent; $x_1 = 0.03$; $x_2 = 0.05$; $x_3 = 0.06$; $x_4 = 0.07$; $x_5 = 0.08$; $x_6 = 0.1$; $x_7 = 0.13$; $x_8 = 0.15$; $x_9 = 0.2$; $x_{10} = 0.25$; $x_K = x_{11} = 0.35$.

income recipient is asked to answer the survey questions, I rely only on the information provided by the head of the household. The reason is that in most cases information on income recipients other than the head is lacking. As I explain below, subjective expectations are also used to construct a measure of income risk, and use of data on income recipients other than the head would require making difficult assumptions about risk sharing arrangements within the household.

Table 1 compares nominal earnings expectations with realizations by demographic and household-income groups. In comparing expectations with realizations, it must be stressed that respondents report forecasts for the 12 months following the day of the interview. Interviews were taken between May and July of 1990 for the 1989 survey, and between May and October 1992 for the 1991 survey,¹⁵ whereas income realizations refer to the calendar years 1989, 1991 and 1993. Thus I use as instrument the one-year forecast of income growth given in May-July 1990 for the growth rate of earnings between 1989 and 1991 and the one-year forecast of earnings given in May-October 1992 for the growth rate of earnings between 1991 and 1993. This implies that expectations and realizations do not coincide in time, and are not immediately comparable.

In an instrumental variables context, this is not a concern. All that is needed is that the expectation be correlated with actual income growth and uncorrelated over time with the innovation term of the Euler equation (1). Under the null hypothesis of the permanent income model, the latter condition is met. My approach is valid even if individuals underestimate or overpredict future income: all I need is that expected income growth helps predicting actual income growth. In the next section I show that income expectations are indeed strongly correlated with realizations. Here I limit myself to a descriptive analysis.

Only if incomes grew steadily over the two-year span one would expect subjective predictions to mirror half of the actual income growth rate. The last row of

¹⁵SHIW interviews usually start in May, with households asked about their income, assets and consumption of the previous calendar year. The reason 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.

Table 1 suggests that while in 1990-91 expectations are quite close to realizations (5.7 against 5.2 percent), in 1992-93 expectations overpredict realizations (3.6 against 2 percent). Subjective expectations can be criticized because respondents may not fully understand the survey questions: households with better education might therefore give more accurate income forecasts simply because they understand the survey questions better. However, individuals with less education do not appear to answer the survey questions less accurately than those with more. For instance, in 1989 individuals with junior high school or less report an average expectation of 5 percent (*vis-à-vis* a realization of 5.5 percent), while individuals with college degrees overpredict income growth (7.5 percent *vis-à-vis* 5 percent). In 1991 it is the group with higher education that makes better forecasts. One explanation of the discrepancy between expectations and realizations is the sharp and largely unanticipated 1993 recession. The explanation usually offered for the recession was strong fiscal contraction and pension reform enacted by the Government in the Fall of 1992 (after the survey was completed), raising taxes, cutting pension benefits and increasing contributions. The recession had different effects for various population groups, hitting particularly the self-employed and the residents of the South. As will be seen, I will exploit knowledge of the groups that suffered mostly from the recession in modelling the structure of the forecast error.

The pattern of expectations and realizations by population groups are also of interest. The young expect earnings to grow faster than the middle-aged and the elderly. Also employees predict their earnings growth more accurately than self-employed in both surveys. In part this is due to the fact that the self-employed have greater income volatility. Yet, comparison between subjective expectations and realizations for the self-employed is difficult, because this group experienced an income decline of 12 percent in 1992-93, due to the 1993 recession and tax increases. Finally, expectations by income quartile do not indicate that rich households predict earnings better than poor ones.

Table 2 displays inflation expectations. In both surveys, average expected inflation is roughly 7 percent, quite close to the forecasts in 1990 (for 1991) and

1992 (for 1993) of sophisticated econometric models and international institutions. Respondents average expectation for 1990-91 (7.2 percent) comes closer to the realized value of 6.8 percent than OECD's forecast for June 1990-June 1991 (5.4 percent). Results are reversed for the June 1992-June 1993 period: OECD projections are closer to realizations (4.2 percent and 4.8 percent, respectively), while individuals overestimate the actual rate (with average expectations of 7.2 percent).¹⁶ An interesting feature is that these average subjective inflation expectations do not in fact mask a great number of implausible extreme values. More than 50 percent of the sample bunches the entire probability distribution for inflation between 5 and 7 percent. Finally, there is no clear pattern of subjective expectations by region, age, education or income.

3.2 Income risk

In the Euler equation it is the term $var_{it}(\Delta \ln c_{it+1} - \rho^{-1}r_{it+1})$ that affects consumption growth. I assume that the only non-insurable risk faced by individuals is income risk, thus neglecting such other possibilities as rate of return and health risks. The subjective variance of the growth rate of real earnings is $\sigma_{i,g}^2 = \sigma_{i,z}^2 + \sigma_{i,\pi}^2 - 2\phi\sigma_{i,z\pi}$. I have data on the marginal distributions of z and π , but lack information on ϕ , the correlation coefficient between nominal earnings shocks and inflation shocks. Thus in the empirical analysis I rely mainly on the subjective variance of the growth rate of nominal earnings ($\sigma_{i,z}^2$) as my preferred proxy for expected consumption risk. One justification for this choice is that it avoids arbitrary assumptions about the value of ϕ ; furthermore, indexation clauses in labor contracts often provide insurance against inflation increases.

Only if utility is exponential and income is a random walk there is a one-to-one correspondence between income risk and consumption risk in the Euler equation (see chapter 1). Otherwise, the relation between the two is non-linear, depending

¹⁶One possibility for the larger gap between expectations and realizations in 1992 is that individuals were surprised by the implementation of income policies in July of 1992. These income policies are generally thought to have been effective in reducing the actual inflation rate. An alternative possibility is that consumers form adaptive expectations (in both 1989 and 1991 the inflation rate was 6.3 percent).

on the utility function and the income process. For this reason one cannot give a structural interpretation of the estimated coefficients, i.e. in terms of prudence or underlying preference parameters. I am also aware that my measure of income risk is open to criticism.¹⁷ For instance, I rule out the potential effect of other non-insurable risks faced by households. And yet if income risk is poorly measured, or if income risk is only poorly correlated with consumption risk, one should find no statistical relation between consumption growth and the subjective variance of income.

4 Sample and specification issue

The panel component of the SHIW includes 1,137 households interviewed in 1989 and 1991, 2,420 households interviewed in 1991 and 1993, and 1,050 households interviewed in 1989, 1991, and 1993. Defining an “observation” as two years of data, this corresponds to 5,657 potential observations (2,187 in the 1989-91 panel, and 3,470 in the 1991-93 panel). I drop cases in which the household head changed (355 observations); those with inconsistent data on age, sex, or education (515 observations); those lacking data on subjective expectations (1,123 observations); and those lacking data for other variables used in the empirical analysis (130 observations). The final sample therefore includes 3,534 “observations” (1,102 for 1989-91, and 2,432 for 1991-93). Since in most cases I have only one observation per household, I test primarily if the cross-sectional variation in consumption growth is explained by the cross-sectional variation in predicted income growth. I explain below how I deal with this problem.

As in previous studies, I control for individual preferences with age and change in family size.¹⁸ Testing for non-separabilities in the utility function is interesting in its own right and ensures that excess sensitivity does not arise from preference misspecification. Given that in my sample virtually no head is unemployed,

¹⁷Given the wording of the questions, the probability of low income states, such as unemployment, may not be reported.

¹⁸I also tried changes in other demographic variables, such as the number of adults or the number of children. In no case were the main results affected.

I introduce in the Euler equation the change in the employment status of the spouse. As mentioned, omitting labor supply indicators can bias upward the coefficient of expected income growth of the household head. The problem is not as serious than if I had total household earnings (employment is almost surely positively correlated with predicted income growth). However, the earnings of the head may still be correlated with the working spouse dummy because common macroeconomic shocks affect the probability of working and income prospects in the same direction. Other labor supply indicators - such as the change in the number of income recipients - were either not significantly different from zero or did not alter the results.¹⁹

As mentioned, one should control for the structure of aggregate shocks, particularly in short panels. Even though forecast errors in consumption are unobservable, I do observe the cross-sectional pattern of income innovations. This can be used to extract potentially useful information about the structure of forecast errors in consumption, which depends on the income innovations.²⁰ For instance, in the absence of common shocks, time dummies should not explain the forecast error. If instead macroeconomic shocks are important, time dummies will be correlated with the innovation in income and in consumption, and therefore cannot be used as instruments to predict income. Rather, one should allow for time effects in the Euler equation.

Preliminary analysis indicates that the income innovation ($\Delta \ln y_{it+1} - g_{it}^e$) is correlated not only with time dummies, but also with education, and dummies for occupation and region. Given the characteristics of the recessionary episode of 1993, I find it plausible to assume that the forecast error contains an aggregate

¹⁹Estimating the elasticity of intertemporal substitution has proven to be extremely difficult with panel data. Even in long panels - such as the PSID - the coefficient of the real interest rate is often poorly determined or implausible. Initially, I constructed a measure of the household-specific real interest rate, subtracting inflation expectations from the nominal rate on Treasury bills. However, the coefficient of the elasticity of intertemporal substitution thus obtained was not significantly different from zero and theoretically implausible. In the end, I decided to drop the interest rate from the regressions: using two-year consumption changes with one-period ahead inflation expectations, it is simply impossible to get the timing of the interest rate right.

²⁰See chapter 3.

component which is unevenly distributed across population groups and an idiosyncratic component that averages out in the cross-section.²¹ Tax increases for the self-employed or a stronger effect of the 1993 recession in the South would have such an effect (see also Miniaci and Weber, 1996). This implies that group dummies (such as region and employment status) should not be used as excluded instruments to predict actual income growth.

Table 3 reports the first-stage coefficients obtained by regressing actual income growth on expected income growth, time dummies, education, regional dummies and employment status interacted with year dummies, lagged employment status of the spouse, age, family size, and income risk. Overall, the first stage regression has good predictive power (the adjusted R^2 statistics is 0.07). The coefficient of expected income growth is 0.5 and significantly different from zero at the 1 percent level.²² A conventional F-test on the excluded instruments (expected income and lagged employment status of the spouse) yields a p -value below 1 percent, confirming the validity of the instruments.

In the following section I thus present instrumental variable estimates of the following Euler equation:

$$\begin{aligned} \Delta \ln c_{it+1} = & \alpha_1 age_{it+1} + \alpha_2 \Delta \ln FS_{it+1} + \eta \sigma_{i,z,t}^2 + \\ & \gamma \Delta ww_{it+1} + \beta \Delta \ln y_{it+1} + \theta_j \mu_{t+1} + \nu_{it+1} \end{aligned} \quad (2)$$

where FS_{it+1} denotes family size, Δww_{it+1} is the change in a dummy for spouse working full-time, $\sigma_{i,z,t}^2$ denotes the expected variance as of time t of nominal income growth, j the population groups affected by macroeconomic shocks, and θ_j captures the effect of unevenly distributed aggregate shocks μ_{t+1} on the forecast error in consumption.²³ In the empirical application I will also present estimates replacing predicted income growth $E_{it} \Delta \ln y_{it+1}$ with the subjective expectation

²¹In the empirical specification I thus assume that the forecast error in consumption growth can be decomposed as $\varepsilon_{it+1} = \theta_j \mu_{t+1} + \nu_{it+1}$, where ν_{it+1} denotes the idiosyncratic component.

²²Our instrument predicts well both income increases and income decreases. The first stage coefficients of expected income growth are, respectively, 0.45 and 0.64 in the samples expecting positive and negative income growth.

²³My identifying assumption is therefore $\text{plim}_{N \rightarrow \infty} N^{-1} \sum_i \nu_{it+1} = 0$, where N is the number of households.

of income growth .

5 Euler equation estimates

The results of estimating equation (2) are reported in column 1 of Table 4. The coefficients of the demographic variables are well determined and have the “right” sign. The positive and significant coefficient of the change in the spouses employment status indicates that expecting to work more in the future reduces current consumption. This will indeed be the case if leisure and consumption are non-separable. The coefficients of the group dummies are not reported for brevity.

The proxy for consumption risk is positive and significantly different from zero at the 1 percent level, and supports the theory of precautionary saving. Since what I measure is not the expected variance of consumption but the expected variance of income growth, the coefficient has no structural interpretation. Nevertheless, its size (5.67) is most suggestive. With isoelastic utility, prudence equals one plus relative risk aversion, and reasonable values for risk aversion vary between 1 and 10.

It is important to note that ignoring the group dummies induces a correlation between the cross-sectional variation in consumption growth and the cross-sectional variation in income growth leading to spurious evidence in favour of excess sensitivity. In fact, if one assumes that the forecast errors can be decomposed into an aggregate shock and an idiosyncratic shock, as in most of the literature (though I know it cannot from the pattern of the innovation in income), introducing time dummies in the Euler equation should provide consistent estimates. If education and dummies for region and occupation, in addition to expected income, are then used as instruments for income growth, one does find excess sensitivity (a coefficient of 0.32 with a t -statistics of 5). However, when the time dummies and their interactions with group dummies are added to the Euler equation (thus controlling for the structure of the forecast error) such evidence vanishes, as in Table 4. Note also that excluding the dummy for

working wife and the variance of income growth does not affect the excess sensitivity coefficient. Thus, in my sample there is no excess sensitivity even when the Euler equation is misspecified.

How should one interpret the role of group dummies and education in the Euler equation? Even though they were introduced as a device to eliminate the inconsistency of IV estimates in short panels, at least two other interpretations are possible. First, group dummies may account for preference shifts and for this reason should not be omitted from the Euler equation, otherwise income growth will simply proxy for the omitted variables (absent group dummies, excess sensitivity is just a signal of misspecified preferences). The second possibility is that there is a subtler form of excess sensitivity, arising not from the correlation between consumption and income, but from the correlation between consumption and income predictors. To clarify this point, suppose that (low) education, residence in the South and self-employment are predictors of the probability of being liquidity constrained in period t . If so, one may expect them to predict higher consumption growth between period t and $t + 1$. However, in the regressions of Table 4 the dummies for South and self-employment are negative, while the coefficient of education is positive (with the exception of the dummy for South in 1993, the other interaction terms are not statistically significant). While alternative explanations for the effect of group dummies are therefore possible, I find it more plausible to attribute their role to the effect of unexpected aggregate shocks.

So far, my sample has included farmers and the self-employed (854 observations). There are several reasons why it may be desirable to test the robustness of the results when these observations are excluded: reported income for the self-employed income is severely under-estimated (Brandolini and Cannari, 1994); some individuals may have chosen self-employment, a more risky occupation, because they are less risk averse than the rest of the population, inducing sample selection; for farmers it is not easy to measure income or to distinguish it from consumption. The first-stage regression excluding farmers and the self-employed is reported in column 2 of Table 3. The coefficient of expected income growth

increases to 0.67, indicating again that this variable is a powerful instrument to predict actual income growth. Column 2 of Table 4 replicates the regressions of column 1 using the restricted sample. There is again no evidence of excess sensitivity (column 2), and the other coefficients are only marginally affected.

An excess sensitivity coefficient of zero may hide possible asymmetric responses of consumption growth to predicted income growth. The well-known approach of Zeldes (1989a) is to split the sample according to the asset-income ratio. If liquidity constraints are the only source of failure of the model, one would find excess sensitivity in the low-asset but not in the high-asset group, in that affluent households can always overcome borrowing constraints by drawing on assets, while the less wealthy cannot. In Table 4 households are defined as “poor” if total net worth (including real estate wealth) does not exceed twice annual income. The sample split thus places about 30 per cent of the sample in the low-asset group and 70 per cent in the high-asset group. It is apparent that I find no evidence of excess sensitivity in either group (two insignificant coefficients of 0.23 and -0.03 in the low-asset and high-asset groups, respectively).²⁴

Under liquidity constraints the response of consumption to predictable income growth should be asymmetric (Altonji and Siow, 1987). If consumers expect their income to increase, they would like to borrow but are prevented from doing so: consumption growth will then respond to predicted income growth. If instead consumers expect income to fall, they will save, not borrow: in this case the liquidity constraint is not binding, and one should not find a violation of the orthogonality conditions.

My instrument for income growth offers an opportunity to test for the potential asymmetric response of consumption to expected income growth. For comparison with previous estimates, in column 5 of Table 4 I replace (instrumented)

²⁴Results are qualitatively unaffected if I split the sample according to the ratio of financial assets to income or if I vary (upwards or downwards) the threshold used to split the sample. In all cases the low-asset group tends to be younger, less educated, with fewer self-employed and lower income than the high-wealth group. Given that reducing the threshold used to split the sample reduces the group of low-asset households, the estimated coefficients tend to be less precisely estimated.

actual income growth with expected income growth. Given the endogeneity of Δww_{it+1} the equation is again estimated by instrumental variables, and the previous results are confirmed.²⁵ I then capture the potential non-linear effect of expected income growth estimating:

$$\begin{aligned} \Delta \ln c_{it+1} = & \alpha_1 age_{it+1} + \alpha_2 \Delta \ln FS_{it+1} + \eta \sigma_{i,z,t}^2 + \\ & \gamma \Delta ww_{it+1} + \beta_1 g_{it}^{e+} + \beta_2 g_{it}^{e-} + \theta_j \mu_{t+1} + \nu_{it+1} \end{aligned} \quad (3)$$

where g_{it}^{e+} denotes positive (or zero) expected income growth, and g_{it}^{e-} denotes negative expected income growth. In column 6 of Table 4 I do not find evidence of asymmetric effects: the coefficients of positive and negative expected income growth are 0.07 and -0.06, respectively, and are not significantly different from zero or from each other. The asymmetry test was replicated also splitting the sample by assets. Under liquidity constraints one should find excess sensitivity mainly in the group of poor households that expect an increase in income. However, even in this case I cannot reject the null hypothesis of no asymmetric effects (whether or not the self-employed are included in the sample).²⁶

I performed several tests to check the robustness of the results. Here I briefly comment on higher moments of the expected income growth variable, sample selection arising from non-responses, the definition of the sample, and alternative instruments to predict income growth.²⁷ The survey questions allow me to estimate higher moments of the conditional distribution of expected income growth, not just the variance, which is only a valid indicator of risk under restrictive assumptions. For instance, households may react more strongly to the risk of low

²⁵Since expectations are available only about bands of possible income and inflation values, my measure of income risk will entail a certain amount of measurement error. I replicate regression 5 in Table (4) by OLS, omitting the change in the employment status of the wife, with results basically unaffected. Since in an OLS context measurement error in an independent variable tends to bias the coefficients towards zero, I take this as an indication that measurement error cannot explain, alone, a significant coefficient of income risk. For the same reason, I cannot rule out that measurement error in expected income biases the excess sensitivity coefficient towards zero in columns (1) to (4).

²⁶Those who expect their income to decline are less wealthy, less educated, and more likely to be near to the retirement (or already retired).

²⁷For brevity these results are not reported.

income realizations. I thus introduced an index of asymmetry of the distribution of income growth and dummies for households that expected with relatively high probability (more than 20 percent) a large decline in income (more than 5 percent). These variables were not significantly different from zero.

My estimates may be criticized on the ground that the respondents reporting expectations presumably understand the survey questions better than those who do not. A formal test of this hypothesis can be made by controlling explicitly for selection bias arising from non-responses. I thus run a probit regression for the probability of response, assuming that the probability is related to demographic and economic variables (income, education, age, occupation, industry, and region of residence). The implied Mills ratio was then added as a regressor to the Euler equation. The ratio was not significantly different from zero and results were again similar to those reported in the basic specification, suggesting that this effect is not important. I also checked the stability of the coefficients with respect to several sample exclusions: individuals older than 40 or 50, households with more than two income recipients, and households whose head is a pension income recipient. In no case did the pattern of results change appreciably.

Finally, my conclusions are qualitatively unchanged if I use lagged income growth, rather than expected income growth, to predict actual income growth. For this purpose I must use the sub-sample of households surveyed in 1989, 1991, and 1993. Here I find again evidence for excess sensitivity if I do not control for the stochastic structure of the forecast error (a coefficient of 0.19 with a t -statistics of 2.1), but no excess sensitivity when education and group dummies (interacted with time) are introduced as additional regressors to the Euler equation (a statistically insignificant coefficient of -0.01). The problem with using lagged income growth is that if income is measured with error, the first lag of income growth is not a valid instrument, as measurement error violates the orthogonality conditions. The advantage of using expected income growth is that the instrument is valid whether or not income is measured with errors.

6 Conclusions

After more than a decade of studies testing the theory of households' intertemporal choices on panel data, the evidence is mixed (Browning and Lusardi, 1996). In this chapter I have tested for excess sensitivity using a 1989-93 panel of Italian households that provides measures of income and inflation expectations and income risk. The expectations are used as an instrument for predicting income growth. Controlling for income risk, predictable changes in employment status of household members, and for aggregate shocks that affect differently population groups, I find that consumption growth is uncorrelated with the expected earnings growth of the household head. I also find that predictable proxies of changes in labor supply and expected income risk affects positively consumption growth. To the extent that income risk is correlated with expected consumption risk, this finding supports the theory of precautionary saving.

My results are robust to a variety of experiments such as asymmetric response of consumption to positive or negative expected income growth and sample splits by assets. It is worth stressing that my result of no excess sensitivity depends on the validity of subjective income expectations to predict income growth. The correlation between the two is statistically significant, but the instrument may not be powerful enough to capture small departures from the permanent income hypothesis.

Given the severe imperfections of the Italian credit markets by the standards of other industrialized countries and the pervasiveness of various liquidity constraints, particularly in the mortgage market (Guiso, Jappelli and Terlizzese, 1994), the fact that I do not find excess sensitivity may come as a surprise, since often excess sensitivity has been linked to liquidity constraints. But it is precisely for this reason that Italian households are high savers, and even at young ages have accumulated considerable assets to buffer income fluctuations. This indicates that excess sensitivity tests have limited power against models in which borrowing constraints play an important role. For instance, prudent consumers will save in anticipation of future constraints, and may never exhibit excess sensitivity to predicted income growth. Consumers who are saving to purchase a house are globally constrained because they must meet a down-payment, but

the orthogonality condition does not fail, except perhaps at the time of the purchase. Thus my results should not be viewed as a contradiction that borrowing constraints play an important role in the Italian economy; rather, as evidence confirming how difficult it is to detect liquidity constraints in structural models of intertemporal choices by conventional excess sensitivity tests.

Table 1
Comparing expectations and realizations of nominal income growth²⁸

	1990-91 Expectation	1990-91 Average Realization	1991-92 Expectation	1991-92 Average Realization
<i>Age group</i>				
<35	0.0758	0.0588	0.0521	0.0361
35-55	0.0640	0.0475	0.0399	0.0178
>55	0.0426	0.0543	0.0306	0.0248
<i>Education</i>				
Junior high-school or less	0.0498	0.0559	0.0333	0.0101
High-school	0.0667	0.0428	0.0439	0.0317
University degree or more	0.0754	0.0493	0.0446	0.0716
<i>Occupation</i>				
Employed	0.0565	0.0623	0.0371	0.0446
Self-employed	0.0607	0.0043	0.0345	-0.1207
<i>Region of residence</i>				
North	0.0576	0.0435	0.0342	0.0246
Center	0.0541	0.0869	0.0306	0.0469
South	0.0582	0.0443	0.0429	-0.0049
<i>Household income</i>				
Iwe quartile	0.0455	0.1048	0.0381	0.0422
IIwe quartile	0.0620	0.0689	0.0351	0.0113
IIIwe quartile	0.0548	0.0310	0.0377	0.0171
IV quartile	0.0637	0.0129	0.0360	0.0111
<i>Total sample</i>	0.0573	0.0565	0.0367	0.0201

²⁸The table compares expectations and realizations of nominal income growth. The realization is the average growth rate over the two years. Expectations are given in May-July of 1990 (column 1) and May-October 1992 (column 3) for the subsequent 12 months. Income is defined as after-tax earnings and pension benefits of the household head.

Table 2
Inflation expectations²⁹

	1991-92	1992-93
	(1)	(2)
<i>Age group</i>		
<35	0.0719	0.0704
35-55	0.0722	0.0747
>55	0.0715	0.0698
<i>Education</i>		
Junior high-school or less	0.0732	0.0717
High-school	0.0693	0.0742
University degree or more	0.0714	0.0712
<i>Occupation</i>		
Employed	0.0720	0.0720
Self-employed	0.0700	0.0734
<i>Region of residence</i>		
North	0.0698	0.0749
Center	0.0663	0.0704
South	0.0760	0.0700
<i>Household income</i>		
Iwe quartile	0.0745	0.0744
IIwe quartile	0.0751	0.0730
IIIwe quartile	0.0679	0.0705
IV quartile	0.0708	0.0711
<i>Total sample</i>	0.0719	0.0722
<i>OECD Projection (Consumer prices)</i>	0.0540	0.0415
<i>Realization (Consumer prices)</i>	0.0680	0.0480

²⁹Inflation expectations are given in May-July of 1990 (column 1) and May-October 1992 (column 2) for the subsequent 12 months. OECD inflation projections are 12-month forecasts given in June 1990 (column 1) and June 1992 (column 2). Inflation projections and realizations refer to the same time periods. Source for inflation projections: *OECD Economic Outlook*, June 1990, vol. 46, p. 125-27; and June 1992, vol. 51, p. 127-29. Source for inflation realizations: *OECD Economic Outlook*, June 1994, vol. 55, p. 68-70.

Table 3
Predicting actual income growth³⁰

	Total sample	Excluding self-employed and farmers
	(1)	(2)
Expected income growth	0.5003 (0.1099)	0.6660 (0.1133)
Education*1991	-0.0040 (0.0029)	-0.0020 (0.0027)
Education*1993	0.0104 (0.0021)	0.0104 (0.0019)
Resident in the South*1991	-0.0813 (0.0370)	-0.1053 (0.0345)
Resident in the North*1993	-0.1110 (0.0248)	-0.0842 (0.0231)
Resident in the South*1991	-0.0840 (0.0372)	-0.1048 (0.0340)
Resident in the North*1993	-0.0404 (0.0239)	-0.0331 (0.0220)
Self-employed*1991	-0.1242 (0.0334)	
Self-employed*1993	-0.3372 (0.0248)	
Farmer*1991	0.0842 (0.0696)	
Farmer*1993	0.0291 (0.0349)	
Working spouse	0.0156 (0.0171)	0.0027 (0.0160)
Sample size	3,534	2,680
Adj. R^2	0.0708	0.0270
Adj. R^2 on excluded instruments	0.0055	0.0109
F-test	10.78	15.78
(degrees of freedom)	(2; 3,531)	(2; 2,677)

³⁰The dependent variable is the growth rate of real after-tax earnings and pensions of the household head. Standard errors are reported in parenthesis. Each regression also includes a constant term, a time-dummy, age, change in family size and the variance of income growth. Column 2 excludes farmers and the self-employed.

Table 4
Euler equation estimates³¹

	Baseline specification		Splitting the sample by the wealth-income ratio		Using exp. income as a regressor	
	Total sample	Excluding self-empl. and farmers	Low wealth	High wealth	Expected income replaces actual income	Asymmetry test
	(1)	(2)	(3)	(4)	(5)	(6)
Age	0.0009 (0.0005)	0.0013 (0.0006)	0.0007 (0.0007)	0.0011 (0.0007)	0.0013 (0.0005)	0.0009 (0.0005)
$\Delta \ln(\text{family size})$	0.3405 (0.0533)	0.3334 (0.0583)	0.3706 (0.0810)	0.3375 (0.0656)	0.3442 (0.0510)	0.3441 (0.0510)
$\Delta \text{working spouse}$	0.3391 (0.0693)	0.4156 (0.0814)	0.2977 (0.1028)	0.3314 (0.0902)	0.3447 (0.0667)	0.3444 (0.0667)
Var.of inc. growth	5.6719 (1.9744)	5.9123 (1.7715)	9.1301 (20.4515)	5.3033 (1.8285)	5.6442 (1.9273)	5.3578 (1.9566)
$\Delta \ln y_{it+1}$	-0.0835 (0.1928)	-0.0514 (0.1469)	0.2309 (0.7714)	-0.0340 (0.1532)		
Exp.inc. growth					-0.0418 (0.0924)	
Exp.inc. increase						0.0522 (0.2563)
Exp.inc. decline						-0.0687 (0.1051)
Sample size	3,534	2,680	1,108	2,426	3,534	3,534

³¹The dependent variable is the growth rate of non-durable consumption expenditures. $\Delta \ln y_{it+1}$ is the after-tax real growth rate of earnings and pensions of the household head. Each regression also includes time dummies, interaction of education with year, and interactions of year and dummies (dated t) for region, self-employed and farmer (omitted in column 2). In columns 1-4 the instruments used are expected income growth and the lagged employment status of the spouse. In columns 3 and 4 an observation is included in the low-asset group (high-asset group) if, at the beginning of the period, the wealth-income ratio is smaller (greater) than 2 (wealth is real estate plus financial assets less household debt). Standard errors corrected for heteroscedasticity of unknown form are reported in parenthesis.

3

Superior information, income shocks and the permanent income hypothesis

1 Introduction

According to the textbook version of the permanent income hypothesis, household consumption responds on a one-for-one basis to permanent income shocks but is nearly insensitive to transitory income shocks. Equivalently, households save for a rainy day the transitory component of the income innovation and consume entirely the permanent one. By and large, testing for the separate effect of income shocks on consumption or saving has proved difficult; the main problem is that while the agent may be *subjectively* able to discriminate between a transitory and a permanent shock, the econometrician is not. As a result, econometric identification of separate income shock components remains infeasible.³²

In this chapter I show that combining subjective income expectations with income realizations can help to identify *separately* and *exactly* the transitory and the permanent shock to income. This allows not only to examine the cross-section

³²Attempts 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 US data.

distribution of separate income shocks (a possibility that would be unthinkable in the presence of data on income realisations alone), but it also provides a neat test of the theory.³³ In particular, I test whether households “save for a rainy day” using data available for a sample of Italian households drawn from the 1989-1991 Bank of Italy Survey of Household Income and Wealth (SHIW).

My estimation strategy has two advantages *vis-à-vis* previous empirical studies. First, 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. Second, the direct observability of one’s superior information set minimizes the problem of low power of the instruments used to test the theory.

To assess the validity of the permanent income hypothesis, I 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 preferences admit prudence, precautionary saving can represent a likely source of failure of the theory. In fact, 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 their realization. But this also suggests that one might test for the deviation from the certainty equivalence assumption augmenting the “saving for a rainy day” equation with the subjective variance term which can be constructed from the data: if the permanent income hypothesis with quadratic preference is true, the subjective variance should not explain saving.

The rest of the chapter is organised as follows. Section 2 presents a formal decomposition of the income shocks into a permanent and a transitory component and shows how subjective expectations of income can help to identify separately the two components and provide a simple test of the permanent income hypothesis. Section 3 presents the data used in this study, while in section 4 I examine

³³Campbell (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 the one available to the econometrician. While consistent under some regularity conditions (see the discussion in Deaton, 1992a, and Flavin, 1993), estimates based on the econometricians information set are inefficient.

the empirical distribution of shocks in my sample. Section 5 presents the results of the empirical analysis. When heterogeneity in income growth is ignored, the evidence is weakly supportive of the permanent income hypothesis with a precautionary motive for saving. In particular, savings do respond to transitory income shocks, but also to permanent income shocks and higher moments of the distribution of earnings. However, when heterogeneity is accounted for, the results are not literally consistent with the permanent income hypothesis. In particular, the effect of transitory shocks on saving is much tinier, a finding that can be reconciled either with high real interest rates or with the existence of binding liquidity constraints affecting a sizeable proportion of the population. In section 6 I test for Chamberlain’s *critique*. Section 7 concludes. See chapter 2 for more details concerning the wording of the survey question and the procedure I adopted to construct the variables used in the empirical analysis.

2 The estimation strategy

In this section I 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 for a rainy day” equation (Campbell, 1987). I also discuss identification and consider some extensions.

2.1 Income shocks decomposition

Suppose current income (in level)³⁴ admits the following canonical decomposition (as in Muth, 1960, and Blundell and Preston, 1997):

$$y_{it} = p_{it} + \varepsilon_{it} \tag{1}$$

where p_{it} is the permanent component of income and ε_{it} a transitory shock. For the sake of simplicity, I assume that the latter is i.i.d. with constant variance σ_ε^2 . The permanent component of income follows a random walk process of the

³⁴The measure of income that would be more appropriate to test the permanent income hypothesis is disposable family income net of asset income.

form:

$$p_{it} = p_{it-1} + \zeta_{it} \quad (2)$$

where ζ_{it} is the permanent shock; this is assumed to be i.i.d. with constant variance σ_u^2 .³⁵ I also assume that the transitory and the permanent shocks are orthogonal to each other at all leads and lags. Combining (1) and (2), one obtains:

$$\Delta y_{it} = \zeta_{it} + \Delta \varepsilon_{it} \quad (3)$$

where Δ is the first-difference operator.

2.2 The effect of transitory and permanent income shocks on savings

In chapter 1 I have shown that when the income process is characterized by equations (1)-(2), the change in consumption equals the sum of the permanent shock and the annuity value of the transitory shock (see equation 14, chapter 1). Similar implications can be derived for household savings. 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 interest rate, and the absence of credit market imperfections), one obtains the following saving “for a rainy day” equation:

$$s_{it} = - \sum_{\tau=1}^{\infty} \frac{1}{(1+r)^\tau} E(\Delta y_{it+\tau} | \Omega_{it}) \quad (4)$$

which implies that savings mirror the present discounted value of expected income declines. Using the income process (1)-(2), equation (4) simplifies to:

$$s_{it} = \frac{1}{1+r} \varepsilon_{it} \quad (5)$$

³⁵The assumption of constant variance for both the transitory and the permanent income shock can be removed without altering the estimation strategy.

The implications one derives from equation (5) are well known: permanent shocks do not matter (because under the conditions above the optimal rule is to consume them all), and only transitory income shocks explain saving. Certainty equivalence also implies that higher moments of the distribution of income shocks do not affect savings. Provided transitory and permanent income shocks were separately identifiable, one could implement the following regression:

$$s_{it} = \beta_0 + \beta_1 \text{transitory shock}_{it} + \beta_2 \text{permanent shock}_{it} + \text{error term}_{it} \quad (6)$$

where the error term reflects reporting error in saving, and test whether $\beta_1 \simeq 1$ and $\beta_2 = 0$. This is the main implication of the permanent income hypothesis I will test in the empirical analysis. Tests of the “saving for a rainy day” based on microeconomic data have been performed, among others, by Deaton (1992a) and Alessie and Lusardi (1997).³⁶

2.3 Identification

Let us define $E(x_{it}|\Omega_{it-1})$ the subjective expectation of x_{it} given the individual’s information set at time $t - 1$.³⁷ 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. Using (3) and the assumption of rational expectations, the transitory shock at time t can be *exactly* identified by:

$$\varepsilon_{it} = -E(\Delta y_{it+1}|\Omega_{it}) \quad (7)$$

Using (3) and (7), the permanent shock at time t is *exactly* identified by the expression:

$$\zeta_{it} = \Delta y_{it} - E(\Delta y_{it}|\Omega_{it-1}) + E(\Delta y_{it+1}|\Omega_{it}) \quad (8)$$

³⁶In both cases the authors had available short panels (2 years in Deaton and 3 years in Alessie and Lusardi). Thus their estimates are likely to suffer from the problem of inconsistency firstly remarked by Chamberlain (1984), even if aggregate shocks are controlled for through the use of time dummies (see chapter 2 for more details on this point).

³⁷Throughout, it is assumed that agents form rational expectations.

e.g., 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 $t + 1$. Thus, given the income process (1)-(2), the transitory and the permanent shock to income can always be identified provided 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.³⁸

The 1989 and 1991 SHIW data provide a unique opportunity to perform the tests of the permanent income hypothesis implied by equation (6). However, 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 (see section 3 for more details). More precisely, the SHIW data provide information on 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 seen that, given equations (7) and (8), the expressions:

$$-E(\Delta y_{it+1}|\Omega_{it}) = \varepsilon_{it} \text{ and } -E(\Delta y_{it}|\Omega_{it-2}) = \varepsilon_{it-2}$$

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

$$(y_{it} - y_{it-2}) - E(\Delta y_{it-1}|\Omega_{it-2}) + E(\Delta y_{it+1}|\Omega_{it}) = (\zeta_{it} + \zeta_{it-1})$$

³⁸Note that the identification of the income shocks carries over exactly as in (7) and (8) if income in levels includes fixed unobservable heterogeneity, i.e. if: $y_{it} = \lambda_i + p_{it} + \varepsilon_{it}$. This is because income shocks are identified from income changes. This also implies that $\varepsilon_{it} = -E(\Delta y_{it+\tau}|\Omega_{it})$ for all $\tau > 0$, as future transitory and permanent shocks have all conditional mean zero. The identification strategy is also robust to an income process of the form: $y_{it} = g_{it} + m_t + p_{it} + \varepsilon_{it}$, where g_{it} and m_t capture stochastic life-cycle and business-cycle effects, respectively (i.e., $E(g_{it+\tau}|\Omega_{it}) = E(m_{t+\tau}|\Omega_{it}) = 0$ for all $\tau > 0$). In this case, it is easy to show that: $s_{it} = \frac{1}{1+\tau}(g_{it} + m_t + \varepsilon_{it})$. The subjective expectation $-E(\Delta y_{it+1}|\Omega_{it})$ would now identify the composite error term ($g_{it} + m_t + \varepsilon_{it}$), with the permanent shock still identified by (8). A regression of s_{it} on the subjective expectation $-E(\Delta y_{it+1}|\Omega_{it})$ would still provide an estimate of the marginal propensity to save out of a transitory shock to income, the parameter I am interested in. Note also that deterministic life-cycle effects (a polynomial in age, for instance) can be easily accommodated by noting that $age_{it+\tau} = age_{it} + \tau$.

identifies the *sum* of the permanent shocks at time t and $t - 1$. Since under the null of the permanent income hypothesis savings depend only upon transitory innovations, that is all I need to implement the estimation of equation (6). The strategy I use to test for the null hypothesis of no effect of permanent shocks on savings is described below.³⁹

2.4 Consistency

The consistency of my saving equation estimates relies on a large cross-section dimension, rather than on a large time-series dimension, as is usually required. This is simply because I *do observe* the innovation in savings, e.g. I can condition on them. Indeed, under the null hypothesis of the PIH, the residual term of equation (6) is assumed to reflect only (additive) measurement error in saving. Hence, the consistency of my estimates rests only on the weak assumption that the cross-section dimension of the sample is large and that (additive) 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 makes practically irrelevant Chamberlain's *critique* (1984).⁴⁰

2.5 Testing for quadratic preferences

If I relax 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 their realizations; if this is the case, estimates will prove inconsistent. But this also suggests that one can test for the validity of the quadratic preferences assumption by augmenting the “saving for a rainy day” equation by higher moments of the distribution of income;

³⁹It turns out that if transitory shocks are serially correlated, subjective expectations no longer identify income shocks. This is my main identifying assumption.

⁴⁰Chamberlain's *critique* states that optimization errors average out over time but not necessarily across households in a cross-section.

under the null hypothesis of the permanent income hypothesis with quadratic preferences, higher moments should not explain saving.⁴¹

2.6 Chamberlain's *critique*

It is worth stressing at this point that the tests I perform are neither standard excess sensitivity tests (where consumption changes are regressed on expected income changes) nor orthogonality tests (where consumption changes are regressed on lagged income). See Browning and Lusardi, 1996, for a comprehensive survey of the empirical literature. Nonetheless, my data can be used to test for one of the more controversial assumptions made in the empirical literature, namely that the cross-section average of consumption innovations approaches zero when the cross-section dimension gets large; similarly, one can easily test for the supposed lack of correlation - here again, in the cross-section dimension- of consumption innovation with lagged instruments.⁴² Under the null of the permanent income hypothesis with quadratic preferences, infinite horizon and $r = \delta$, the consumption innovation equals: $\frac{r}{1+r}\varepsilon_{it} + u_{it}$ (see equation 14, chapter 1). The appropriate orthogonality condition for this problem is then: $E_{\tau} \left[\left(\frac{r}{1+r}\varepsilon_{it} + u_{it} \right) Z_{it-1} \right] = 0$, where the subscript “ τ ” attached to the symbol of expectation denotes the time-series dimension for each individual in the population and Z is an instrumental variable. In other words, the permanent income hypothesis with rational expectations implies:

$$p \lim_{T \rightarrow \infty} T^{-1} \sum_{t=1}^T \left[\left(\frac{r}{1+r}\varepsilon_{it} + u_{it} \right) Z_{it-1} \right] = 0 \text{ for all } i = 1, 2, \dots, N \quad (9)$$

Applied researchers who have not available a long time-series of individual data make use instead of the orthogonality condition: $E_n \left[\left(\frac{r}{1+r}\varepsilon_{it} + u_{it} \right) Z_{it-1} \right] =$

⁴¹Note that a regression of savings on the transitory shock and the conditional variance of income can be seen as a generalization of Caballero's model (1990) with CARA preferences and homoscedasticity. See chapter 1.

⁴²The lack of cross-section correlation between the consumption innovation and lagged instruments is invoked in tests of the excess sensitivity of consumption with respect to expected income changes. The latter is obtained as the projection of actual income changes on lagged instruments.

0, where the subscript “ n ” denotes the cross-section dimension. In other words, they assume that:

$$p \lim_{N \rightarrow \infty} N^{-1} \sum_{i=1}^N \left[\left(\frac{r}{1+r} \varepsilon_{it} + u_{it} \right) Z_{it-1} \right] = 0 \text{ for all } t = 1, 2, \dots, T \quad (10)$$

I will perform the test implied by equation (10) by fitting it to the timing of my data, i.e. I will consider the empirical validity of the orthogonality condition:

$$E_n \left[\left(\frac{r}{1+r} \varepsilon_{it} + (u_{it} + u_{it-1}) \right) Z_{it-2} \right] = 0 \quad (11)$$

which is the cross-section equivalent of the orthogonality condition:

$$E \left[\left(\frac{r}{1+r} \varepsilon_{it} + (u_{it} + u_{it-1}) \right) | \Omega_{it-2} \right] = 0 \quad (12)$$

with $Z_{it-2} \in \Omega_{it-2}$. The main advantage of this test is that under the null of the permanent income hypothesis with quadratic preferences I do observe the innovation in the change in consumption (as given by equation 14, chapter 1).⁴³ The main drawback of such test is the lack of generality: rejecting the null for a given time period does not imply a rejection for all time periods.

3 The data and the actual implementation of the test

I estimate various versions of the “saving for a rainy day” equation using the 1989-1991 panel section of the Bank of Italy Survey of Household Income and Wealth (SHIW). As shown in chapter 2, one of the main features of this data set is that it collected subjective information on future income in both 1989 and 1991.⁴⁴ The 1989 and 1991 SHIW have been used by Guiso, Jappelli and

⁴³As noted above, the nature of the SHIW data implies that - rather than observing $\frac{r}{1+r} \varepsilon_{it} + u_{it}$ (the true consumption innovation implied by the permanent income hypothesis)- I observe $\Phi_{it} = \frac{r}{1+r} \varepsilon_{it} + (u_{it} + u_{it-1})$. It follows that $E_{t-1}(\Phi_{it}) \neq 0$ but $E_{t-2}(\Phi_{it}) = 0$. Therefore, I can test for the supposed lack of cross-section correlation between Φ_{it} and the instruments dated $t - 2$. See also section 3 for more details about the data.

⁴⁴Subjective expectations are also asked in the 1995 SHIW, but differ quite radically from those I use in this paper.

Terlizzese (1992), and Lusardi (1997) to test various hypothesis related to the life-cycle permanent income hypothesis.

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 labour 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 (see chapter 2 for more details about the wording of the survey question and the construction of the variables of interest).

A problem with these data is that subjective expectations are not reported as for 1989 (1991), but in the following year, usually between March and September, although income, consumption and wealth data refer to the previous calendar year. The reason for that 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. I thus need to assume that people do not update their information set between the end of 1989 (1991) and the date of the interview,⁴⁵ 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; on the other hand, it is worth noting that in Italy labour contracts are renewed in the Autumn (usually between October and December).

As noted in chapter 2, SHIW respondents report one-year-ahead expectations referring to the rate of growth of their nominal earnings net of taxes and contributions and inflation expectations. Let $E_{it}z_{it+1}$ denote the expected growth rate of nominal earnings or pension income between t and $t + 1$, $E_{it}\pi_{it+1}$ the expected rate of inflation and $g_{it+1}^e = E_{it}z_{it+1} - E_{it}\pi_{it+1}$ the expected growth rate of real earnings (where $t = 1989$ or 1991). To obtain the one-year-ahead expectations of changes in earnings that would identify the transitory earnings shocks, I simply solve for the expected change in earnings. Given the assumptions on the timing of the expectations, the computation of the latter is simple. In fact, $E(\Delta y_{it+1}|\Omega_{it}) = y_{it}g_{it+1}^e$.

⁴⁵The same assumption has been made implicitly in all the papers quoted in this section.

Although each labour income recipient is asked to answer the survey question, I 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.⁴⁶

4 The empirical distribution of the income shocks

Tables 1 allows to examine the cross-section distribution of income shocks for the sample that includes heads or spouses (1,102 households).⁴⁷ For the sake of comparison, income shocks are divided by current earnings; hence, they can be interpreted in relative terms. Since I have only available the sum of permanent shocks in 1990 and 1991, the figures in the first column should be read as the ratio of average permanent shock between 1990 and 1991 and earnings in 1991. The next two columns focus on the relative transitory shocks in 1991 and 1989, respectively.

In 1991 average earnings featured a negative innovation of about 1.3 percent in real terms; the decomposition into transitory and permanent shocks, however, shows that while the permanent component plays a negative role (−4.5 percent on average), the transitory shock is positive (+3.2 percent on average).

Permanent shocks are negative for all population groups; however, the effect is stronger for the self-employed, the middle aged, the more educated, and the poor (as measured by family income quartiles). As for the transitory shocks in 1991, these are higher for those approaching the retirement, with few years of schooling, and living in the north. While still positive on average, the transitory earnings shock in 1989 is not as large as in 1991 (1.4 percent *vis-à-vis* 3.2 percent), and it is even negative for few population groups (the very young and the most educated).

An average permanent shock of −4.5 percent is not negligible. On the other

⁴⁶In other words, I regress saving on the head's earnings shocks, rather than on the shocks referring to disposable family income.

⁴⁷For 95 percent of our sample, I use information directly pertaining to the head of the household.

hand, there is an increasing body of evidence (Miniaci and Weber, 1996; Bertola and 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 changes, 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, pension and labour market reforms, etc.

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 labour market regulations. It has been argued that the former had the effect of increasing earnings inequality after decades of compression in the earnings differentials (Manacorda, 1997), while the latter had the effect of increasing earnings uncertainty because of job instability (Bertola and Ichino, 1996).

The income policy experiments were introduced as transitory measures aimed at freezing pay rise after years of unnecessary adjustments; ex post, some of these measures seem to have permanently reduced wages purchasing power (income policy agreements are actually still in force in 1999).

In my 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 pieces of legislation 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, labour market and pension reforms were accompanied by an increase in taxation. The population group that is likely to have suffered more from the introduction of new fiscal measures is the self-employed. While 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

changes in political attitudes towards tax non-compliance 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 I only observe 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 study only in 1989 and 1991.

5 Empirical results

Table 2 presents the results of estimating the “saving for a rainy day” equation for the sample of heads and spouses (1,102 households). I estimate three basic regressions: (i) the one implied by equation (5), with only the transitory income shock included as an explanatory variable, and then including: (ii) the permanent income shock, and (iii) the conditional variance of income. The latter can be easily derived from subjective expectations data. Note that the OLS regression 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)-(3) of table 2. I trim the sample at the bottom and top percentile of the distribution of saving to avoid my estimates being contaminated by influential outliers. Standard errors are robust to the presence of heteroscedasticity of unknown form. Saving is defined as the difference between family disposable income and non-durable consumption; I also experimented defining saving as the difference between family disposable income and total consumption and found no appreciable difference in the pattern of results.

The results of estimating equation (6) are supportive of the permanent income hypothesis with rational expectations. Column (1) shows that savings strongly

react to transitory income shocks (a point estimate of 0.71). The hypothesis that the propensity to save out of a transitory earnings shock is one has a p -value of 27 percent. The null hypothesis that the coefficient on the transitory shock equals $(1 + r)^{-1}$ can also be tested by considering a grid of possible values for the real interest rate ranging from 0 to 10 percent;⁴⁸ in no case did I reject the null hypothesis.

In column (2) I add to the main specification in levels the sum of the permanent income shocks in periods t and $t - 1$. This raises the problem that I cannot separately identify the effect of the shocks at two different dates. The problem can be handled by noting that under the null hypothesis of the permanent income hypothesis with quadratic preferences, both coefficients are zero, and so should be the coefficient attached to the sum of the current and past permanent income shocks. The results show that the null hypothesis is to be rejected: permanent income shocks are significant predictors of household savings. Nonetheless, the size of the coefficient is tiny: the null hypothesis that the transitory shock and the permanent shock equally affect savings is strongly rejected (a p -value of 0.32 percent). Taken at face value, these results suggest that households save not only the transitory income shocks in their entirety, but also a sizeable portion of their permanent shocks. Therefore, the certainty equivalence model seems to fail in the sense of predicting saving rates that are too low (in absolute value) with respect to the available evidence. In passing, these results are also against the full consumption insurance hypothesis, for which idiosyncratic (transitory or permanent) shocks should have no effect whatsoever on intertemporal consumption choices. Note that my test of full consumption insurance is in strict agreement with the theory, for I observe directly idiosyncratic shocks rather than a proxy of them.

Yet, too much savings can be reconciled with the existence of a precautionary motive for saving. A piece of evidence strongly in support of the latter is re-

⁴⁸The average real interest rates in 1991 were: 0.58 percent (deposits), 5.58 percent (Treasury bonds), and 4.32 percent (other assets, including shares). Interest rates in 1989 were very similar to those for 1991.

ported in column (4); here I include the conditional variance of head's earnings⁴⁹ alongside the transitory and the permanent income shock. The version of the permanent income hypothesis I have tested so far might fail because preferences are not quadratic. If individual utility admits a positive third derivative (e.g., if consumers are prudent in the sense clarified by Kimball, 1990), 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 realizations. The test I 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 rejected: the conditional variance of earnings has the expected sign (more uncertainty should in fact increase 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 from Italy (Guiso, Jappelli and Terlizzese, 1992; chapter 2 of this dissertation).

To confirm the robustness of my findings, I have re-estimated the “saving for a rainy day” equation by accounting for preference heterogeneity. In particular, I assume that the bliss point of household utility is a function of age, age squared and family size.⁵⁰ Results are presented in table 3, columns (1) to (3). As is clear, the results are not much affected by the introduction of bliss point heterogeneity.

I have also experimented by excluding the elderly (those aged more than 65, the standard retirement age for males) and the self-employed. 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 the retirement; the reason to exclude the self-employed is

⁴⁹This is defined as $var(\Delta y_{it+1}|\Omega_{it}) = var(y_{it+1}|\Omega_{it})$. Note that I cannot distinguish between the variance of the transitory shock and the variance of the permanent shock. The variance term is obtained from the conditional variance of the rate of growth of future earnings by noting that: $var\left(\frac{\Delta y_{it+1}}{y_{it}}\right)|\Omega_{it} = y_{it}^{-2} var(\Delta y_{it+1}|\Omega_{it})$. See chapter 2 for more details on the construction of the subjective variance.

⁵⁰The inclusion of additional variables (education, region dummies, etc.) does not affect the results.

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 labour income from asset income. The results obtained from excluding these two groups are presented in table 4. It is worth noting that the magnitude of the coefficients is not much affected by such exclusions; on the other hand, the precision of the estimates suffers from keeping out either population group.

Heterogeneity in savings may also arise from unobservable individual effects in the income process. Suppose to rewrite equation (3) as:

$$\Delta y_{it} = \zeta_{it} + \Delta \varepsilon_{it} + \phi_i \quad (13)$$

where ϕ_i is an idiosyncratic deterministic trend in the income process.⁵¹ Given (13), the “saving for a rainy day” equation rewrites:

$$s_{it} = \frac{1}{1+r} \varepsilon_{it} + \lambda_i + v_{it} \quad (14)$$

where $\lambda_i = -r^{-1}\phi_i$, and v_{it} is a measurement error in savings.

The presence of unobservable fixed heterogeneity in the saving function might invalidate simple least squares estimates, in particular if $E(\varepsilon_{it}\phi_i) \neq 0$. However, first-differences estimates are still consistent because the individual effect is washed out when transforming the data.⁵²

In columns (4) and (5) of table 2 I present the results of estimating equation (14) by first-differences. Accounting for heterogeneity in the income process leads has a dramatic effect. As shown in column (4), the saving equation displays point estimates that are much lesser than those in column (1). The effect of the transitory shock is now only 0.45 (with a robust standard error of 0.22). In columns (5), I present the results of controlling for both the unobservable

⁵¹See Blundell and Preston (1997) for details on the assumptions generating the income process (13).

⁵²Note that, from income process (13), subjective expectations no longer identify exactly the transitory income shock, because: $E(\Delta y_{it+1}|\Omega_{it}) = -\varepsilon_{it} + \phi_i$. However, in order to estimate (14) in first-differences, I only need to identify the change in the transitory income shock, which is provided by first-differencing the subjective expectations, i.e. (taking into account the timing of the SHIW data), from: $E(\Delta y_{it+1}|\Omega_{it}) - E(\Delta y_{it-1}|\Omega_{it-2}) = -(\varepsilon_{it} - \varepsilon_{it-2})$.

heterogeneity and the conditional variance of earnings. Here again, the results are against the certainty equivalence model, confirming the practical importance of the heterogeneity in the income process and the precautionary motive for saving. The specification tested in columns (4) deliver point estimates that are not literally consistent with the permanent income hypothesis. In particular, the permanent income hypothesis implied by equation (5) is consistent only with interest rates of about 30 percent in real terms. A possible explanation is that agents discount the future at rates higher than the ones prevailing in the credit market as a device to reduce the uncertainty related to future resources.⁵³

However, a modest effect of transitory shocks on saving may hide binding liquidity constraints. Flavin (1993) proposes the consumption function to be written as:

$$c_{it} = y_{it}^P + \gamma (y_{it} + r a_{it} - y_{it}^P) \quad (15)$$

where $y_{it}^P = \frac{r}{1+r} \left(a_{it} + \sum_{\tau=0}^{\infty} \frac{E(y_{it+\tau} | \Omega_{it})}{(1+r)^\tau} \right)$.⁵⁴ According to Deaton (1992a), the parameter γ represents “the extent to which consumption responds to current income over and above the amount that is warranted by the permanent income hypothesis”. Thus, the finding that $\gamma > 0$ can be interpreted as a symptom of liquidity constraints, as liquidity constrained households can increase their consumption only when income is directly available. It is easy to show that, using (15), the “saving for a rainy day” equation rewrites:

$$s_{it} = -(1 - \gamma) \sum_{\tau=1}^{\infty} \frac{1}{(1+r)^\tau} E(\Delta y_{it+\tau} | \Omega_{it}) \quad (16)$$

and then, using again the income process (13), that:

$$s_{it} = \frac{1 - \gamma}{1 + r} \varepsilon_{it} + \lambda_i$$

⁵³Similar results are obtained when preference heterogeneity is accounted for (see columns (4) and (5) of table 3), although estimates are slightly less precisely measured.

⁵⁴This is the standard definition of “permanent income” (see Deaton, 1992b, and chapter 1). This should not be confounded with the permanent component of income defined in equation (2)

where $\lambda_i = -(1 - \gamma)r^{-1}\phi_i$. A positive γ is consistent with the estimates in first-differences reported in columns (4). With a real interest rate of 5 percent, say, γ is roughly 0.5, implying that liquidity constraints are likely to be playing an important role.⁵⁵ Given the severe imperfections of the Italian credit markets by the standards of other industrialised countries and the pervasiveness of various liquidity constraints, particularly in the mortgage market (Guiso, Jappelli and Terlizzese, 1994), this result does not come as a surprise. It is worth pointing out that this evidence is not in contrast with that reported in chapter 2, where I could not find excess sensitivity. The reason is that it can be hard to detect the presence of liquidity constraints from Euler equation estimates, essentially because of the lack of power of the test; in contrast, the test presented in this chapter is more robust because it uses information that are very rarely available to the econometrician.

5.1 Testing for Chamberlain's *critique*

I conclude this section by presenting the results of testing for the lack of cross-section correlation between the innovation in consumption growth and lagged instruments. The motivation for this test are reported in section 2.6. I focus on a constant term, and lagged consumption, disposable family income and head's (or spouse's) earnings, as these are the instruments usually considered in the empirical literature. Note that the test is robust to the presence of measurement errors in the subjective expectation variables.

Consumption innovations do not average out in the cross-section; the null hypothesis that the cross-section average of consumption innovations is zero (which is a test for the lack of cross-section correlation between consumption innovations and a constant term) is strongly rejected (the standard normal test statistics displaying a value of -3.5). Table 5 reports the pairwise coefficients of correlation between the consumption growth innovation and the three variables I focus on.

⁵⁵The magnitude of this coefficient is very similar to that obtained when testing for excess sensitivity on time series data. Recall that under the null of the permanent income hypothesis with quadratic preferences, estimates based on time series data are consistent (because of the large T argument) and do not suffer from aggregation bias.

The null hypothesis is strongly rejected. This confirms the practical importance of Chamberlain's *critique* (1984) and casts some serious doubts on the consistency of estimates derived in almost the whole empirical literature (with the not surprising exception of Hayashi, 1985, who uses subjective expectations of consumption and income to test the permanent income hypothesis).

6 Conclusions

This study has presented simple tests of the permanent income hypothesis with quadratic preferences and infinite horizon. I have shown that the availability of subjective income expectations 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 for a rainy day", namely that saving reacts only to transitory shocks.

According to the empirical analysis, this version of the permanent income hypothesis should be rejected. I have shown that when heterogeneity is accounted for, savings do react to transitory income shocks, but the magnitude of the effect is tiny, a finding that can be reconciled only with very high real interest rates or binding liquidity constraints. In addition, I have shown that the assumption of quadratic preferences is inappropriate: higher moments of the distribution of earnings should not matter whereas 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.

There are various theoretical reasons why the permanent income with quadratic preferences may fail to hold. A good account of the literature is given in Deaton (1992b), Browning and Lusardi (1996) and Attanasio (1999). In addition to theoretical explanations, I cannot rule out the possibility that the permanent income hypothesis is true, but my data are inappropriate to test it. For instance, measurement error in the independent variables might invalidate my estimates.

It is known that measurement errors in the independent variable tend to bias the OLS estimate towards zero. This could explain the very low coefficient attached to the transitory shock in the saving equation in first-differences, a data transformation strategy that exacerbates the downward bias of OLS estimates. In ordinary circumstances, a variable measured with error should be instrumented in order to eliminate the bias. But the variable I am dealing with is the transitory income innovations, e.g., a variable that - by definition of innovation process - is assumed to be unpredictable on the basis of past information.⁵⁶

⁵⁶To some extent, subjective expectations could be free from measurement error problems if they partly reflect the ignorance that respondents have about their future. On the other hand, it has been argued that subjective expectations are more prone than, say, past earnings, to measurement error problems. Yet, in the absence of administrative data both are elicited measures.

Table 1
The distribution of income shocks

	<u>Average permanent shock</u> <i>Earnings in 1991</i>	<u>Transitory shock in 1991</u> <i>Earnings in 1991</i>	<u>Transitory shock in 1989</u> <i>Earnings in 1989</i>
<i>Age in 1991</i>			
≤ 35	-0.0308 (0.0154)	0.0251 (0.0046)	-0.0009 (0.0044)
35-55	-0.0537 (0.0110)	0.0275 (0.0034)	0.0073 (0.0021)
>55	-0.0392 (0.0098)	0.0404 (0.0041)	0.0282 (0.0031)
<i>Education</i>			
Compulsory	-0.0374 (0.0093)	0.0345 (0.0028)	0.0218 (0.0021)
High school	-0.0576 (0.0114)	0.0289 (0.0048)	0.0060 (0.0029)
University	-0.0496 (0.0194)	0.0257 (0.0077)	-0.0080 (0.0063)
<i>Region</i>			
North	-0.0515 (0.0102)	0.0364 (0.0042)	0.0126 (0.0028)
South	-0.0538 (0.0115)	0.0261 (0.0027)	0.0166 (0.0025)
<i>Occupation</i>			
Employed	-0.0288 (0.0067)	0.0297 (0.0026)	0.0092 (0.0022)
Self-employed	-0.0865 (0.0237)	0.0304 (0.0076)	0.0085 (0.0053)
<i>Sector</i>			
Private	-0.0481 (0.0093)	0.0299 (0.0029)	0.0097 (0.0028)
Public	-0.0439 (0.0100)	0.0355 (0.0041)	0.0080 (0.0030)
<i>Family income</i>			
1 st quartile	-0.0769 (0.0163)	0.0312 (0.0047)	0.0263 (0.0032)
2 nd quartile	-0.0443 (0.0131)	0.0313 (0.0035)	0.0155 (0.0040)
3 rd quartile	-0.0311 (0.0113)	0.0317 (0.0055)	0.0062 (0.0027)
4 th quartile	-0.0266 (0.0131)	0.0335 (0.0047)	0.0070 (0.0034)
<i>Whole sample</i>	-0.0447 (0.0068)	0.0319 (0.0023)	0.0138 (0.0017)

Table 2
The “saving for a rainy day” equation⁵⁷

	(1)	(2)	(3)	(4)	(5)
Transitory shock _t	0.7105 (0.2632)	1.2782 (0.3038)	1.2823 (0.3037)	0.4530 (0.2179)	0.4774 (0.2153)
Permanent shock _t		0.1528 (0.0763)	0.1479 (0.0758)		
Conditional variance _t			0.0010 (0.0004)		0.0007 (0.0004)
# of observations	2,204	1,102	1,102	1,102	1,102
R^2	0.0085	0.0414	0.0474	0.0070	0.0117

⁵⁷The sample includes 1,044 heads (94.74 percent of the sample) and 58 (5.26 percent) spouses. Standard errors robust to heteroscedasticity of unknown form are reported in parenthesis. All regressions, except the ones in columns (4) and (5), include a constant.

Table 3**The “saving for a rainy day”: accounting for preference heterogeneity⁵⁸**

	(1)	(2)	(3)	(4)	(5)
Transitory shock _t	0.6673 (0.2624)	1.2531 (0.2998)	1.2566 (0.2995)	0.3918 (0.2258)	0.4147 (0.2234)
Permanent shock _t		0.1643 (0.0750)	0.1591 (0.0745)		
Conditional variance _t			0.0011 (0.0004)		0.0006 (0.0004)
Age	452.47 (109.27)	520.27 (146.35)	504.64 (145.90)	11.22 (23.03)	10.37 (23.07)
Age ²	-4.19 (1.05)	-4.71 (1.37)	-4.51 (1.37)		
Family size	767.10 (218.46)	1130.50 (262.02)	1159.41 (260.04)	2302.57 (895.13)	2223.11 (898.83)
# of observations	2,204	1,102	1,102	1,102	1,102
R ²	0.0261	0.0726	0.0788	0.0138	0.0177

⁵⁸The sample includes 1,044 heads (94.74 percent of the sample) and 58 (5.26 percent) spouses. Standard errors robust to heteroscedasticity of unknown form are reported in parenthesis. All regressions, except the ones in columns (4) and (5), include a constant.

Table 4

The “saving for a rainy day” equation : sensitivity analysis⁵⁹

	Excluding those aged over 65			Excluding the self-employed		
	(1)	(2)	(3)	(1)	(2)	(3)
Transitory shock _t	0.6197 (0.2868)	1.1565 (0.3230)	0.6598 (0.2797)	0.6127 (0.2747)	1.0677 (0.3688)	0.6261 (0.2746)
Permanent shock _t		0.1694 (0.0832)			0.1945 (0.0724)	
Conditional variance _t			0.0011 (0.0005)			0.0007 (0.0004)
# of observations	1,834	899	1,834	1,800	903	1,800
R ²	0.0069	0.0417	0.0136	0.0058	0.0272	0.0087

⁵⁹Standard errors robust to heteroscedasticity of unknown form are reported in brackets. All regressions include a constant.

Table 5
Testing for Chamberlain's critique⁶⁰

Real interest rate	Disposable family income	Earnings	Consumption
	(1)	(2)	(3)
0.05	-0.3298 [0.0000]	-0.5519 [0.0000]	-0.2276 [0.0000]

⁶⁰This table reports the pairwise coefficients of correlation between the consumption innovation and the variables in columns (1), (2) and (3) dated $t - 2$. The p -value for the null hypothesis of no correlation is reported in square brackets under the correlation coefficient. The consumption innovation is calculated under the null hypothesis of the permanent income hypothesis with quadratic preferences as: $\frac{r}{1+r}\varepsilon_{it} + (u_{it} + u_{it-1})$.

4

Consumption insurance or consumption mobility?

1 Introduction

A large body of literature in industrialized and developing countries alike has proposed tests of full or partial consumption insurance (Cochrane, 1991; Townsend, 1994). The main implication of full consumption insurance is that the cross-sectional distribution of consumption over any group of households is constant over time. Therefore under complete markets, consumption growth is uncorrelated with changes in individual endowments. Of course aggregate consumption can increase or decrease, so that consumption growth for any household can be positive or negative, but the relative position of each individual in the cross-sectional distribution is preserved both in the short and the long run. Consumption insurance thus implies strong predictions about the entire consumption distribution, not just its mean or variance.⁶¹

⁶¹Deaton and Paxson (1994) show that the certainty equivalence version of the permanent income hypothesis implies that the cross-sectional dispersion in consumption of any given cohort should increase over time. They also note that full consumption insurance implies that the cross-sectional variance of consumption of the same cohorts should be constant over time.

In particular, the theory implies the total absence of consumption mobility between any two time periods, a much stronger proposition than is usually addressed by tests of consumption insurance. It follows that if one observes individuals moving up and down in the consumption distribution one must conclude that some people are not insulated from idiosyncratic shocks, which contradicts the assumptions of full consumption insurance. Although this implication of consumption insurance was mentioned in a theoretical paper by Banerjee and Newman (1991), to our knowledge it has never been explored in empirical analysis.

To test for the invariance of the consumption distribution one needs panel data. I construct a transition matrix for the distribution and apply non-parametric statistical tools to test the hypothesis of absence of consumption mobility between time periods. The empirical analysis is conducted on a panel of households drawn from the Bank of Italy Survey of Household Income and Wealth for the years 1987 to 1995.

With respect to previous studies that found overwhelming evidence against full consumption insurance (Cochrane, 1991; Attanasio and Davis, 1996) my contribution relates to both method and substance. On the methodological side, I analyse the transition matrix for household consumption and can therefore characterize the entire distribution of consumption rather than just its mean. Since I use a non-parametric index of market completeness, the statistical procedure is not sensitive to the particular utility function used, e.g. relative or absolute risk aversion.

On substance, examining the entire consumption distribution avoids arbitrary identifying assumptions. In fact, the statistical tests of consumption insurance used so far are tightly parametrised. To test the prediction that idiosyncratic shocks are uncorrelated with consumption growth, they rely on univariate regressions of consumption growth on aggregate variables and idiosyncratic shocks (such as change in household resources, unemployment hours, days of illness, etc.). Finding appropriate and exogenous proxies for the shocks is difficult in the extreme. My procedure has several advantages: (i) I need not rely on any

parametrised form for the utility function, (ii) I need not identify any of these shocks; (iii) I need not assume that they are uncorrelated with unobservable or omitted preference shocks, including household fixed effects. Furthermore, the statistical test naturally provides an index of market completeness that measures the deviation of the actual consumption distribution from the distribution predicted by complete markets. This index can be used to compare the evolution of market completeness over time and check whether different population groups experience different degrees of consumption mobility.⁶² Such information can be important for policy purposes. Consider for instance the possibility of a switch from a less to a more redistributive tax system and recall that tax progressivity provides implicit insurance to consumers. The effect of such policy change depends upon the amount of risk sharing already available in the economy. If private insurance markets are absent or largely incomplete, the policy change I am examining generates a welfare gain because it provides consumers with additional insurance; however, if consumers can fully insure the idiosyncratic shocks they face through private insurance markets, the policy change plays no role. As argued by Bogarde and Perri (1999), such a policy may even turn into a welfare loss if the provision of public insurance through progressive taxes crowds out private insurance schemes.

In Section 2 I review the model of consumption insurance and set out the basic intuition underlying my procedure. Section 3 presents the non-parametric test of consumption insurance and the mobility index. In Section 4 I explore the robustness of the test with respect to preference specification of the utility function and measurement error in consumption. The data and the empirical results are presented in Sections 5 and 6, respectively. I strongly reject full consumption insurance, in both the short and the long run and for each sample group that I analyse. The rejection of consumption insurance is not due to preference specification or measurement error. Section 7 concludes.

⁶²Hayashi, Altonji and Kotlikoff (1996) recommended that “future research should be directed to estimating the extent of consumption insurance over and above self-insurance” (p. 290). This paper is a step in this direction.

2 Consumption insurance

I have already reviewed the model's main insight in chapter 1; this is repeated here for convenience. The argument I develop in this chapter does not rest on the specific form of the utility function; however, as a matter of convenience, I proceed on the assumption that households have identical preferences of the *CRRA* type, $u(c) = (1 - \gamma)^{-1}c^{1-\gamma}$. If the social planner maximizes a weighted sum of individual households' utilities, the Lagrangian of the problem can be written (Deaton, 1997):

$$L = \sum_h \lambda_h \sum_s \sum_t \pi_{s,t} u(c_{h,s,t}) + \sum_s \sum_t \mu_{s,t} \left(C_{s,t} - \sum_h c_{h,s,t} \right)$$

where h , s and t are subscripts for the household h in the state of nature s in period t , λ_h is the social weight for household h , $\mu_{s,t}$ is the Lagrange multiplier associated with the resource constraint, $\pi_{s,t}$ the probability of state s in time period t , and $C_{s,t}$ aggregate consumption in state s and time t .

The first order condition can be written in logarithms as:

$$-\gamma \ln c_{h,s,t} = \ln \mu_{s,t} - \ln \lambda_h - \ln \pi_{s,t} \quad (1)$$

To obtain the rate of growth of consumption, one subtracts side-by-side from the expression at time $t + 1$:

$$\Delta \ln c_{h,t+1} = -\gamma^{-1} \Delta \ln \mu_{t+1} + \gamma^{-1} \Delta \ln \pi_{t+1} \quad (2)$$

where I drop the subscript s because only one state is realized in each period. The two terms on the right-hand-side of equation (2) represent aggregate effects. The first is the growth rate of the Lagrange multiplier, the second is the growth rate of the state probabilities. Note that first-differencing has eliminated all household fixed effects.

Equation (2) states that the rate of growth of consumption of each household is the same. This implies that the initial cross-sectional distribution of consumption levels is a sufficient statistic to describe all future distributions:

since all households have the same rate of growth of consumption, their relative position is stationary. Note that the stationarity of the cross-sectional distribution is directly implied by the assumption that insurance markets fully insulate households from idiosyncratic shocks.

The statistical counterpart of consumption insurance is that the transition matrix for household consumption is an identity matrix. In the next section I show how to construct such a transition matrix and how the matrix can be summarized by an appropriately designed mobility index. This index can be used to test the null hypothesis of no consumption mobility, which is implied by the theory of consumption insurance.

3 Consumption mobility

In order to summarize the transition matrix for consumption through an appropriate index of mobility, I build on an approach proposed by Shorrocks (1978). Assume that \mathbf{P} is an unobservable $q \times q$ stochastic transition matrix of household consumption; q is the number of quantiles that summarize the distribution. For notational simplicity let's consider transition probabilities from period t to period $t + 1$; it is then straightforward to extend the argument to transition probabilities in periods $t + 2$, $t + 3$, and so on. The generic element of the \mathbf{P} matrix is p_{ij} , the probability of moving from quantile i in period t to quantile j in period $t + 1$. Define n_{ij} as the number of households that move from quantile i in period t to quantile j in period $t + 1$ and $n_i = \sum_j n_{ij}$ as the total number of observations in each row i of the \mathbf{P} matrix. The maximum likelihood estimator of the first-order Markov transition probabilities is $\hat{p}_{ij} = \frac{n_{ij}}{n_i}$. The Shorrocks index of mobility is then defined as:⁶³

$$S(\mathbf{P}) = \frac{q - \text{trace}(\mathbf{P})}{q} \quad (3)$$

If the probability of being in quantile i in period t is independent from the

⁶³In its original formulation, the index is divided by $(q - 1)$ rather than by q . We use this slight modification to bound the index between 0 and 1.

probability of being in quantile j in period $t+1$, the typical entry of the transition matrix is $p_{ij} = q^{-1}$ for all i and j . It follows that $\text{trace}(\mathbf{P}) = 1$ and $S(\mathbf{P}) = (q - 1)/q$. Under consumption insurance the probability of being in quantile i in period t equals the probability of being in quantile i in period $t + 1$ and the probability of moving to a different quantile is zero. In this case the transition matrix is an identity matrix:

$$p_{ij} = \begin{cases} 0 & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases}$$

so that $\text{trace}(\mathbf{P}) = q$ and the index reaches its lower bound, $S(\mathbf{P}) = 0$. Since $0 \leq \text{trace}(\mathbf{P}) \leq q$, the mobility index satisfies the inequalities $0 \leq S(\mathbf{P}) \leq 1$.⁶⁴ $S(\mathbf{P})$ can be interpreted as the proportion of households moving across the consumption distribution between t and $t + 1$.

The central limit theorem implies that $\text{trace}(\hat{\mathbf{P}}) \overset{a}{\sim} N\left(\sum_i p_{ii}; \sum_i \frac{p_{ii}(1-p_{ii})}{n_i}\right)$ so that $S(\hat{\mathbf{P}})$, the maximum likelihood estimator of $S(\mathbf{P})$, is asymptotically normally distributed (Schluter, 1998):

$$S(\hat{\mathbf{P}}) \overset{a}{\sim} N\left(\frac{q - \sum_i p_{ii}}{q}; \frac{1}{q^2} \sum_i \frac{p_{ii}(1-p_{ii})}{n_i}\right)$$

One can therefore test the null hypothesis of full consumption insurance, $S(\mathbf{P}) = 0$, using the statistic:

$$Z_1 = \frac{\frac{q - \sum_i \hat{p}_{ii}}{q}}{\sqrt{\frac{1}{q^2} \sum_i \frac{\hat{p}_{ii}(1-\hat{p}_{ii})}{n_i}}} \sim N(0, 1) \quad (4)$$

The test is simple and powerful: the data requirement are minimal, because the test requires knowledge only of the consumption distribution, and there is no need to identify exogenous idiosyncratic shocks to test full consumption insurance. An important feature of the test is that it does not rely on any specific form for the utility function. As the ordering of household consumption is invariant to monotonic transformation of the utility function, so are quantile probabilities.

⁶⁴The upper bound is a case in which all households move to a different quantile so that $\text{trace}(\mathbf{P}) = 0$ and $S(\mathbf{P}) = 1$.

It is often claimed that some population groups are more insulated than others from idiosyncratic shocks, or that households are more protected by idiosyncratic shocks in some periods than in others. To evaluate if consumption mobility differs statistically over time or between population groups one can construct a test of no differential mobility between two groups or time periods, based on the statistic:

$$Z_2 = \frac{S(\widehat{\mathbf{P}}_d) - S(\widehat{\mathbf{P}}_k)}{\sqrt{s.e.(S(\widehat{\mathbf{P}}_d))^2 + s.e.(S(\widehat{\mathbf{P}}_k))^2}} \sim N(0, 1) \quad (5)$$

where d and k are appropriately defined to allow comparisons over time or between population groups. Under the null hypothesis of no differential mobility, the statistic (5) is also asymptotically distributed as a standard normal.

4 Extensions

The mobility index is derived assuming that the utility function is the same for all households and that there is no measurement error in consumption. In practice the index could potentially be upward biased by idiosyncratic preference shifts, preference heterogeneity or reporting errors. Supposing that demographic variables, household composition and labor supply affect marginal utility and not just consumption, the latter might rise or fall as these variables change over time. Part of the change in the consumption distribution as measured by the mobility index may therefore reflect genuine choices by households rather than uninsurable shocks. Likewise, consumption trajectories may differ because people have different preference parameters.

Measurement errors too can produce apparent consumption mobility. If households report their consumption with errors, one will find units moving up and down even with consumption insurance; hence, the index will tend to report higher mobility. I address these two problems in turn.

4.1 Preference specification and heterogeneity

Equation (1) suggests that the ratio between the marginal utilities of house-

holds h and h' is stationary. This does not always imply that the ratio of consumption levels is stationary, nor that the growth rate of consumption is the same for all households. Consider a case where the isoelastic utility function is augmented by a multiplicative preference shift θ :

$$u(c, \theta) = \theta \frac{c^{1-\gamma}}{1-\gamma}$$

It can be immediately shown that the growth rate of consumption for household h can be written as:

$$\Delta \ln c_{h,t+1} = -\gamma^{-1} \Delta \ln \mu_{t+1} + \gamma^{-1} \Delta \ln \pi_{t+1} + \gamma^{-1} \Delta \ln \theta_{h,t+1} \quad (6)$$

Equation (6) states that, over and above the effect of aggregate components, part of the cross-sectional movement in consumption growth is due to household-specific preference shifts (with the arrival of children, changes in household composition, age, and so on). If θ changes over time, the consumption distribution will no longer be stationary and the mobility index will be greater than zero even under consumption insurance. In the empirical section I therefore check for the robustness of the mobility index using *per capita* consumption and consumption *per adult equivalent*; I also experiment with a measure of consumption adjusted by a larger set of preference shifts.⁶⁵

A related problem is the possibility that consumption and leisure may not be separable.⁶⁶ Although the implications of consumption insurance are unaffected when consumption and leisure are not separable, the right-hand-side of equation (2) includes another term, the rate of growth of the Lagrange multiplier of aggregate leisure. Cochrane (1991) points out that this term will vary between individuals except under the highly unrealistic assumption that the planner can freely transfer leisure across households. If the assumption is discarded, standard

⁶⁵In conventional tests of consumption insurance, preference shifts pose a rather different problem. If idiosyncratic shocks are correlated with omitted preference shifts, the estimated coefficients of the shock variables are biased, the direction of the bias depending on the correlation between preferences and shocks.

⁶⁶I do not focus on non-separabilities between different consumption goods because in the empirical application I use a measure of total non-durable consumption.

tests and my own procedure produce spurious evidence against consumption insurance. To address this problem, in the empirical section I augment the vector of preference shifts with the household head's leisure.

Note that my test is asymmetrically robust. The absence of consumption mobility (a result that does not reject consumption insurance) must imply that the preference shifts for which I do not control are not important determinants of the growth of marginal utility. In other words, the lack of consumption mobility cannot reflect estimator bias, as in more standard tests of the theory. Moreover, my test is robust in circumstances in which standard tests are not. The latter rely on regressions of consumption growth on idiosyncratic shocks and reject consumption insurance when the coefficients of the shock variables are significantly different from zero. But if the shocks are affected by measurement error, the OLS estimates are biased towards zero, providing spurious evidence in favour of consumption insurance. In contrast, my index will still report mobility because it does not require identifying idiosyncratic shocks in the first place.

An alternative way of introducing heterogeneity is to assume that the (unobservable) parameters of the utility function, say the degree of relative risk aversion, vary across individual:

$$u(c) = \frac{c^{1-\gamma_h}}{1-\gamma_h}$$

implying that the growth rate of consumption for household h is:

$$\Delta \ln c_{h,t+1} = -\gamma_h^{-1} (\Delta \ln \mu_{t+1} - \Delta \ln \pi_{t+1}) = -\gamma_h^{-1} \cdot \kappa_{t+1}$$

Substituting in the expression above the growth rate of consumption in period t :

$$\Delta \ln c_{h,t+1} = \frac{\kappa_{t+1}}{\kappa_t} \cdot \Delta \ln c_{h,t}$$

Even if individual growth rates may be different, the period t ordering of growth rates will be identical in period $t + 1$. While preference homogeneity implies that the initial cross-sectional distribution of consumption levels is a sufficient statistic for all future distributions, preference heterogeneity implies

that the initial cross-sectional distribution of consumption growth rates is a sufficient statistic for all future distribution of growth rates implying that one should not observe mobility in the transition matrix for consumption growth between period t and $t + 1$.

4.2 Measurement error

In the absence of preference shifts, consumption insurance delivers the following transition rule for *true* log-consumption:

$$\ln c_{h,t+1} = m_{t+1} + \ln c_{h,t} \quad (7)$$

where $m_{t+1} = -\gamma^{-1} (\Delta \ln \mu_{t+1} - \Delta \ln \pi_{t+1})$. Now suppose that consumption is measured with a multiplicative error:

$$\ln c_{h,t+1}^* = \ln c_{h,t+1} + v_{h,t+1} \quad (8)$$

$$\ln c_{h,t}^* = \ln c_{h,t} + v_{h,t} \quad (9)$$

where $\ln c^*$ is *measured* consumption and v is a classical measurement error satisfying the assumption $v \sim i.i.n.d. (0, \sigma_v^2)$ (for simplicity, I also assume that its distribution is stationary). The transition law for log-consumption can be rewritten as:

$$\ln c_{h,t+1}^* = m_{t+1} + \ln c_{h,t}^* + v_{h,t+1} - v_{h,t}, \quad (10)$$

which implies that individual consumption growth is no longer a constant but an $MA(1)$ process with a time-varying drift, m_{t+1} . Measurement error therefore biases the mobility index $S(\hat{\mathbf{P}})$ upwards: rejecting the null hypothesis $S(\mathbf{P}) = 0$ no longer implies that consumption insurance is violated.

The bias can be handled by noting that measurement error effectively increases the lower bound of the “true” mobility index $S(\mathbf{P})$. To see why, note first that regardless of consumption insurance the cross-sectional mean of $\ln c^*$ equals that of $\ln c$ because measurement errors average out. Note also that the

difference between $var(\ln c^*)$ and $var(\ln c)$ depends on the variance of the measurement error. Since $\ln c^* = \ln c + v$, it follows that $var(\ln c^*) = var(\ln c) + \sigma_v^2$, or $\sigma_v^2 = \left[1 - \frac{var(\ln c)}{var(\ln c^*)}\right] var(\ln c^*) = \alpha \cdot var(\ln c^*)$.

The parameter α indicates the fraction of the cross-sectional variance of measured consumption that is contaminated by measurement error, ranging from 0 in absence of measurement error to 1 when the variance of measured consumption is entirely explained by measurement error. To get a feeling for how measurement error affects the statistical test, I use the variance-covariance matrix of consumption growth to estimate realistic values for α . I then perform a Montecarlo simulation under the null hypothesis of consumption insurance and measurement error. For each value of α I show how to generate different lower bounds of mobility, and then compare the actual mobility index with the theoretical index obtained under the joint hypothesis of consumption insurance and measurement error.

5 The data

The statistical test requires panel data on consumption. I use the 1987-1995 panel of the Italian Survey of Household Income and Wealth (SHIW). As explained in chapter 1, the data set contains measures of consumption, income, and demographic characteristics of households. The SHIW provides a measure of total non-durable consumption, not just food, thus overcoming one of the main limitations of other panels, such as the PSID, that have been used to test for consumption insurance. See chapter 1 for more details about the survey and definition of the variables.

To minimize measurement error I exclude cases in which the head changes over the sample period or gives inconsistent age figures. The total number of transitions is 10,508. After the exclusions, the sample has 9,214 transitions. Table 1 reports sample statistics of log consumption and other household characteristics. All statistics are computed using sample weights. The panel is relatively stable over the sample period. Consumption grows considerably between 1987 and

1989 and is stable afterwards. Over time, family size declines while the number of income recipients increases. Other demographic characteristics remain roughly unchanged. The fall in self-employment is paralleled by an increase in public employees.

6 Empirical results

I first present full-sample results. I then address the issue of preference specification and measurement errors in consumption. Finally, I focus on consumption mobility in specific population groups.

6.1 Full sample estimates

There are two methods to construct a transition matrix. One can keep the width of the interval in which consumption is discretized constant and let vary the number of observations within each interval. The alternative is to keep constant the marginal probabilities and let change the interval width, for instance dividing the distribution in discrete quantiles. The second method is more standard, and I proceed using quartiles throughout. Results with deciles are qualitatively similar and are not reported. In what follows, I focus on the distribution of the logarithm of non durable consumption, but results are identical for consumption levels or for any monotonic transformation of consumption.

Table 2 reports the transition matrix when all transitions are pooled over all years. Recall that the SHIW is run every two years, so I observe transitions from period t to period $t+2$. The elements of the main diagonal report the proportion of households that did not change quartile. For instance, the entry in the top left of the table indicates that 66 percent of the households in the first quartile at time t were still in that quartile two years later. Off-diagonal elements signal consumption mobility. For instance, the second entry in the first row indicates that 25 percent of households moved from the first quartile in period $t-2$ to the second quartile in period t . Overall, the table indicates that a substantial amount of consumption mobility takes place over the sample period. About

one third of households in the first quartile moves upwards in the consumption distribution, about one third in the fourth quartile moves downwards, and more than half in the third and fourth quartiles move either upwards or downwards.⁶⁷

Further insights about the evolution of the cross-sectional distribution of consumption can be gained by examining the probability that households move to another quartile in the sample period. In Figure 1 I denote these values as “mobility probabilities”. The probability of moving to a different quartile is relatively high in the second and third quartiles (about 60 percent) and lower in the top and bottom quartiles (between 30 and 40 percent). The figure indicates not only that there is substantial consumption mobility in all quartiles, but also that the mobility is persistent in all survey years. As we shall see, the results of the descriptive evidence is confirmed by the statistical test.

The mobility index corresponding to the elements of the matrix in Table 2 is reported in the first row of Table 3. The statistic has a value of 0.47, with a standard error of 0.005. The null hypothesis of consumption insurance, $S(\mathbf{P}) = 0$, is therefore overwhelmingly rejected. This finding confirms previous studies for the United States that reject the hypothesis of consumption insurance. The other rows of Table 3 report mobility for selected periods of my sample, which is characterized by economic expansion in the early years and by the deep 1991-93 recession. Overall, the results suggest that there has not been a great variability in the consumption mobility (the index ranges from 0.44 to 0.51). In the long-run mobility is still as high as 0.40.⁶⁸

The descriptive and statistical analysis suggest that between 1987 and 1995 the Italian economy was characterized by a substantial amount of consumption mobility. On average, half of the population moves up or down in the consump-

⁶⁷The symmetry of the transition matrix can be tested using the maximum likelihood test suggested by Bishop, Fienberg and Holland (1988). The statistic is of the form $\Psi = \sum_{i>j} \frac{(p_{ij}-p_{ji})^2}{p_{ij}+p_{ji}} \sim \chi_{q(q-1)/2}^2$. The p -value of the test is close to 1, and does not reject the hypothesis that the transition matrix is symmetric.

⁶⁸A Markov process is a stochastic process in which the probability of entering a certain state depends only on the previous state and on the matrix governing the process. If these assumptions hold for the stochastic transition matrix \mathbf{P} , it is possible to determine the limit (or long-run) state as the eigenvector of the matrix \mathbf{P} associated to the eigenvalue 1.

tion distribution every two years, a result that is strongly at variance with full consumption insurance. The counterpart of this finding is that half of the households is unable to insure the idiosyncratic shocks by formal or informal market arrangements. If consumption is regarded as a proxy for permanent income, the results imply that permanent income is not that permanent after all.

Deaton and Paxson (1993) have pointed out that consumption insurance implies that the cross-sectional variance of log-consumption is constant over time. In my sample this hypothesis is not rejected (the p -value associated with this hypothesis is 0.73).⁶⁹ However, it should be clear that the stationarity of the cross-sectional variance does not imply absence of consumption mobility and cannot be used as evidence in favor of consumption insurance. Since the results indicate that consumption is mobile but the variance is roughly constant, it must be the case that the variance of the cross-sectional distribution is not an adequate statistic to measure consumption mobility. This is one case in which simple measures of dispersion must be supplemented by careful analysis of the entire distribution.

6.2 Preference specification and heterogeneity

As mentioned in Section 4.1, the marginal utility of consumption is potentially affected by demographic or labor supply variables that change over time. I would then observe mobility even in the absence of non-insurable shocks. One of the most important demographic variables that can affect preferences is certainly the changing composition of the household. For instance, the arrival of children changes family needs and therefore consumption allocation. I thus compute mobility defining transitions in terms of *per capita* consumption and consumption *per adult equivalent*; the latter is more appropriate in the presence of economies of scale.⁷⁰ The results in Table 4 indicate that using *per capita* consumption makes no difference with respect to Table 3 and that using consumption *per*

⁶⁹More precisely, this is the p -value of a test that $s.d.(\ln c_t) = s.d.(\ln c_{t-1})$.

⁷⁰The number of adult equivalent is defined as: $1+0.8(\text{Number of adults}-1)+0.25(\text{Number of children})$. Data for 1987 are not used because information on the number of children is missing.

adult equivalent increases only slightly the mobility index.

Defining consumption per adult equivalent eliminates just one of the possible sources of predictable consumption variability. In order to account for a larger set of demographic variables that can potentially affect marginal utility, rewrite equation (6) as:

$$\Delta \ln c_{h,t} - \gamma^{-1} \Delta \ln \theta_{h,t} = m_{t+1} \quad (11)$$

where as before $m_{t+1} = -\gamma^{-1} (\Delta \ln \mu_{t+1} - \Delta \ln \pi_{t+1})$. Equation (11) implies that the ratio of marginal utilities for any two households in the cross-section is stationary after controlling for preference shifts. My procedure consists in two steps. In the first step I impute a measure of consumption adjusted for demographic effects, $\ln \tilde{c}_{h,t} = \ln c_{h,t} - \tilde{\gamma}^{-1} \ln \theta_{h,t}$, where $\tilde{\gamma}$ is the OLS estimate of a regression of $\ln c_{h,t}$ on $\ln \theta_{h,t}$. The θ variables that I use are family size, age, age squared, number of children and number of income recipients. I then obtain a measure of consumption where demographic effects have been filtered out. In the second stage I construct transitions on the generated variable $\ln \tilde{c}_{h,t}$ and test the absence of consumption mobility. The resulting mobility index is again quite close to the index estimated without controlling for demographic effects (0.53 with a standard error of 0.005).⁷¹

Leisure is another factor that potentially affects the marginal utility of consumption. In Figure 2 I plot the empirical distribution of annual hours of household heads.⁷² The graph displays the expected concentration of observations at 0 (unemployment or retirement) and 2080 hours (a standard working week of 40 hours). The low variability of the distribution reflects a well-known feature of the Italian labor market where part-time jobs are not widespread. The limited flexibility of hours is *prima facie* evidence that changes in leisure should not be a major factor in explaining consumption mobility.

⁷¹We experiment also with a wider set of demographic variables (number of children in various age bands, education, region of residence, city size). The mobility index is virtually unchanged.

⁷²The density function is estimated non-parametrically by a standard kernel method. We use the optimal bandwidth suggested by Silverman (1986). The 1987 distribution is omitted because data on labor supply are not available.

A more formal test of the effect that leisure has on consumption mobility consists in including leisure in the first-stage regression described above. In this case the mobility index increases to 0.60, regardless of whether log-leisure is instrumented with past values or not. If leisure was responsible for some of the consumption transitions one should observe a decline, not an increase, in the mobility index.⁷³ Therefore, I conclude from this section that misspecification of preferences explains little or nothing of the consumption mobility that I observe in my sample.

As a final check of the potential impact of preference heterogeneity on mobility, I construct a transition matrix for consumption growth rates.⁷⁴ The associated mobility index is 0.81 with a standard error of 0.07, confirming that my sample displays substantial consumption mobility. The finding of higher mobility in growth rates than in levels suggests that also preference heterogeneity is unlikely to explain the rejection of consumption insurance. Therefore, I conclude from this section that misspecification or heterogeneity of preference explains little or nothing of the consumption mobility that I observe in my sample.

6.3 Measurement error

In Section 4.2 I define α as the proportion of the variance of measured log-consumption due to measurement error. Clearly the bias in mobility increases with α . Here I provide evidence on the size of α and the likely impact of measurement error on the estimate of mobility; I also provide bounds of the estimator of mobility in the presence of measurement error in consumption.

Even in the presence of measurement error, complete markets impose strong restrictions on the covariance matrix of consumption. In fact, writing equation (10) as: $\Delta \ln c_{h,t+1}^* = m_{t+1} + \Delta v_{h,t+1}$, omitting the aggregate component, the following testable restrictions are implied by the autocovariance matrix of

⁷³The mobility index can be biased downward if leisure or preference shifts increase consumption needs and if they are negatively correlated with the rank in the consumption distribution, i.e. if they affect more strongly households in the bottom part of the consumption distribution.

⁷⁴This requires at least three years of observations. The sample size for this experiment is therefore reduced to 3,341 transitions.

$(\Delta \ln c^*)$:

$$\begin{aligned} E \left[\left(\Delta \ln c_{h,\tau}^* \right)^2 \right] &= 2\sigma_v^2 \\ E \left[\left(\Delta \ln c_{h,\tau}^* \right) \left(\Delta \ln c_{h,\tau-1}^* \right) \right] &= -\sigma_v^2 \\ E \left[\left(\Delta \ln c_{h,\tau}^* \right) \left(\Delta \ln c_{h,\tau-j}^* \right) \right] &= 0 \text{ for all } j \geq 2 \end{aligned}$$

I am interested in identifying $\alpha = \frac{\sigma_v^2}{\text{var}(\ln c^*)}$. To estimate σ_v^2 , I first define a mean zero measure of *per capita* log consumption adjusted for aggregate shocks:

$$\xi_{h,t} = \ln c_{h,t}^* - \ln \theta_{h,t} - \tilde{\Gamma}_t$$

where θ includes only family size and $\tilde{\Gamma}$ is the cross-sectional mean of consumption per capita. The covariance matrix of $\Delta \xi_{h,t}$ is given in Table 5. At first sight, the pattern is not inconsistent with the restrictions implied by consumption insurance and measurement error: the first order autocovariances are negative and statistically significant, second and higher order autocovariances are small, not statistically significant different from zero. Note also that the pattern of autocovariances denies the possibility of persistent measurement error.

At face value, the covariance matrix suggests that the variance of measurement error is on the order of 0.06 (average over all years). Since the overall variance of consumption is about 0.29 (average over all years), measurement error explains roughly one fifth of the overall variance ($\alpha=0.06/0.29 \simeq 0.2$). By comparison with the PSID, where researchers have found much larger estimates of α (between 70 and 90 percent), the covariance matrix suggests that the SHIW data on total non-durable consumption are of much better quality than the PSID data on food consumption.

Even though $\alpha=0.2$ is not a high number, it must be regarded as an unlikely upper bound for the fraction of the variance explained by measurement error. Recall that this value is obtained on the hypothesis of full consumption insurance. Suppose, however, that consumption insurance does not hold and that an idiosyncratic shock $\eta_{h,t}$ affects consumption growth, $\Delta \ln c_{h,t+1}^* = m_{t+1} + \Delta v_{h,t+1} + \eta_{h,t}$. Assume that $\eta_{h,t}$ is uncorrelated with measurement error

at all leads and lags. The restrictions on the covariance matrix of the adjusted measure of consumption growth can then be rewritten:

$$\begin{aligned} E \left[\left(\Delta \ln c_{h,\tau}^* \right)^2 \right] &= 2\sigma_v^2 + \sigma_\eta^2 \\ E \left[\left(\Delta \ln c_{h,\tau}^* \right) \left(\Delta \ln c_{h,\tau-1}^* \right) \right] &= -\sigma_v^2 \\ E \left[\left(\Delta \ln c_{h,\tau}^* \right) \left(\Delta \ln c_{h,\tau-j}^* \right) \right] &= 0 \text{ for all } j \geq 2 \end{aligned}$$

Note that the restrictions now imply that the variance exceeds, in absolute value, twice the covariance, $E \left[\left(\Delta \ln c_{h,\tau}^* \right)^2 \right] > -2E \left[\left(\Delta \ln c_{h,\tau}^* \right) \left(\Delta \ln c_{h,\tau-1}^* \right) \right]$.⁷⁵ A test that $E \left[\left(\Delta \ln c_{h,\tau}^* \right)^2 \right] = -2E \left[\left(\Delta \ln c_{h,\tau}^* \right) \left(\Delta \ln c_{h,\tau-1}^* \right) \right]$ against the one-sided alternative $E \left[\left(\Delta \ln c_{h,\tau}^* \right)^2 \right] > -2E \left[\left(\Delta \ln c_{h,\tau}^* \right) \left(\Delta \ln c_{h,\tau-1}^* \right) \right]$ rejects the null (the t -statistic is 2.75 with a p -value of 0.003).⁷⁶ The rejection is also apparent from the pattern of covariances reported in Table 5, particularly for 1989-91 and 1991-93. This example indicates that the autocovariance matrix is affected by something other than measurement error alone. To estimate the consumption variability that cannot be attributed to measurement error, note that:

$$\sigma_\eta^2 = E \left[\left(\Delta \ln c_{h,\tau}^* \right)^2 \right] + 2E \left[\left(\Delta \ln c_{h,\tau}^* \right) \left(\Delta \ln c_{h,\tau-1}^* \right) \right] \quad (12)$$

One possibility is to choose values of σ_η^2 that are consistent with a significance level of 5 percent or higher. In my sample the null hypothesis that $\sigma_\eta^2 = 0.025$ has a p -value of 0.049 and therefore cannot be rejected. This implies $\sigma_v^2 = 0.035$ (averaged over all years) and $\alpha = 0.12$. In more realistic examples, first-order autocovariances alone are not sufficient to identify σ_v^2 , so that α is likely to be lower than 0.12. For instance, if the idiosyncratic shock η is persistent one cannot disentangle the fraction of the variance due to measurement error from that due to shocks.⁷⁷ From the foregoing, I conclude that $\alpha=0.12$ is an upper bound to measurement error and that more realistic values of α range from 0.05 to 0.10.

⁷⁵Equivalently, under the hypothesis maintained, this implies $\sigma_\eta^2 > 0$.

⁷⁶This test is pooled over all years. For single years, the null hypothesis is rejected for 1991 and 1993 but not for 1995.

⁷⁷To take one example, if $\eta_{h,t} = \psi_{h,t} - \rho\psi_{h,t-1}$, the restrictions can be rewritten as:

The next step is to assess how measurement error affects the mobility index under the null hypothesis of consumption insurance. For this purpose, I design a Montecarlo simulation based on 100 replications, using *per capita* consumption throughout. In each year I choose a sample size identical to the number of transitions (for instance, it is 3,211 for 1993-95). Measurement errors at times t and $t - 2$ are drawn from a normal distribution with mean zero and variance α times the variance of measured consumption at t and $t - 2$. True consumption $\ln c_{t-2}$ is drawn from a normal distribution with mean equal to the mean of measured consumption and variance of $(1 - \alpha)$ times the variance of measured consumption at $t - 2$. Under the null hypothesis of consumption insurance, $\ln c_t = m_t + \ln c_{t-2}$, where m_t is the aggregate consumption growth, estimated as the average of individual consumption growth rates between $t-2$ and t . Given my assessment of the likely magnitude of measurement error, I choose values for $\alpha = \{0.05, 0.1, 0.12\}$ and simulate the mobility index $S(\tilde{\mathbf{P}})$.

The results of the simulation are reported in Table 6. The first column reproduces the actual mobility index $S(\hat{\mathbf{P}})$ of consumption *per capita*, the same as in Table 4, column 2. If $\alpha=0.05$ the simulated index $S(\tilde{\mathbf{P}})=0.26$ in 1987-89, against $S(\hat{\mathbf{P}})=0.47$. The fraction of mobility that cannot be attributed to measurement error is $\frac{S(\hat{\mathbf{P}})-S(\tilde{\mathbf{P}})}{1-S(\tilde{\mathbf{P}})}=0.29$. If $\alpha=0.10$ this fraction is 0.19; even in the most unfavorable case of $\alpha=0.12$ the fraction of “true” mobility is 0.15. To summarize, in 1987-89 the fraction of households that move across the consumption distribution for reasons other than measurement error ranges from 15 ($\alpha = 0.12$) to 47 percent ($\alpha = 0$). Similar results are obtained for transitions in other years.⁷⁸

$$\begin{aligned}
 E \left[(\Delta \ln c_{h,\tau}^*)^2 \right] &= 2\sigma_v^2 + (1 + \rho^2)\sigma_\psi^2 \\
 E \left[(\Delta \ln c_{h,\tau}^*) (\Delta \ln c_{h,\tau-1}^*) \right] &= -\sigma_v^2 - \rho\sigma_\psi^2 \\
 E \left[(\Delta \ln c_{h,\tau}^*) (\Delta \ln c_{h,\tau-j}^*) \right] &= 0 \text{ for all } j \geq 2
 \end{aligned}$$

and identification would no longer be possible.

⁷⁸Focusing on households with low rates of growth of consumption growth ($-1 \leq \Delta \ln c \leq 1$) to minimize the impact of measurement errors has virtually no effect on the results (the mobility index is 0.48).

6.4 Sub-sample estimates

My statistical test allows me to examine which population groups are more exposed to idiosyncratic shocks. To evaluate if there are differences in consumption mobility I use the statistic on difference of means discussed in Section 3. I use a measure of consumption *per capita* throughout.

Table 7 reports consumption mobility for households living in different regions, occupations (public vs. private and self-employed vs. employee), education, year of birth and income recipients groups. The index shows that mobility is larger in the North than in the South (0.50 against 0.48), a difference possibly explained by the greater social insurance role offered by the family and the presence of informal market arrangements in the South. Mobility is also higher in the private sector than in the public sector (0.49 against 0.45), a reflection of the fact that in Italy public sector employees enjoy stable earnings tied to strict seniority rules rather than performance and face little unemployment risk. The employee, which face less income risk than the self-employed, also exhibit lower consumption mobility (0.47 against 0.50). The difference between households with compulsory education is not statistically significant from that of households with college degrees. The comparison by year of birth suggests that younger cohorts are progressively more able to smooth away idiosyncratic shocks.

Common sense suggests that households with multiple earners can insure income shocks better than single earners. I distinguish three groups: households with no change in number of earners, and households with negative or positive change, respectively. The results indicate that mobility attains its maximum among those that experience a decline or an increase in the number of income recipients (0.55 and 0.56 respectively).

Overall, I find plausible and significant variation of mobility by occupation and demographic groups. Some of the differences between groups can be tied to specific hypothesis about the actual working of credit, insurance and informal markets. This is certainly the case for the relative low amount of mobility of public sector employees and households that experience no variability in the

number of income recipients. However, contrary to my expectations, overall I find a surprisingly small amount of variability between different groups. Table 7 indicates that even if in most cases the mobility indexes are statistically different from each other, the p -values associated with the difference in means generally indicate marginal rejection of the null hypothesis of equality between groups. Furthermore, the differences of the mobility index between groups are generally not large in absolute value.

7 Conclusions

Consumption insurance implies that in any time period the initial cross-sectional distribution of consumption is a sufficient statistic for all future distributions. This implication of consumption insurance is as yet unexplored. I construct a transition matrix for total non-durable consumption using the 1987-95 panel contained in the Bank of Italy Survey of Household Income and Wealth. I then summarize the transition matrix of consumption by an appropriate mobility index. The test of consumption insurance I propose is simple and powerful. Most importantly, the non-parametric test proposed does not depend on functional form, identification assumptions about the source of idiosyncratic shocks, or their potential correlation with omitted preference shifts.

I find that roughly 50 percent of households move up or down in the consumption distribution between any two periods, in both the short and the long run. This constitutes very strong evidence against consumption insurance. There are interesting and expected variations in the mobility patterns within different population groups, but overall the inter-group variation in mobility is not large. The mobility observed is unlikely to be explained by the effect of preference shifts. Consumption *per capita* and per adult equivalent exhibit mobility comparable to that of total non-durable consumption. When I control for other potentially important observable preference shifts (such as family size, age, education, and leisure) mobility actually increases. Finally, I find substantial mobility in consumption growth, not only in consumption levels, implying that preference

heterogeneity does not explain rejection of consumption insurance.

Part of the consumption mobility observed in the sample may be due to measurement error. I show that in my data measurement error is unlikely to explain a large fraction of the total cross-sectional variance of consumption. To assess the impact of measurement error on the mobility index I then perform a Montecarlo experiment. The simulation shows that my test rejects the hypothesis of consumption insurance even in the most unfavourable case, one in which measurement error has the highest impact on mobility. I conclude that in Italy a great deal of consumption mobility is explained by idiosyncratic shocks that households are unable to insure. Even though my test is powerful disproof of consumption insurance, it does not constitute evidence either for or against the permanent income hypothesis; these “are distinct propositions, and each may hold independently of the other” (Cochrane, 1991). In future research I plan to use transition matrices to model consumption and income mobility jointly and look into their implications for smoothing idiosyncratic shocks across different states of nature.

Table 1
Descriptive statistics⁷⁹

	1987	1989	1991	1993	1995	All years
$\ln c_t$	9.90	10.08	10.02	10.01	10.00	10.02
$var(\ln c_t)$	0.26	0.26	0.29	0.29	0.27	0.28
South	0.41	0.37	0.34	0.36	0.39	0.37
North	0.43	0.46	0.48	0.47	0.43	0.46
Family size	3.15	3.12	3.04	3.07	3.01	3.07
Self-employed	0.20	0.17	0.17	0.16	0.15	0.16
Public employee	0.17	0.18	0.27	0.23	0.25	0.23
Years of schooling	7.38	7.97	8.19	8.03	8.10	8.03
Born \leq 1927	0.33	0.29	0.26	0.24	0.24	0.26
Born 1928-1937	0.20	0.21	0.20	0.18	0.19	0.19
Born 1938-1947	0.26	0.25	0.24	0.23	0.22	0.55
Income recipients	1.63	1.72	1.72	1.74	1.78	1.73
# of obs.	1,097	2,717	4,036	4,006	3,211	15,067

⁷⁹Cross-sectional means and variances are computed using sample weights.

Table 2
The transition matrix of consumption⁸⁰

Quartile at time $t - 2$	Quartile at time t			
	1 st	2 nd	3 rd	4 th
1 st	0.66	0.25	0.07	0.02
2 nd	0.24	0.41	0.26	0.09
3 rd	0.09	0.26	0.41	0.24
4 th	0.02	0.10	0.25	0.63

Table 3
Mobility index⁸¹

Panel	Number of transitions	$S(\hat{\mathbf{P}})$	$s.e.(S(\hat{\mathbf{P}}))$
All years	9,204	0.4729	0.0051
1987-1989	1,097	0.5066	0.0146
1989-1991	1,914	0.4621	0.0110
1991-1993	2,982	0.5060	0.0090
1993-1995	3,211	0.4367	0.0085

⁸⁰The table reports consumption transitions from period t to period $t + 2$. Transitions are pooled over all sample years.

⁸¹The table reports the mobility index computed using the transition matrix pooled over all sample periods and for separate sample periods.

Table 4
Computing mobility with different consumption measures⁸²

	Number of transitions	Consumption per capita	Consumption per adult equivalent	Consumption filtered with demographic variables
	(1)	(2)	(3)	(4)
All years	8,107	0.4713 (0.0050)	0.5005 (0.0054)	0.5306 (0.0054)
1987-89	1,097	0.4702 (0.0146)	n.a.	n.a.
1989-91	1,914	0.4705 (0.0110)	0.5117 (0.0110)	0.5439 (0.0110)
1991-93	2,982	0.5029 (0.0089)	0.5273 (0.0089)	0.5594 (0.0089)
1993-95	3,211	0.4432 (0.0085)	0.4689 (0.0086)	0.4962 (0.0085)

⁸²In column (1) the number of transitions for the row "All years" is 9,204. In 1987-89 the index cannot be computed because information on the number of children is missing in 1987. Standard errors are reported in parenthesis.

Table 5**The autocovariance matrix of consumption growth⁸³**

	1989	1991	1993	1995
1989	0.1405 (0.0071)			
1991	-0.0443 (0.0081)	0.1398 (0.0056)		
1993	-0.0061 (0.0143)	-0.0643 (0.0070)	0.1748 (0.0064)	
1995	0.0121 (0.0162)	0.0049 (0.0055)	-0.0637 (0.0050)	0.1304 (0.0047)

Table 6**Correcting for measurement errors⁸⁴**

Year	Actual results		Montecarlo results		
	$S(\hat{\mathbf{P}})$	$s.e.(S(\hat{\mathbf{P}}))$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.12$
	$S(\hat{\mathbf{P}})$	$s.e.(S(\hat{\mathbf{P}}))$	$S(\tilde{\mathbf{P}})$	$S(\tilde{\mathbf{P}})$	$S(\tilde{\mathbf{P}})$
1987-89	0.4702	0.0146	0.2572	0.3495	0.3749
1989-91	0.4705	0.0110	0.2650	0.3609	0.3882
1991-93	0.5029	0.0089	0.2616	0.3550	0.3825
1993-95	0.4432	0.0085	0.2564	0.3482	0.3772

⁸³Standard errors are reported in parenthesis.⁸⁴The Montecarlo simulation is described in Section 5.

Table 7
Computing mobility by demographic groups⁸⁵

Group		Group mobility		Diff. of means		
		$S(\hat{P})$	$s.e.(S(\hat{P}))$	<i>Gr.2</i>	<i>Gr.3</i>	<i>Gr.4</i>
1	North	0.5066	0.0079	2.51	0.92	
2	South	0.4782	0.0080		1.10	
3	Center	0.4937	0.0115			
1	Public	0.4470	0.0113	2.90		
2	Private	0.4865	0.0076			
1	Self-employed	0.5016	0.0183	1.75		
2	Employee	0.4681	0.0056			
1	Compulsory ed.	0.4881	0.0062	2.81	0.66	
2	High school	0.4524	0.0111		1.78	
3	College	0.4750	0.0188			
1	Born \leq 1927	0.5002	0.0102	0.81	2.14	3.01
2	Born 1928-37	0.4881	0.0110		1.26	2.03
3	Born 1938-47	0.4690	0.0104			0.73
4	Born >1947	0.4589	0.0092			
1	No change in earners	0.4639	0.0055	4.62	4.80	
2	Pos. change in earners	0.5615	0.0204		0.51	
3	Neg. change in earners	0.5480	0.0166			

⁸⁵The entries in the columns labelled *Gr.2*, *Gr.3*, and *Gr.4* report the *Z*-statistic associated with the test that mobility for the group in that row equals mobility for groups 2, 3, and 4, respectively.

5

Income risk dynamics and heterogeneity

1 Introduction

At present, there is little or no evidence for the existence of formal or informal mechanisms allowing full insurance of idiosyncratic shocks (Cochrane, 1991; Attanasio and Davis, 1996).⁸⁶ As emphasized in recent theoretical work, the absence of insurance markets for idiosyncratic shocks introduces additional difficulties when calibrating models of both general and partial equilibrium, in particular that of measuring microeconomic uncertainty. This has been recently argued by Browning, Hansen and Heckman (1998). In their own words, calibrating model economies under imperfect insurance “requires a measure of the magnitude of microeconomic uncertainty, and how that uncertainty evolves over the business cycle [...]. This introduces the possibility of additional sources of heterogeneity because different economic agents may confront fundamentally different risks. To calibrate the macroeconomic model it becomes crucial to measure the distribution of individual shocks”.

These remarks have implications for many areas of research. Measuring uncer-

⁸⁶See chapter 4 for the implicationd of full consumption insurance.

tainty is crucial when trying to determine whether and to what extent households save for precautionary motives (Deaton, 1992b). Variations in individual uncertainty are also shown to affect the width of the inaction band in Ss models of durable demand and investment (Eberly, 1994). Measurement of individual risk is important in models of general equilibrium that depart from the full insurance assumption and assume limited commitment or private information (Alvarez and Jermann, 1998; Cole and Kocherlakota, 1997). Under background uncertainty (Kimball, 1990), the possibility of confronting a substantial amount of uninsurable risk may lead prudent individuals to demand a higher risk premium on risky assets, thus providing an explanation for the equity premium puzzle. Background uncertainty can also generate more subtle business-cycle propagation effects, as recently argued by Costain (1998). With imperfect insurance and counter-cyclical cross-sectional variation (i.e., microeconomic uncertainty raising during economic slumps), a fall in output increases individual risk; this leads to a portfolio re-allocation towards riskless assets and away from risky investments; if the latter proxies for the firms' productive capacity, a further decline in output may occur. Thus, in general one would expect fluctuations in individual uncertainty to generate important fluctuations in aggregate saving, investment and growth. Undertaking the task of modelling individual risk in a world of incomplete markets thus seems critical. This is the task that motivates this study and that I address empirically.

I consider the possibility that the income process is actually more complicated than the one considered so far in the literature, with both heterogeneity and higher moment dynamics shaping the evolution of the distribution of income over time as well as across individuals. As done elsewhere in the literature, I posit a model where earnings shift over time because of transitory and permanent unanticipated fluctuations. Part of the transitory fluctuation is in reality measurement error, but part of it is genuine innovation. Transitory shocks to individual productivity include overtime labour supply, piece-rate compensation, bonuses and *premia*, etc.; in general, such shocks are mean reverting, e.g. their effect does not last long. On the other hand, some of the innovations to earn-

ings are highly persistent or non-mean reverting, e.g. their effect cumulates over time. Example of permanent innovations are associated to job mobility, long-term unemployment or promotions.

I consider two different models of earnings determination. In the first model I adopt the canonical decomposition of unobservable components of earnings between a permanent time-invariant effect and a transitory effect endowed with some persistence. The second model is more complicated, and specifies earnings as the sum of a martingale component and a persistent transitory disturbance. This implies that shocks of different nature are not separately identifiable (unless one combines subjective expectations with actual realizations, as shown in chapter 3). One of the interesting features of this study is that I do not assume *a priori* the validity of a model of earnings determination, but explicitly test for the presence of a permanent component in income using a simple statistic. This is accomplished by using the theoretical restrictions on the autocovariance matrix of the residuals of the earnings function in first differences. Moreover, I consider the possibility that individuals with different education levels face different income processes, both at the mean and the variance level; while imperfect, this is the strategy I choose to control for the possibility of changing returns from observable skills remarked in the literature on earnings inequality.

To face the challenge of modelling income risk dynamics, I focus on models for the conditional variance of income shocks, that is, the appropriate measure of risk emphasized by the theory. In particular, I propose an econometric strategy that identifies the parameters of the truly stochastic second moments of income. Allowing for stochastic risk, as I do, is far from being a mere statistical exercise. As showed by Caballero (1990), once higher moments uncertainty is taken into account, reasonable parametric assumptions provide an explanation for the empirical puzzles emerging in tests of the permanent income hypothesis, such as the excess smoothness and the excess sensitivity puzzle (see Attanasio, 1999, for a recent survey). In a different context, Hassler (1996) has showed that the introduction of stochastic risk in a partial equilibrium-*Ss* model of durable demand (or investment) can explain excess delay in the adjustment to the optimal

stock.

Identifying the conditional variances of unobservable error components (which occurs in the second model I examine) is not a trivial task, and considerable effort must be paid to discuss such issue. To pre-empt, I have to combine information on the variance and on the autocovariances at various length of the earnings residuals in first difference to achieve identification. This is because the variance of earnings growth is a mixture of the variance of shocks of different nature occurring at different dates, so that the variance of the permanent shock can be identified by filtering out the variances of the transitory shock from the overall variance of earnings growth. Similarly, since growth rates of earnings covary over time only because of a common transitory component, the variance of the latter is identified by using information on the autocovariances alone.

Drawing partly from the empirical literature that models the uncertainty related to changes in security prices (Engle, 1989), I then assume that the variance of income innovations follows an autoregressive process with both observable and unobservable heterogeneity (i.e., income innovations are *ARCH* processes). Individual heterogeneity has the role of generating cross-sectional, long-term differences in risk across individuals, while aggregate effects capture fluctuations in risk over the business cycle. Finally, persistence in variance provides a simple markovian rule relating income variability at two different calendar years.

The rest of the chapter is organized as follows. In section 2, I discuss and motivate my approach for modelling the evolution of the income distribution. I then show how the various components are identified under the two different models. In section 3 I illustrate the data used in the empirical application; as in many other studies, I conduct my analysis on the Panel Study of Income Dynamics (heretofore, PSID). The long period of observation plays a double role. On one hand, it gives us the opportunity to distinguish between a fixed effect model of earnings, where income depends on a set of fixed but unobservable attributes, and a model with permanent shocks and continuously evolving unobserved productivity effects; on the other hand, it provides the time series variability required for earnings shocks to converge to their steady state configuration. In section 4 I

present and discuss my empirical findings for the two proposed models, while in section 5 I analyse the implications of the empirical results, focusing in particular on the implications for precautionary savings. Section 6 concludes.

2 Modelling the distribution of earnings

2.1 A brief survey of the literature

Models for the evolution of the earnings distribution in the US have been proposed by several authors. Here I will focus on the contributions by Lillard and Willis (1978), MaCurdy (1982), Abowd and Card (1989), Gottschalk and Moffitt (1996) and Baker (1997). This is for two main reasons: first, all the papers above contain (explicitly or implicitly) a separation between shocks that are mean reverting and shocks that are not; second, all of them focus on PSID data.

Lillard and Willis study earnings mobility in a sample of 1,144 males drawn from the 1967-1973 PSID. Individual heterogeneity is randomly distributed in the population and transitory effects are assumed to follow a stochastic AR(1) process (the “auto-correlated individual component model”). When a limited set of control variables is used, they find that 73 per cent of total variance is due to permanent earnings differences, 22 per cent to purely stochastic variation and 5 per cent to serial correlation. Assuming that permanent and transitory components are jointly normally distributed, they find a great deal of earnings mobility and conclude that poverty is not a permanent status (55 per cent of whites and 35 per cent of blacks leave poverty status between any two adjacent years).

MaCurdy’s influential paper applies standard time series techniques to model the serial correlation matrix of earnings residuals using a sample of 531 white males from the 1967-1976 PSID. To avoid his estimates being contaminated by influential outliers, he uses stringent sample selection criteria. He admits the presence of individual heterogeneity in levels and an individual random effect

in growth rates (the “random growth rate model”). Estimating the parameters of the earnings process by *quasi* maximum likelihood, he reach the conclusion that the earnings shock can be modelled as a non-stationary $ARMA(1,2)$ or $ARMA(2,1)$ process. Given the evidence for (roughly linear) non-stationarity, he favours the unit root model rather than the random growth model.

Abowd and Card study simultaneously movements in earnings and movements in hours over time. They find that a component-of-variance model with three sources of earnings variation successfully summarizes the covariance structure from three different survey samples: 1,448 males from the 1969-1979 PSID, 1,318 males from the 1966-1975 National Longitudinal Survey of Men 45-59, and 560 males drawn from the Seattle-Denver Income Experiment. The first component of earnings variation is a time-stationary serially uncorrelated measurement error. The second component is a stationary component common to both earnings and hours changes, while the third component is a transitory one affecting only the variances and the contemporaneous covariances of earnings and hours. As in MaCurdy, they cannot reject the unit root model.

Gottschalk and Moffitt present an approach that is very close in spirit to the one I use in this study. Using again the PSID (1967-1987), they model earnings as the sum of a transitory component (which follows a low order serially correlated process) and a permanent component that is a random walk stochastic process. They also attempt to provide a structure for the finding of non-stationarity; their main conclusion is that in the period under observation the variance of the transitory component has grown at a rate similar to the variance of the permanent component.

Using a sample of 534 males drawn from the 1967-1986 PSID and following the separate approach of MaCurdy and Lillard and Willis, Baker (1997) models two different earnings functions in a nested framework. In the first case, the earnings shocks are assumed to be the sum of a component reflecting heterogeneity (both in levels and growth rates) and a serially uncorrelated transitory component (the “random growth rate” model). In the second case, earnings shocks are specified as random walks to allow for permanent effects. According to his empirical

analysis, the preferred model of the data includes person-specific earnings profiles and a component that captures serially correlated measurement error or short-lived innovation to earnings. Yet, the rejection of the unit root model is probably due to the bad properties of his estimators in the presence of a relatively small sample size.

2.2 Aspects of the distribution of earnings

To understand how microeconomic uncertainty evolves when insurance markets for idiosyncratic shocks are missing, I need to characterize various aspects of the distribution of earnings. For instance, I need to know which part of the year to year income variance is due to observable or unobservable heterogeneity and which to shocks to productivity. I also need to know what proportion of the variance of the innovation is due to permanent shocks and what to transitory shocks. The welfare implications of these two types of shocks are obviously very different. From the point of view of consumption choices, a permanent shock is incorporated in its entirety on permanent income, while only the annuity value of a transitory shock matters (see chapter 1, and Deaton, 1992b). Following the most recent literature, I distinguish between permanent and transitory shocks but provide empirical tests that allow me to test the null hypothesis of permanent shocks.

As pointed out above, a critical aspect of my analysis is to assess whether the variance of the shocks to income is itself state dependent, and to what extent it depends on (observed and unobserved) individual characteristics. Such information is important in modelling the behaviour of rational agents in a risky environment. In principle, all the moments of the conditional distribution of income enter the intertemporal optimality condition for consumption.⁸⁷ While desirable, modelling third or higher moments is complicated, essentially because of identification problems, and is beyond the scope of this study.⁸⁸

⁸⁷This can be seen when taking a Taylor expansion of the Euler equation $E_t u'(c_{t+1}) = \frac{1+r}{1+\delta} u'(c_t)$ around $E_t c_{t+1}$.

⁸⁸Yet, if normality holds, the first two moments characterize the entire distribution of earnings.

2.3 The conditional mean of earnings

As in previous empirical work, I posit the following model for the conditional mean of log earnings:

$$A^e(L, p)y_{it} = m_t^e + \beta^e X_{it} + \delta^e c_i + u_{it} \quad (1)$$

where y_{it} is the log of real earnings, $A^e(L, p)$ a lag polynomial of order p (i.e., $A^e(L, p)y_{it} = \sum_{j=0}^p \alpha_j^e y_{it-j}$, with $\alpha_0^e \equiv 1$) with education-varying coefficients (the superscript “e” standing for education), m_t^e a year effect,⁸⁹ X_{it} a vector of time-varying observable characteristics (such as a polynomial in age to capture deterministic life-cycle effects), c_i includes year of birth cohort, race and all observable time invariant individual attributes, and u_{it} is a shock to earnings. A distinguishing feature of this model is the assumption that the returns from individual attributes vary across education groups in a very general way.⁹⁰ Economy-wide shocks are estimated by directly inserting year dummies in (1).

I consider two different models of earnings determination. In the first model (model A) I assume that:

$$u_{it} = p_i + \varepsilon_{it} \quad (2)$$

where p_i is a time invariant permanent component, and ε_{it} a transitory fluctuation. While this specification is often rejected by the data (see Moffitt and Gottschalk, 1994), it serves the scope of introducing my argument in a straightforward fashion. Thus, I shall not emphasize it as the most appropriate way to specify the evolution of the earnings distribution.

⁸⁹By inserting time dummies to control for aggregate effects we are explicitly removing aggregate risk. In principle, it would be possible to follow alternative strategies, such as identifying aggregate shocks from the composite error term.

⁹⁰A possible extension to this model is to assume that the return from individual attributes varies over time as well as across education. Given the evidence on increasing returns to both observed and unobserved individual skills in the US, this is probably an extension that would be worth considering. While interesting, I decided not to pursue such line of research in this paper.

In the second model of earnings determination (model B), I assume that income shocks are decomposed into a transitory innovation with low persistence and a martingale permanent component. Therefore:

$$u_{it} = p_{it} + \varepsilon_{it} \quad (3)$$

with ε_{it} being the transitory shock and p_{it} the permanent component of income. I assume that the latter, if present, follows the process:

$$p_{it} = p_{it-1} + \zeta_{it} \quad (4)$$

I show below how to test for the presence of a permanent shock in the conditional mean process. As far as the transitory shock is concerned, I assume that it can be modelled as an $MA(q)$ process, with the order q of the process to be determined at empirical level. This is explained below. I assume that the permanent shock and the transitory shock are uncorrelated at all leads and lags. Now I turn to the identification of the parameters of the conditional mean process.

For model A, identification does not pose many problems. I follow the standard practice of filtering out deterministic business-cycle effects (the year dummies) and life-cycle effects (the age polynomial) and then estimating the parameters of the autoregressive polynomial $A^e(L, p)$ on the resulting earnings residuals obtained after first-differencing. This gets rid of the permanent component p_i and provide consistent estimates of $A^e(L, p)$.⁹¹ In particular, I use a *GMM* procedure (Arellano and Bond, 1991a) based on the orthogonality condition $E(\Delta\varepsilon_{it} | \Omega_{t-2}) = 0$, where Ω_{t-2} represents the history of the income process up to the period $t-2$, and Δ is the first-difference operator, i.e. $\Delta a_{it} = a_{it} - a_{it-1}$. I assume that there is no further serial correlation over and above that allowed by $A^e(L, p)$, but provide formal tests for it. The individual-specific time average of the estimate of $(A^e(L, p)y_{it} - m_t^e - \beta^e X_{it}) = \tilde{y}_{it}$ provides a consistent estimate of the individual specific component $(\delta^e c_i + p_i) = \tilde{p}_i$. Finally, a consistent estimate of the transitory shock is obtained as the difference: $(\tilde{y}_{it} - \tilde{p}_i)$. Note that in this

⁹¹I choose a two-step procedures for the aim of identifying transitory shocks in levels. Ideally, one could estimate the parameters of (1) from a specification in first difference.

model I need to invoke the large- T assumption at individual level; this is relaxed in the other models I estimate.

Identification in model B is slightly more complicated. In the presence of a martingale permanent component and/or serially correlated transitory effects, estimates in levels are inconsistent. First differencing the data allows one to get rid of this problem and yields:

$$A^e(L, p)\Delta y_{it} = \Delta m_t^e + \beta^e \Delta X_{it} + v_{it} \quad (5)$$

where $v_{it} = \Delta u_{it}$. Note that with $p = 1$ and $q = 0$ the rate of growth of earnings follows an $ARMA(1, 1)$ process, a finding on which various studies agree (in particular, MaCurdy, 1982, and Abowd and Card, 1989). As I shall see, this assumption is consistent also with the PSID sample I work with. The only relevant distinction is that in my case the autoregressive bit is imposed directly on earnings rather than on the transitory shock (i.e., as in partial adjustment models), while the moving average bit $(\zeta_{it} + \varepsilon_{it} - \varepsilon_{it-1}) = \nu_{it} - \epsilon\nu_{it-1}$, with the implicit MA coefficient ϵ being a function of the variance of the transitory shock and the variance of the permanent shock (Deaton, 1991). The orthogonality condition that identifies the parameters of (5) is therefore:⁹²

$$E(v_{it} | \Omega_{t-q-2}) = 0 \quad (6)$$

In the absence of structural restrictions, it is practically impossible to identify separately the order of the AR process from the order of the MA process for the transitory shock. I thus follow the strategy of setting *a priori* the order of the AR process and choose the order of the MA process that is statistically consistent with (6). For instance, the over-identifying restrictions implied by (6) should hold with respect to information dated $t - 2$ if transitory shocks are not serially correlated (i.e., if $q = 0$).

Next, I need to test for the presence of a permanent shock. Note that under

⁹²Notice that the presence of a permanent shock does not alter the stochastic lag structure of the orthogonality condition (6), although of course the residual of the model in first differences will differ in the presence of a permanent shock.

the null of no permanent shocks the error term in the model in first differences (5) captures the growth in the transitory shock, i.e.:⁹³

$$v_{it} = \Delta\varepsilon_{it} \quad (7)$$

whereas under the alternative:

$$v_{it} = \zeta_{it} + \Delta\varepsilon_{it} \quad (8)$$

Let us assume for simplicity that the transitory shocks are not serially correlated. It is worth noting that information on the variance or the autocovariance of v_{it} alone will not be sufficient to distinguish empirically (7) from (8). The reason is that in both cases v_{it} follows an $MA(1)$ process. However, combining information on both the variance and the autocovariance of v_{it} allows to disentangle empirically (7) from (8). One can show that whether or not permanent shocks are present the following restriction applies:

$$E(v_{it+1}v_{it}) = -E(\varepsilon_{it}^2) \quad (9)$$

This is because the (adjusted) growth rates of earnings covary over time only because of a common transitory component. Under the null hypothesis of no permanent shocks, this would imply:

$$E(v_{it}^2) = E(\varepsilon_{it}^2) + E(\varepsilon_{it-1}^2) \quad (10)$$

Using (9) and (10) one yields:

$$E(v_{it}^2 + v_{it+1}v_{it} + v_{it}v_{it-1}) = 0 \quad (11)$$

Under the alternative hypothesis, the same expression equals the variance of the permanent shock, namely:

⁹³The null hypothesis of my test encompasses two different assumptions concerning the structure of the error term: either that a permanent component is absent altogether or that is time invariant (the model A). The alternative hypothesis is that the permanent component follows a martingale process (the model B).

$$E\left(v_{it}^2 + v_{it+1}v_{it} + v_{it}v_{it-1}\right) = E\left(\zeta_{it}^2\right) \quad (12)$$

In words, equation (12) states that since the variance of (adjusted) earnings growth is a mixture of the variance of shocks of different nature occurring at different dates, the variance of the permanent shock can be obtained by filtering out the variances of the transitory shock from the overall variance of earnings growth. Since I do not observe the true shocks to earnings, I am bound to use the predicted composite residual from (5), i.e.: $\hat{v}_{it} = \hat{A}^e(L)\Delta y_{it} - \Delta \hat{m}_t^e - \hat{\beta}^e \Delta X_{it}$, to implement the test discussed above.⁹⁴ The test is based on the unconditional distribution of the earnings composite residuals in first differences. By the central limit theorem, the test statistic is asymptotically standard normal.⁹⁵

2.4 The conditional variance of earnings

Before detailing the identification strategy for the conditional variance of earnings shocks, it is important to describe how current empirical research approaches the problem of measuring microeconomic uncertainty. This research is particularly active in the precautionary savings literature, and I will mainly refer to that. Various studies calibrate model economies under partial or general equilibrium using as a measure of uncertainty the variance of earnings innovations estimated in studies such as MaCurdy (1982), Hall and Mishkin (1982), and

⁹⁴It is worth noting that the only assumption I need to invoke in order to implement the test above is the absence of serial correlation for the transitory component (this is consistent with the results of our empirical analysis, see section 4.2). If serial correlation did in fact exist, a more complicated test statistics could be derived, but in this case I would need to invoke the much stronger assumption of covariance stationarity. Note also that using predicted residuals in the place of the true disturbances is an asymptotically valid operation as long as the fourth moments exist and these are constant across individuals (MaCurdy, 1982).

⁹⁵This is a standard one-sided test. Note that although I may accept the null hypothesis, I may still have permanent time-invariant effects in levels (i.e. $u_{it} = p_i + \varepsilon_{it}$), as in model A. Using the test outlined in the main text with a comparison between estimation in levels and after transforming the data I can distinguish between three models of interest: no permanent component, a time-invariant permanent component (model A), and a stochastically evolving permanent component (i.e., a martingale in levels, as in model B).

Lillard and Willis (1978). They all estimate *ARMA* processes for earnings or family income using microeconomic data drawn from the PSID. However, as measured in these studies, microeconomic uncertainty varies neither over time nor across individuals. Recent studies allow for considerably more heterogeneity when measuring risk (Carroll, 1992; Hubbard, Skinner and Zeldes, 1995). Yet, the approach followed in these studies is not without problems, in particular when attempts are made to calculate comparative statics results. First, dynamics in income volatility is neglected; moreover, the focus of the analysis is on unconditional variances; finally, all variability in income is assumed to translate on a one-for-one basis in consumption variability. The assumption of i.i.d. income innovations of course has the advantage of simplifying the search for the numerical solution; however, as shown in studies in the labour economics literature, it is hardly realistic. The approach taken so far is tantamount to assigning a time-invariant measure of income risk to each household, assuming away variation over the business cycle and over the life cycle. However, uncertainty does decline with age and during economic booms. The focus on unconditional variances is moreover at odds with the theory. In fact, what generates precautionary savings is the perception of future income variability, not variability itself. Ideally, one would like to use as a proper measure of uncertainty the subjective variance of future income, a measure that is rarely available, at least for the US (see Dominitz and Manski, 1998). Finally, the possibility of insuring some of the idiosyncratic shocks through various informal or formal mechanisms implies that not all the variability I observe or measure is necessarily uncertainty as perceived by the household. As remarked by Zeldes (1992) the representative paper in this literature “assumes that individuals have always faced –and always will face– the same uncertainty”. Comparative statics results are obtained by allowing for changes in uncertainty and examining the resulting dynamic effects. “That is, people face a change to which they had ex-ante assigned zero probability. While this approach used to be a commonly performed comparative statics exercise, it is an inappropriate experiment to perform. It is especially problematic” in models such as the permanent income hypothesis, “where the whole point of

the exercise is that people are forward-looking and take into account all possible contingent outcomes. The correct way to do the calculations [...] is to let people face a probability distribution for [...] uncertainty, and then examine the optimal response to various realizations". I follow Zeldes' suggestion in considering a stochastic process for the variance of earnings shocks.

The problem that arises when modelling the variance function is that, unlike the mean process, the variance process lacks a well defined theoretical structure.⁹⁶ Here I follow the strategy of modelling the conditional variance of earnings shocks as a parsimonious autoregressive process with unobservable heterogeneity, but acknowledge that different structures might be admissible as well. I draw such structure from the empirical literature on *ARCH* processes, although I must adjust it to the fact that I deal with low frequency longitudinal data.⁹⁷

Also in the context of the variance function, I consider two models concerning the nature of income shocks. The main distinction between the two models is that in the first case I can identify income shocks as the residual of the estimated income process (see above), and therefore all higher moments directly. In contrast, in the second model no separate estimate of income shock is available if a permanent component is present (unless one combines subjective expectations with actual realizations, as in chapter 3 of this dissertation). This implies a much more complicated process of identification. A further complication, present in both models, is that since earnings innovation might be conditionally heteroscedastic because workers receive shocks from different distribution based on their skills and other unobservable characteristics, I need to control for unobserved heterogeneity in the variance function too.⁹⁸

⁹⁶In fact, the modelling of the conditional mean process is based on the human capital theory (Becker, 1966) and the hypothesis of learning (Farber and Gibbons, 1997), where employers learn workers' productivity progressively. In the latter case income residuals follow a martingale.

⁹⁷A challenge for future research is undoubtedly to find a theoretical justification for this or alternative specifications. An autoregressive process can somehow be rationalized by the idea of bayesian probability updating.

⁹⁸Up to my knowledge, Banks, Blundell and Brugiavini (1998) is the only attempt to estimate the conditional variance of income shocks according to an *ARCH* process. However, the authors

To start with, let us consider the implications of model A. An *ARCH*(1) structure applied to transitory income innovations writes:⁹⁹

$$E_{t-1}(\varepsilon_{it}^2) = \kappa_t^e + \gamma^e \varepsilon_{it-1}^2 + \lambda_i \quad (13)$$

where the parameters are education-varying, κ_t^e is a year effect, and λ_i captures unobservable individual heterogeneity in the variance function. If γ^e is zero the variance function does not exhibit state dependence; in this case, persistence in the earnings distribution can only be explained by persistence in mean (or by persistence in conditional higher moments that I do not model explicitly). Nevertheless, heteroscedasticity might still be an issue because of the role played by individual heterogeneity and aggregate effects. Note that I do not impose non-negativity constraints on the parameters of the *ARCH* process to bound the conditional variance away from negative values. Of course, in the absence of such constraints, *GMM* estimation does not necessarily deliver a positive estimate of the risk term, as it should. While inconvenient, this strategy provides an informal way to check the plausibility of my specification. The parameters of the *ARCH* process can be identified *via* the moment condition:

$$E_{t-2}(\Delta \varepsilon_{it}^2 - \Delta \kappa_t^e - \gamma^e \Delta \varepsilon_{it-1}^2) = 0 \quad (14)$$

Let us turn to model B. Suppose for simplicity that transitory shocks are serially uncorrelated. I assume that the transitory shock follows the *ARCH*(1) process given by equation (13). I show below how to identify the parameters of the variance function under fairly general conditions. First, note that using the conditional expectation equivalent of (9) and the assumption of no serial correlation in the transitory component yields:

$$E_{t-1}(v_{it+1}v_{it}) = -E_{t-1}(\varepsilon_{it}^2) \quad (15)$$

Taking the first lag of the variance function (13), using (15), and then applying the law of iterated expectations, yields:

do not allow for the distinction between transitory and permanent disturbances.

⁹⁹For notational simplicity, we define: $E(\cdot | \Omega_{t-j}) = E_{t-j}(\cdot)$.

$$E_{t-2} (v_{it+1}v_{it} + \kappa_t^e - \gamma^e v_{it}v_{it-1} + \lambda_i) = 0 \quad (16)$$

Finally, taking first differences to eliminate the unobservable heterogeneity component, obtains the orthogonality condition:

$$E_{t-3} (\Delta v_{it+1}v_{it} + \Delta \kappa_t^e - \gamma^e \Delta v_{it}v_{it-1}) = 0 \quad (17)$$

that can be used to identify the parameters of interest. It is worth noting that if there is no unobservable heterogeneity, then first differencing is not required. The relevant orthogonality condition will hold in levels (equation 16) rather than in first differences. This suggests that one might test for the presence of unobservable heterogeneity in the variance function by comparing estimates in first differences with estimates in levels. Under the null of no individual effects, they should coincide, although those in levels should be more efficient than those in first differences.¹⁰⁰

The variance of the permanent shock, if present (model B), is assumed to follow an *ARCH*(1) process with unobservable heterogeneity of the form:

$$E_{t-1} (\zeta_{it}^2) = \phi_t^e + \theta^e \zeta_{it-1}^2 + \eta_i \quad (18)$$

To identify the parameters of the variance function (18) I make use of the fact that (under the null of no serial correlation in the transitory shock), from the conditional expectation equivalent of (12), I have that:

$$E_{t-2} (v_{it}^2 + v_{it+1}v_{it} + v_{it}v_{it-1}) = E_{t-2} (\zeta_{it}^2) \quad (19)$$

It is then easy to prove that the relevant orthogonality condition for this problem is:

$$E_{t-4} \left[\begin{array}{l} \Delta (v_{it}^2 + v_{it+1}v_{it} + v_{it}v_{it-1}) - \Delta \phi_t^e \\ -\theta^e \Delta (v_{it-1}^2 + v_{it}v_{it-1} + v_{it-1}v_{it-2}) \end{array} \right] = 0 \quad (20)$$

¹⁰⁰An interesting point to be made is that the stochastic structure of the orthogonality condition (17) remains unchanged once I allow for a permanent shock.

The flexibility I impose on the modelling structure of the earnings process under model B comes to a cost. This refers to data requirement. Equations (17) and (20) require the availability of instruments (in levels or first differences) dated $t-3$ and $t-4$, respectively. Since the composite error term v_{it} depends on the autoregressive properties of the mean income process, data requirement will become increasingly stringent as higher is the order of the autoregressive process. For instance, if $A^e(L, p)$ is of the second order, the construction of a consistent estimate of v_{it} will require the availability of four years of consecutive individual data, and estimation of (17) would require at least seven years of consecutive individual data. However, the model does not require a balanced panel data set.

2.5 The validity of the instruments

Recently, some concern has been raised regarding the bias of instrumental variables estimators when instruments are only weakly correlated with the right-hand-side variable (see Bekker, 1994, and Staiger and Stock, 1995). If this is the case, instrumental variables estimates will be biased towards OLS. In practice, one can check this issue out by evaluating the power of the excluded instruments in the reduced form, or comparing instrumental variables estimates with simple OLS estimates. This is after accounting for the time effects and age, i.e. for the variables included in the conditional equations. One also needs to evaluate the validity of the over-identifying restrictions; this can be done by computing the relevant Sargan statistic.

The conditions for identification of the parameters of the *ARCH* process under model B need to be examined more carefully. As said above, the difficulties arise from the fact that under model B the permanent and the transitory shocks are not directly observable, so that the estimation of the parameters has to rely on the autocovariance properties of the residuals from (5).

Consider first the identification of γ^e in (17). If $0 < |\gamma^e| < 1$, values of $v_{it-3}v_{it-4}$ will be correlated with the endogenous regressor, $\Delta v_{it}v_{it-1}$, and uncorrelated with the error term, thus allowing identification. Identification becomes problematic at two points in the parameter space: $\gamma^e = 0$ and $\gamma^e = 1$.

The former is a problem exclusively because I do not observe directly ε_{it}^2 , which forces me to use information lagged by one extra period. Lacking information on the true squared residuals, inference ought to be based on $\Delta v_{it+1}v_{it}$. In the case $\gamma^e = 0$:

$$E(v_{it-3}v_{it-4}\Delta v_{it}v_{it-1}) = 0 \quad (21)$$

The model is clearly unidentified from moment conditions (17) because one of the two identification conditions required for instrumental variables estimation is violated. Consider now the case when $\gamma^e = 1$. Using the same reasoning, it is possible to show that the lagged dependent variable in (17) may result uncorrelated with *all* past information, again leading to lack of identification.

These remarks have two important implications. First, it is not possible to test the null hypothesis that $\gamma^e = 0$ or that $\gamma^e = 1$ using the estimates obtained by solving the moment conditions in (17). This is because these moment conditions do not identify γ^e under either null. Second, it is important to evaluate the validity of the instruments by examining the correlation between the endogenous regressor and the elected instrumental variables from the available data. This is equivalent to examining the rank condition for identification. In this simple model, if I accept the null hypothesis that $E(v_{it-3}v_{it-4}\Delta v_{it}v_{it-1}) = 0$ based on the result of a Wald-test from the reduced form, then I need to distinguish between the two possible parameter configurations: $\gamma^e = 0$ or $\gamma^e = 1$. This can be done by examining the serial correlation properties of $\Delta v_{it+1}v_{it}$. Under the null hypothesis that $\gamma^e = 1$, $\Delta v_{it+1}v_{it}$ should be serially uncorrelated. If this is rejected, an $MA(1)$ error structure would imply $\gamma^e = 0$. The null hypothesis of $MA(1)$ should then be accepted.

A similar discussion applies to the identification of the $ARCH(1)$ coefficient in the variance function of the permanent shock. I omit it because it follows exactly the same line of reasoning.

2.6 Asymmetry and the lack of identification

An interesting possibility allowed in $ARCH$ models for time-series data is

that of asymmetry of response to shocks. In other words, the conditional variance function is allowed to respond asymmetrically to positive and negative past shocks. This could be interesting here as well, for a considerable amount of asymmetry in the distribution of earnings is related to unemployment. Caballero (1990) shows that asymmetric distributions enhance the need for precautionary savings. One model embedding the notion of asymmetry has been advocated by Sentana (1995), and is here reproduced as:

$$E_{t-1} \left(\varepsilon_{it}^2 \right) = \kappa_t^e + \gamma^e \varepsilon_{it-1}^2 + \rho^e \varepsilon_{it-1} + \lambda_i$$

It is easy to show that the moment condition that identifies the parameter of this model is:

$$E_{t-2} \left(\Delta \varepsilon_{it}^2 - \Delta \kappa_t^e - \gamma^e \Delta \varepsilon_{it-1}^2 + \rho^e \varepsilon_{it-2} \right) = 0 \quad (22)$$

Note that the latter is identifiable under model A, but not under model B. The reason is that I directly observe (a consistent estimate of) the shock only in the former case. Under model B transitory and permanent income shocks are not separately observable; since my identification strategy consists of using repeatedly the law of iterated expectations, the parameter ρ^e remains unidentified.

3 The data

The data are drawn from the 1967-1991 family and individual-merged files of the PSID (waves I through XXV). See chapter 1 for more details about the survey and Hill (1982).

The PSID started in 1968 collecting information on a sample of roughly 5,000 households. Of these, about 3,000 were representative of the US population as a whole (the core sample), and about 2,000 were low-income families (the Census Bureau's SEO sample). In the empirical analysis I use both the core sample and the SEO sample. It is fair to say that few authors (Lillard and Willis, 1978) have suggested not to use the SEO low-income sample because of endogenous selection. In other words, an initial condition problem arises for those in poverty at the

beginning of the survey period. However, given linearity, the initial condition problem is taken care of by the presence of the permanent component. To put it differently, I deal with the problem by estimating models for the growth rate rather than specifications in levels.

The selection stage aims at obtaining a sample of individuals with at least nine years of usable earnings data on which to perform the estimation. The final sample includes male heads aged 25 to 55 in all years who reported positive earnings data for at least nine consecutive years.¹⁰¹ Individuals with top-coded earnings in at least one year were discarded; I also dropped those who were self-employed in at least a year. My final sample slightly differs from those of Lillard and Willis (1978), MaCurdy (1982), Abowd and Card (1989), Moffitt and Gottschalk (1994), and Baker (1997), the representative papers of the literature on earnings dynamics. A common difference with all studies is that I have available a longer time-series of individual observations. Abowd and Card and Moffitt and Gottshalk are the only studies where both the core sample and the SEO sample are used. With the exception of Abowd and Card, all studies focus on white males only. MaCurdy restricts his sample to individuals with little year-to-year variability in wages or hours. Finally, Baker works with a balanced panel of 534 individuals continuously surveyed from 1967 to 1986.

Questions referring to labour income are retrospective; thus, those asked in 1968, say, refer to the 1967 calendar year. The earnings variable¹⁰² is the labour portion of money income from all sources; the variable name in the PSID tapes is “head’s money income from labour” and includes the labour part of farm income and business income, wages, bonuses, overtime, commissions, professional practice, labour part of income from roomers and boarders or business income.¹⁰³

¹⁰¹Very few individuals had one or more years of missing data on earnings between two series of usable earnings data. I decided to keep these observations but treat them as belonging to two different individuals.

¹⁰²I deflate the nominal measure of earnings by the GNP personal consumption expenditure deflator (using 1991 as the base year).

¹⁰³As noted by Gottshalk and Moffitt (1993), the measure of labour income available in the PSID has sources that may reflect capital income, such as the labour part of farm income and roomers and boarders. I do not account for this problem.

Education level is computed using the PSID variable with the same name. Since I perform my empirical analysis for different education groups, individuals with inconsistent report on education are discarded. I consider three education groups: high school dropout (those with less than 12 grades of schooling), high school graduate (those with at least a high school diploma, but no college degree), and college graduate (those with college degree or more).

The selection process outlined above leads to a sample of 2,093 individuals and 31,659 individual-year observations. Relevant sample statistics are presented in tables 1A and 1B.¹⁰⁴

It is worth examining the resulting PSID data on earnings. Figures 3 and 4 plot the median and the interquartile range of log earnings for the three education groups I focus on against time. The dispersion in log earnings data displays some clear business cycle effects. Median log earnings are slightly compressed only in the early 1970s; after that period, the gap between education groups widens substantially. This can also be seen from the dynamics of the interquartile range displayed in figure 4. Two facts are worth noting. First, dispersion in log earnings increases for both the high school dropout and the high school graduate, but the pattern for the most educated is less regular, with a decline in the early years, an increase from the mid-1970s to the mid-1980s, and a further slow decline afterwards. Second, as an effect of these diverging patterns, at the end of the survey period there is little or no difference between college graduates and high school graduates as far as earnings dispersion is concerned, while the least educated still face most of the earnings variability we observe at cross-sectional level.

4 Empirical results

4.1 Model A

I firstly estimate the conditional mean process under the simple model A. As

¹⁰⁴For brevity, step-to-step details about sample selection are omitted.

said in section 2.2, I start by filtering out deterministic life-cycle and business-cycle effects from annual earnings. This is accomplished by regressing log real earnings on age, age squared, and a full set of year dummies, and saving the residuals. Regressions are run separately for the three education groups I focus on. From the saved residuals I then estimate $A^e(L, p)$, imposing (arbitrarily) the restriction $p = 1$.¹⁰⁵ Note however that the Sargan test does not reject such restriction. The estimates of the autoregressive coefficients and of the variance of the transitory component are presented in Table 2, separately for each education group.¹⁰⁶ There are two results deserving comment. First, persistence in earnings varies with education, being in general higher for those with more years of schooling. In particular, the same transitory shock has an impact of 1.24 for the least educated, 1.41 for the high-school dropout, and 1.39 for the most educated.¹⁰⁷ Second, the variance of the transitory shock decreases with education, a finding common to Gottschalk and Moffitt (1993) and others.¹⁰⁸

Next, I study the dynamics of the conditional variance by estimating the *ARCH* process (14). I estimate the specification (14) by *GMM* using as instruments lags of the dependent variable dated $t - 2$ and $t - 3$ (see section 4.6 for more details on the estimation procedure). All regressions in level include year dummies and a quadratic in age to capture fluctuations in risk related to the business cycle and ageing. The results are presented in Table 3. While the estimates support the hypothesis of the particular form of conditional heteroscedasticity implied by an *ARCH* process, persistence is not extremely high, and moreover appears quite similar across education groups. The test statistics in general reject neither the specification chosen nor the timing of the instru-

¹⁰⁵These are GMM estimates. I obtained them instrumenting the lagged dependent variable Δy_{it-1} with lags dated $t - 2$ and $t - 3$. See also section 4.2 for more details on the estimation procedure.

¹⁰⁶The variance of the transitory component is obtained using the formula reported in Gottschalk and Moffitt (1994, p. 254), i.e.: $\sigma^2 = N^{-1} \sum_{i=1}^N (T_i - 1) \sum_{t=1}^{T_i} (\tilde{y}_{it} - \tilde{p}_i)^2$, where \tilde{y}_{it} and \tilde{p}_i have been defined in section 2.2. In practice, one averages N individual variances to obtain a single measure over all years and all individuals.

¹⁰⁷This is calculated as $(1 - \text{first lag})^{-1}$.

¹⁰⁸The variance in the whole sample is 0.1042, with a standard error of 0.0014.

ments. Interestingly, risk appears to decline with age, a finding that matches casual observation and intuition.

We also estimate a quadratic *ARCH* model to allow for asymmetry (see section 2.6). I find that once the level of the transitory shock is controlled for, *ARCH* coefficients decline in absolute value and statistical significance. The sign of the coefficient on ε_{it-1} is consistent with the view that bad shocks (for instance, a brief spell of unemployment) generate more uncertainty than good shocks (such as a bonus); this also implies that the conditional variance function is steeper for negative shocks. However, in terms of statistical inference the model is badly rejected according to the Sargan test at least in one case. I thus prefer not to put too much emphasis on such evidence.

Consider now the possibility of using the numbers implied by my empirical exercise in models of general or partial equilibrium. Consider the simplest case of a model that calculates consumption profiles based on numerical solutions, as in Zeldes (1989b), for instance. In such models, one uses the terminal condition and the Euler equation to solve backward induction for the consumption function at all dates. Consumption profiles are then numerically calculated starting from the first period: a value for the income shock is drawn from a distribution with mean zero and variance σ^2 (for instance, the variance of the income innovation as in MaCurdy, 1982). This determines consumption, saving and the assets carried forward to the second period. In the second period, a further draw from the stationary distribution of income shock is taken, and this determines again consumption, saving and assets for that period. The process continues until the last period. Yet, such procedure is restrictive. To account for income risk dynamics and heterogeneity, one could draw the value for the income shock from a distribution that has zero mean but variance determined by the *ARCH* expression: $E_{t-1}(\varepsilon_{it}^2) = h(\text{age}_{it}) + 0.05 \cdot \varepsilon_{it-1}^2$, say, where ε_{it-1} is the value obtained from the previous draw and $h(\text{age}_{it})$ is a function of age that can be estimated from the data. Clearly, some form of discretisation of the space of the shock must be adopted to make the solution computable. While this strategy makes the derivation of the consumption function much more complicated, it may

provide a substantial gain in terms of realism. For the first period, the shock can be drawn from the unconditional distribution. I am not aware of methods of solution based on this simple markovian rule for the conditional variance of income shocks.

While suggestive, the results reported in this section should be taken with caution. The simple model A does not provide a very accurate description of the earnings process. In fact, as I shall see, it is rejected according to the test presented in section 2.2.

4.2 Model B

4.2.1 The conditional mean

From now on, my attention will be focused on model B. The estimates of the conditional mean process for log earnings are presented in table 4, separately for the high school dropout, the high school graduate, and the college graduate. I impose an *a priori* restriction about the order of the lag polynomial $A^e(L, p)$, transform the data to eliminate serial correlation due to the presence of a martingale in levels, and assume that transitory shocks are serially uncorrelated (i.e., $q = 0$). In particular, I set $p = 1$. The estimation strategy follows Arellano and Bond (1991a). I explore model misspecification using a number of conventional tests.

In performing my *GMM* estimates I use only the second and the third lag of log earnings (in levels) as instruments. In particular, I estimate separately for each education group the following specification:

$$\Delta y_{it} = \Delta m_t + \beta_m \text{age}_{it} + \alpha \Delta y_{it-1} + v_{it} \quad (23)$$

for all $i = 1, 2, \dots, n$.¹⁰⁹ Neglecting for simplicity of notation time effects and age, the relevant $(T_i - 3)$ equations for individual i can be written more compactly as:

¹⁰⁹The equation in levels would imply a quadratic polynomial in age, as is standard in the literature and in our model A.

$$\Delta \mathbf{Y}_i = \alpha \Delta \mathbf{Y}_{i,lag} + \mathbf{V}_i \quad (24)$$

where $\Delta \mathbf{Y}_{i,lag}$ is a vector containing the lagged endogenous variable in first-difference. The two-step *GMM* estimator of α is:

$$\hat{\alpha}_{2S} = \left[\left(\sum_i \Delta \mathbf{Y}'_{i,lag} \mathbf{Z}_i \right) \mathbf{A} \left(\sum_i \mathbf{Z}'_i \Delta \mathbf{Y}_{i,lag} \right) \right]^{-1} \left(\sum_i \Delta \mathbf{Y}'_{i,lag} \mathbf{Z}_i \right) \mathbf{A} \left(\sum_i \mathbf{Z}'_i \Delta \mathbf{Y}_i \right) \quad (25)$$

where \mathbf{Z}_i is a matrix of instruments, and $\mathbf{A} = n^{-1} (\sum_i \mathbf{Z}'_i \mathbf{J}_i \mathbf{Z}_i)^{-1}$. In the first step of the estimation, the matrix \mathbf{J}_i weights the covariance matrix of the instruments by a square matrix that has twos in the main diagonal, minus ones in the first subdiagonals, and zeroes otherwise (to mirror the *MA*(1) structure with constant variance of the disturbance in first difference); in the second step, the weighting matrix is defined as: $\mathbf{J}_i = \hat{\mathbf{V}}_i \hat{\mathbf{V}}_i^T$, with $\hat{\mathbf{V}}_i$ being the one-step residuals (i.e., $\hat{\mathbf{V}}_i = \Delta \mathbf{Y}_i - \hat{\alpha}_{1S} \Delta \mathbf{Y}_{i,lag}$). Finally, in my specific case, the instrument matrix \mathbf{Z}_i is to be written as:

$$\mathbf{Z}_i = \begin{bmatrix} y_{i1} & y_{i2} & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & y_{i2} & y_{i3} & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & 0 & \dots & y_{iT_i-3} & y_{iT_i-2} \end{bmatrix} \quad (26)$$

This is not the optimal set of instruments as suggested, among others, by Arellano and Bond (1991a), because it does not exploit all the available linear orthogonality conditions. In fact, I truncate the set of available instruments at the third lag. I decided to do so because using fewer instruments should reduce the bias that arises in the optimal *GMM* estimator when all the available orthogonality conditions are exploited in the estimation (Ziliak, 1997); moreover, it should reduce the bias of the Sargan test towards the null hypothesis.¹¹⁰

¹¹⁰This bias arises from inflating the degrees of freedom of the test with the inclusion of irrelevant instruments.

The results from table 4 suggest, as in the simple model A, that the persistence in the mean process increases with schooling; once permanent effects have been allowed for, the total impact of any transitory shocks on annual earnings is roughly 1.36 for both the high school graduate and the college graduate, but it is only 1.21 for the high school dropout.¹¹¹ These numbers imply that the data are best described by low persistence of transitory shocks and permanent unobserved skill characteristics (see the results of my test below). OLS estimates in levels are clearly upward biased and suggest much more state dependence than those in first differences (and thus false evidence for much more rapid regression towards the mean).¹¹² The null hypothesis that the transitory shock is not serially correlated (over and above the amount of serial correlation guaranteed by the polynomial $A^e(L, 1)$) is tested by considering the value of the Sargan statistic, i.e. the test for the over-identifying restrictions. The one I present is robust to the presence of conditional heteroscedasticity. The p -value of the Sargan test is generally supporting the null of no serial correlation in the transitory shock. This conclusion is also supported by the results of the test of second-order serial correlation for the residuals of the equation in first difference (the m_2 statistic, see Arellano and Bond, 1991a). The null hypothesis can never be rejected at conventional levels of statistical significance.

To evaluate the power of my instruments (see section 2.4), I estimate the specification (4.1) by OLS (results not reported). I find that *GMM* estimates are not biased towards OLS estimates (a formal difference test would reject the null of no difference). I also compute the p -value of the Wald-statistic for the null hypothesis that the excluded instruments in the reduced form are jointly insignificant and the adjusted R^2 for the regression of the endogenous explanatory

¹¹¹These results slightly differ from those of model A because in this case age and time effects are estimated jointly with the autoregressive coefficient. In model A, instead, age and time effects are filtered out before estimating the autoregressive coefficient.

¹¹²MaCurdy (1982) considers an *ARMA*(1,1) process for the transitory component and finds the estimate of the autoregressive coefficient to be 0.22, a value that sits amidst my estimates of 0.17, 0.27 or 0.26 (my estimates for the three education groups considered). His preferred estimate for the variance of the transitory shock is 0.056, while mine is slightly higher (0.068).

variable on the excluded instruments. Here again, it is clear from the results reported in table 4 that my instruments are sufficiently powerful to rule out the concern of poor first-stage correlation.

4.2.2 The structure of the error term

I test for the presence of a permanent shock using the test statistic outlined in section 2.2. I consider the null hypothesis given by equation (11); the latter is valid in the absence of serial correlation in the transitory shock, an assumption that is not rejected by the evidence presented in the previous section. The null hypothesis of no permanent shock is strongly rejected: the value of the test on the whole sample (23,287 individual-year observations) is 5.17, implying a p -value well below 1 percent.¹¹³ Thus, I impose the presence of a martingale permanent component in the mean income process and reject the simple specifications embedded in models A.

I use the time series of equations (9) and (12) to measure the unconditional variance of the transitory and the permanent shock, respectively.¹¹⁴ I estimate the average variance of the permanent shock to be 0.0117, while the average variance of the transitory shock is estimated to be 0.0685. Table 5 presents the estimates for the whole sample and for the three education groups separately. These numbers are very close to those found elsewhere in the empirical literature (see Carroll and Samwick, 1997), although my focus is on earnings rather than family income. The estimated variance of the transitory shock is undoubtedly very high, highlighting the importance of measurement errors. Unfortunately, in my model it is impossible to distinguish between measurement errors and genuine transitory innovations.

Note that once permanent shocks are taken into account, the results of the simple model A demonstrate how different earnings specification may lead to very different conclusions concerning the magnitude of individual risk. With respect

¹¹³The p -values of the test conducted separately for each education group are 0.0010, 0.0002, and 0.0133, respectively for the high school dropout (6,105 observations), the high school graduate (11,910), and the college graduate (5,272).

¹¹⁴These variances are estimated as the average, over all individuals and years, of $-(\widehat{v}_{it+1}\widehat{v}_{it})$ and $(\widehat{v}_{it}^2 + \widehat{v}_{it+1}\widehat{v}_{it} + \widehat{v}_{it}\widehat{v}_{it-1})$, respectively.

to model A, in fact, the variance of the transitory shock declines substantially, implying that such model would over-estimate the variance of the transitory shock. The reason is that in model A I assume that the permanent component is fixed over time. In practice, there I pretend to be estimating transitory fluctuations in earnings from the specification $\tilde{y}_{it} = \tilde{p}_i + \varepsilon_{it}$. In reality, the presence of a martingale permanent component implies that $\tilde{y}_{it} = \tilde{p}_i + \sum_{j=1}^t \zeta_{ij} + \varepsilon_{it}$, where \tilde{p}_i is the starting value of the martingale process. Thus, in model A I assume that the variance of the term $(\tilde{y}_{it} - \tilde{p}_i)$ is estimating the variance of the transitory shock, while in fact it is estimating the variance of $(\sum_{j=1}^t \zeta_{ij} + \varepsilon_{it})$. In practice, the variance of the transitory shock in model A is increased by the cumulative sum of the variances of permanent shocks that are falsely assumed to be absent. While interesting in their own right, these results are not the focus of this study. As remarked in the introduction, I am interested in a different aspect of income variability, that related to the inherent dynamics of risk.

I conclude this section by presenting less structural evidence on the relationship between the variance of income shocks, the business cycle and the life cycle. To this aim, in figures 5 and 6 I plot the unconditional variances against the U.S. male unemployment rate.¹¹⁵ While I do not emphasize this evidence too much, it is of some interest to note that the unemployment rate and the variance of transitory shock series appear to be quite synchronized, especially in the 1980s. This confirms the intuitive claim that uncertainty rises during economic slumps (a finding common to Gottschalk and Moffitt, 1993, and Storesletten, Telmer and Yaron, 1998). Likewise, there seems to be a strong association between the variance of permanent shocks and the rate of male unemployment (recall that both variances have been constructed from income residuals that take already into account economy-wide effects, which means that the counter-cyclicity of variances seems genuine rather than spurious). This counter-cyclicity has been advocated by those who propose a resolution of the equity premium puzzle based on

¹¹⁵A regression of the variance of permanent shocks on male unemployment rate displays a point estimate of 0.46 with a standard error of 0.14; for the transitory shock, the estimated coefficient is 0.82 with a standard error of 0.32.

the negative correlation between aggregate shocks and individual risk (Mankiw, 1982).

In figure 7 I plot the variances of earnings shocks against age. It is interesting to note that both variances appear to decline over the life-cycle; yet, while transitory volatility declines linearly, permanent uncertainty displays a distinctive U-shape. This suggests that permanent income shocks tend to become progressively more important towards the end of the working career, perhaps because facing an involuntary dismissal at 55, say, has more dramatic consequences than facing the same event at 25. In the latter case young workers may choose to re-train or return to school, while for more mature workers the only available option is perhaps early retirement. Similarly, permanent shocks are fairly important at the beginning of the life cycle because of the possibility of mismatch between individual skills and the elected job.

4.2.3 The identification of the ARCH coefficients

As outlined in section 2.4, it is important to check carefully the conditions for the identification of the *ARCH* coefficients in (17) and (20). I shall use two items to perform this exercise: the set of reduced form statistics and the set of autocovariances for the first-differenced squared residuals. A signal for the lack of identification of the model is poor correlation between lagged instruments and the endogenous explanatory variable. To examine this issue, I compute the *p*-value of the Wald-test for the significance of the excluded instruments in the reduced form and the adjusted R^2 of the regression of the endogenous explanatory variable on the excluded instruments (dated $t - 3$ and $t - 4$ for the transitory shock and $t - 4$ and $t - 5$ for the permanent shock; all instruments are in levels). Results are reported in table 6.

I first comment on the results concerning the identification of the *ARCH* coefficient for the transitory shock. With the exception of the high-school dropout (on which I comment later), the *ARCH* coefficient for the variance of the transitory shock appears to be unidentifiable from the moment condition (17). The *p*-values for the excluded instruments in the reduced form are extremely high, and well above standard levels of acceptance. In all cases, the adjusted R^2 is low

or even negative. If one were to estimate the parameters of the *ARCH* specification (13) based on the moment condition (17), statistical inference would be undoubtedly questionable.

Recall from the discussion in section 2.4 that failure to pass the Wald-test in the reduced form implies that the *ARCH* coefficient should lie in the parameter space $\{0, 1\}$. As illustrated in section 2.4, a possible way to discriminate between these two hypotheses is to test whether $\Delta v_{it+1}v_{it}$ is serially correlated. If this is the case, the hypothesis $\gamma^e = 0$ should be accepted, as $\gamma^e = 1$ rules out any kind of serial correlation. In the table 6 I report estimates of $E[(\Delta v_{it+1}v_{it})(\Delta v_{it-j+1}v_{it-j})]$.¹¹⁶ To conserve space, I choose to report results for $1 \leq j \leq 5$ only. It is evident that $\Delta v_{it+1}v_{it}$ is indeed serially correlated. However, as the null hypothesis $\gamma^e = 0$ would imply, the order of the serial correlation is no higher than one. Further lags are sufficiently small and not significantly different from zero. This evidence should be seen as complementary to the results on the rank condition for identification presented above. Based on such results, I believe that for all the education groups there is very weak evidence for the existence of dynamics in the process of the conditional variance of transitory shocks. For the high-school graduates and the college graduates the evidence is confirmed by both reduced form statistics and the pattern of autocovariances; for the least educated, the marginally favourable evidence from the set of first-stage statistics is attenuated by the pattern of autocovariances. At the second order, the autocovariance literally collapses to a steady-state value of zero, a fact that is strongly against the pattern of correlation that an autoregressive process would produce.

Next, I turn to the identification of the *ARCH* coefficient of the permanent shock. The statistics of the reduced form are in support of the hypothesis of identifiability of the *ARCH* coefficient based on the elected instruments. However, I find a somewhat high p -value (17 percent) for the group of high-school graduates. The adjusted R^2 s are not extremely high, but this is not surpris-

¹¹⁶I assume covariance stationarity within education. A test for covariance stationarity has p -values of 0.19, 0.28, and 0.27, respectively. Similar results hold for higher-order autocovariances.

ing for first-difference specifications. The pattern of autocovariances displays an oscillatory pattern around zero. Autocovariance are sufficiently high and statistically significant up to third order.¹¹⁷ Based on such evidence, I may conclude that $0 < |\theta^e| < 1$ for all education groups, with the exception of the high school graduate. For them, in fact, the unfavourable evidence portrayed by reduced form statistics is not attenuated by the pattern of autocovariances. Such pattern is also inconsistent with the hypothesis that $\theta^e = 1$ and an individual effect is present. If this was the case, autocovariances should be of similar magnitude and display the same sign at all lags, a fact that is not consistent with the evidence reported in table 6. Thus, I conclude that -at least for the groups at the two extremes of the distribution of schooling- the *ARCH* process for the variance of the permanent shock is identifiable from moment conditions (20), while for the high-school graduate there is no evidence for ARCH (i.e., idiosyncratic) dynamics in the conditional variance of the permanent shock. In what follows, I shall present the results of estimating the *ARCH* process only for the permanent shock and commenting on the effect of time and of observable and unobservable heterogeneity for the conditional variance of the transitory shock.

4.2.4 The conditional variance of the transitory shock

Recall from (13) that I posit a model for the conditional variance of the transitory shock where the latter depends on time, the square of past innovation, and individual heterogeneity. As showed in the previous section, there seems to be very weak evidence for the presence of dynamics in the conditional variance (apart from the possibility of aggregate dynamics). It is still interesting to assess whether there are any time or individual effects at all. To this aim, I estimate individual-specific intercept regressions (by education group) for the variable $v_{it+1}v_{it}$; I condition the latter on a full set of year dummies, a quadratic in age and a full set of person dummies. This is equivalent to applying a standard

¹¹⁷Also in this case, I assume that unconditional autocovariances are constant over time within the three education groups. A test of equal year effects on the first-order autocovariance has *p*-values of 0.48, 0.06, and 0.42, respectively. *P*-values are somewhat lower (but still above 10 percent) for higher-order autocovariances.

within-group procedure.

The results of these regressions are reported in table 7. I report the coefficients on age and age squared and the p -value of the Wald statistic for the null hypothesis that year dummies and individual dummies are jointly insignificant, respectively. Observable heterogeneity (as measured by an age polynomial) is often poorly measured, while time effects are particularly relevant for those without a college degree (i.e., for those who are perhaps more likely to be exposed to aggregate fluctuations, see my interpretation below). As for unobservable heterogeneity, I find that it matters for all the education groups I consider: the p -value for the null hypothesis of no heterogeneity is in fact consistently below 2 percent in all cases.

The evidence presented in this section is thus in favour of a model where transitory volatility in income is mainly explained by unobserved heterogeneity (this also implies cross-sectional differences in risk). Deviations around idiosyncratic trends are due to aggregate fluctuations; due to the evidence presented in section 5.2.2, transitory volatility in earnings increases during economic recessions and declines during booms. Yet, given that time dummies are not jointly significant for the most educated, this pattern is typical of individuals with little education.

In my view, unemployment risk (e.g., entry and exit from employment) is the leading explanation for such results, especially in the light of the differences across education I have illustrated above. Because of training on-the-job and specific human capital investments, it is perhaps more expensive for firms to fire (and then re-hire) highly educated individuals than individuals with low skills or qualifications.

4.2.5 The conditional variance of the permanent shock

The estimates of the conditional variance function for the permanent shock are based on the orthogonality condition (20). I estimate separately for each education group the following specification:

$$\Delta\psi_{it} = \Delta\phi_t + \beta_v age_{it} + \theta\Delta\psi_{it-1} + \omega_{it} \quad (27)$$

where $\psi_{it} = \hat{v}_{it}^2 + \hat{v}_{it+1}\hat{v}_{it} + \hat{v}_{it}\hat{v}_{it-1}$, and \hat{v} is the two-step residual from the

estimated mean income process in first difference. The equations for individual i can be written more compactly (ignoring again age and time effects) as:

$$\Delta\Psi_i = \theta\Delta\Psi_{i,lag} + \mathbf{W}_i \quad (28)$$

where $\Delta\Psi_i$ is the vector of $\Delta(\widehat{v}_{it}^2 + \widehat{v}_{it+1}\widehat{v}_{it} + \widehat{v}_{it}\widehat{v}_{it-1})$. The two-step *GMM* estimator of θ can be written as:

$$\widehat{\theta}_{2S} = \left[\left(\sum_i \Delta\Psi'_{i,lag} \mathbf{R}_i \right) \Theta \left(\sum_i \mathbf{R}'_i \Delta\Psi_{i,lag} \right) \right]^{-1} \left(\sum_i \Delta\Psi'_{i,lag} \mathbf{R}_i \right) \Theta \left(\sum_i \mathbf{R}'_i \Delta\Psi_i \right) \quad (29)$$

where $\Theta = n^{-1}(\sum_i \mathbf{R}'_i \mathbf{N}_i \mathbf{R}_i)^{-1}$. \mathbf{R}_i is the instrument matrix, and the first-step structure of \mathbf{N}_i mirrors that of \mathbf{J}_i above (see section 4.2); in the second step: $\mathbf{N}_i = \widehat{\mathbf{W}}_i \widehat{\mathbf{W}}'_i$, with $\widehat{\mathbf{W}}_i$ being the first-step residuals. Finally, the instrument matrix \mathbf{R}_i is:

$$\mathbf{R}_i = \begin{bmatrix} \psi_{i1} & 0 & 0 & \dots & 0 & 0 \\ 0 & \psi_{i1} & \psi_{i2} & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & \psi_{iT_i-6} & \psi_{iT_i-5} \end{bmatrix} \quad (30)$$

Results are presented in table 8 separately for the two education groups I focus on. I find the result that at least for the high-school dropout and the college graduate the variance of the permanent shock exhibits a statistically significant negative state dependence;¹¹⁸ this holds true for the specification in level as well as in first difference. However, the test of second-order correlation signals model misspecification when the level equation is used (for the high-school graduate the specification in first difference is accepted only marginally, at the 2 percent level). The latter is consistent with the presence of unobservable heterogeneity in levels, as my specification (19) would imply. Yet, the results of the Sargan test would not reject the null hypothesis of the specification in level; it is likely that

¹¹⁸First-step estimates are similar, although less precisely measured.

this test tends to over-accept the null even when a relatively small number of instruments is used, probably because of small sample size.

Time effects are significant. While the effect of age is consistent with the view of an age-declining risk pattern, it is often poorly measured. Finally, the OLS estimates (results not reported) confirm that the concern of a poor reduced form can be ruled out: OLS estimates and *GMM* estimates for the specification in first difference appear dramatically different (a formal difference test would reject the null of no systematic difference).

5 The implications of the empirical estimates

5.1 A general discussion

A tentative interpretation of the results presented in the previous section is now in order. At the cost of simplifying the analysis by a great deal, I decided to synthesize my results as follows. Based on the results from the reduced form of my *GMM* regressions, the conditional variance of the transitory shock is unlikely to exhibit state dependence; as for the conditional variance of the permanent shock, I find a result of negative state dependence, at least for the individuals at the two extremes of the distribution of schooling. Estimates for the high-school graduate, the largest group in my sample, are instead consistent with the view that permanent risk is independent on past shocks. Time effects are in general important for both conditional variances, while the effect of age is somewhat ambiguous and often poorly measured. Finally, unobservable heterogeneity seems to matter for both the permanent and the transitory uncertainty.

The result of negative state dependence implies that a permanent disturbance to income in a given time period (say, a promotion) leads to a reduction in the conditional variance of the permanent income shock in the following period. Before discussing the implications of this result more in detail, it is interesting to see whether is capable of capturing some of the features of permanent risk determination in the real world. It is an intuitive presumption that permanent

shocks occur at low frequencies. In what follows I will consider a trivial numerical example. Suppose a worker knows his own productivity on a particular job, but the employer is uncertain about it. However, the employer observes the worker's performance and updates her beliefs based on such observation. The employer is willing to hire skilled and unskilled workers for high-quality and low-quality tasks, respectively (i.e., the production function for this firm is $Y = f(U, S)$, where U is unskilled and S skilled work). Initially, all workers are deemed unskilled.

The initial wage paid to the worker when she enters the firm is w . After the first period of work, the employer can take three actions based on performance observation: (a) doing nothing (which reflects the fact that the employer is still trying to ascertain individual productivity, i.e. whether the worker is skilled, unskilled or unfit to the job), (b) fire the worker (which reflects the fact that the employer has ascertained that the worker has skills unfit to the job), or (c) promote the worker by paying her the long-run wage, which is $2w$ (this reflects the fact that the employer has ascertained the good quality of the worker, i.e. the possession of skills for high-quality tasks). After promotion, attachment to the firm is permanent.

Let us assume that after being fired, a worker can always find another job but that prospective employers ignore the fact that she was fired in the past (perhaps some skills are useful to a job and useless elsewhere). Again, lack of better information, the new employer pays w in the first period and follows the same updating rule described above afterwards. The three actions the employer can take initially attract equal probabilities (namely, $1/3$). Pay is received at the beginning of each period, and I assume that a worker fired at the beginning of the second period receives 0 and spends one period unemployed seeking work. The wage in the first period of work is certain. For simplicity, I assume that no worker quits voluntarily.

At the beginning of period one, the conditional mean of future income shock is 0, current income shock is 0 as well, while the conditional variance of future shocks is $0.66w^2$. At the beginning of period two, the conditional mean of income shock is 0 for those fired or promoted (although current shocks are now $-w$ and

w , respectively), and the conditional variance is 0 as well. For those who are still under testing, instead, I assume that the employer has updated her belief in the following way. She weights the probability that the worker is useless or skilled a bit less and the probability that the worker is unskilled, but still useful to the firm in low-quality tasks, a bit more than in the previous period. For instance: the probability of being fired or promoted tomorrow equals 1/4 and the probability of remaining at the same wage equals 1/2. The conditional mean of future income shock is still 0 (as is the current shock), but now the conditional variance is $0.5w^2$.

Over time, information become more complete. For instance, for those remaining employed at a wage of w , the probability of being fired or promoted declines over time (I assume that it does so symmetrically), while the probability of being fit for low-quality tasks increases. Therefore, I may have various different combinations of past shock/conditional variance. If I were to collect data on this particular labour market, I would obtain, for instance:

$$\begin{array}{cccccc} u_{t-1} & -w & w & 0 & 0 & 0 & \dots \\ E_{t-1}(u_t^2) & 0 & 0 & 0.66w^2 & 0.5w^2 & 0.4w^2 & \dots \end{array}$$

A regression of the conditional variance of u_{it} on past shock squared (with all the obvious cautions this exercise implies) would provide a negative *ARCH* coefficient, as is easy to check assigning arbitrary values to w .

My empirical estimates suggest that –given symmetry– rational agents make the quite plausible assumption that the probability of receiving a major permanent shock, of either sign, falls with the magnitude of the earnings shock experienced in the more recent past. In other words, the conditional density function must change to accommodate a decline in the conditional variance brought about by the realization of today’s permanent shocks. For instance, outside my trivial numerical example, it might be quite unlikely for a worker to be promoted every year (perhaps because employers face promotion costs), or to face a firm closure event for two consecutive years (unless bad luck is a very persistent process). The result of negative state dependence in the variance function can then capture the theoretical notion of uncertainty resolution over a worker’s labour market career.

While this result could be somehow intuitively appealing, when taken to the data it does not guarantee that the estimates will be consistent with the very obvious fact that the variance must be non-negative with certainty. Since the contribution of past shocks to the variance is negative, the only way to ensure positive conditional variances is to let the unobserved heterogeneity have a strong balancing positive effect on the variance function, or else my identification strategy is bound to work only for small permanent shocks.¹¹⁹ Therefore, the role of unobserved heterogeneity in the variance function should be paramount.

Alternatively, it is perhaps possible that the linear *ARCH* structure I have adopted is just a local approximation for a highly non-linear specification, where risk depends on features of the distribution of earnings that I ignore or am unable to characterize. In fact, lacking a structural model for the evolution of the conditional variance, a linear *ARCH* process is just one of the possible specification one could use to measure stochastic risk. In other words, I cannot rule out the possibility that my results could be interpreted as signalling stochastic risk model misspecification rather than as portraying genuine features of, say, the functioning of the labour market or pay determination.

5.2 Precautionary savings and insurance

As said in the introduction, papers in the precautionary savings literature assume a relatively simple process for income. A very common assumption is that of i.i.d. income shocks; this is somehow inconvenient, as uncertainty is likely to be itself stochastically determined. As suggested in section 4.1, one possibility would be to estimate a stochastic process for risk (as my *ARCH*) and examine the optimal response to various realizations in models that calibrate the behaviour of optimising consumers. In the remainder of this section, I illustrate the role that state dependence in the conditional variance function may have in a formal model of consumption choice. I stress that the analysis here is conducted on a very informal way, and it is not meant to be exhaustive. It is likely that only numerical simulations may shed clear light on the interaction

¹¹⁹As a partial justification, only individuals with low unemployment are in my sample.

between consumption choice and income risk dynamics and heterogeneity.

Consider Caballero's model (1990) with exponential utility, $u(c) = -(1/\theta)e^{-\theta c}$, a certain and constant interest rate, r , and an income process defined as: $y_t = y_{t-1} + u_t$ (see chapter 1). This income process is slightly simpler than the one I estimated above, and is used only for analytical convenience. Suppose that the conditional variance of income shock can be modelled according to the *ARCH*(1) process: $E_{t-1}(u_t^2) = \sigma_u^2 + \gamma u_{t-1}^2$. Define also $e_t = u_t^2 - E_{t-1}(u_t^2)$ the innovation in the conditional variance of income shocks, with the property that $E_{t-1}(e_t) = 0$. For simplicity, assume that e_t is i.i.d. Caballero shows that under these assumptions the Euler equation can be written as:

$$\Delta c_t \simeq \frac{r - \delta}{\theta} + \frac{\theta}{2} E_{t-1}(v_t^2) + v_t \quad (31)$$

where δ is the intertemporal discount rate, and v_t is the consumption change innovation. Given the income process, it can be shown that the consumption innovation can be written as:

$$v_t = u_t - (\theta/2) \frac{1}{1 + r - \gamma} e_t \quad (32)$$

The role of e_t can be made clear with the following example. Suppose $e_t < 0$, for instance because at time $t - 1$ the agent under-estimated income volatility for time t . Because perceived risk was high at time $t - 1$, in that period prudent agents saved more and consumed less, i.e. c_{t-1} was low relatively to c_t . But at time t a negative variance innovation occurs and this suggests to consume more today, thus generating a positive consumption change between $t - 1$ and t . Note that the negative effect described above is higher, in absolute value, when $\gamma > 0$, i.e. when there is positive state dependence. Thus shocks to the variance reduce consumption growth, but such effect is attenuated by the possibility of negative state dependence.

Apart from this effect, state dependence in the variance function also affects the marginal propensity to consume out of the income shock. Let us consider the following expressions:

$$E(v_t|u_t) = u_t - (\theta/2) \frac{1}{1+r-\gamma} \rho_{u_t e} \frac{\sigma_e}{\sigma_{u_t}} u_t \quad (33)$$

i.e., the expectation of the consumption innovation given the income innovation. In this equation: $\rho_{u_t e} = \frac{\sigma_{u_t e}}{\sigma_{u_t} \sigma_e}$ is the correlation coefficient between the income shock and the variance shock. As showed by Caballero, when income shocks and variance shocks are positively correlated ($\rho_{u_t e} > 0$), the marginal propensity to consume out of an income shock is smaller than in a model with i.i.d. income shocks. Moreover, the marginal propensity to consume is even smaller when $\gamma < 0$. Thus stochastic risk might explain the low propensity to consume out of unanticipated income changes that is usually attributed to liquidity constraints; the possibility of negative state dependence in the variance function reinforces this conclusion, while positive state dependence makes it weaker. Finally, note that if income shocks and income variance shocks are uncorrelated, the standard certainty equivalence result obtains.

Consider now the precautionary saving effect. This effect is complicated by the presence of risk dynamics. From (5.2):

$$E_{t-1}(v_t^2) = E_{t-1}(u_t^2) + (\theta^2/4) \frac{1}{(1+r-\gamma)^2} \sigma_e^2 - \frac{\theta}{(1+r-\gamma)} E_{t-1}(u_t e_t) \quad (34)$$

and using the fact that $E_{t-1}(u_t^2) = \sigma_u^2 + \gamma u_{t-1}^2$:

$$E_{t-1}(v_t^2) = \sigma_u^2 + \gamma u_{t-1}^2 + (\theta^2/4) \frac{1}{(1+r-\gamma)^2} \sigma_e^2 - \frac{\theta}{(1+r-\gamma)} E_{t-1}(u_t e_t) \quad (35)$$

The impact of the conditional variance of income on the change in consumption (as measured by $\frac{\theta}{2}$), is not affected by the sign or magnitude of γ , as can be seen from (5.4). However, from equation (5.5) it is clear that the sign matters for the effect of past shocks squared in the reduced form; in an ARCH model, the volatility of past shocks is a predictor for future volatility (although past shocks do not help predicting future shocks). A shock at time $t-1$, of any sign, signals an increase (a decline) in the risk expected for time t if $\gamma > 0$ ($\gamma < 0$), and therefore increases (reduces) the amount of savings needed for precautionary purposes. Note also that precautionary savings may arise from at least three

other sources: the constant component of the conditional variance of income σ_u^2 (the standard comparative statics result), the variance of the variance shock (σ_e^2), and –provided $E_{t-1}(u_t e_t) > 0$, as assumed by Caballero– the correlation between income shocks and income variance innovations. They all affect the conditional variance of consumption. Of course, the strength of the last two effects will depend on both the sign and the magnitude of γ . For instance, the effect of σ_e^2 on consumption will be smaller when $\gamma < 0$ than when $\gamma > 0$.

A different problem is related to the divergence between the amount of income variability estimated by the econometrician and the uncertainty truly faced at household level. In fact, not all the variance, conditional or otherwise, is necessarily uncertainty when seen from the point of view of the household. One way of accounting for partial consumption insurance has been suggested by Banks, Blundell and Brugiavini (1997), and is a fruitful way to solve the problem of bounding measures of microeconomic uncertainty estimated from panel data. They suggest to test for the existence of partial consumption insurance by comparing the coefficients for the variance of the idiosyncratic shock with the coefficient attached to the variance of the aggregate shock in a standard Euler equation framework. Under the null of no insurance, they should be equal; if partial insurance does exist, the latter should exceed the former. A different way to assess the discrepancy of information between the household and the econometrician is to confront measures of uncertainty obtained *via* estimation of dynamic income processes with measures of risk recovered from subjective expectations data. Data on the subjective distribution of future incomes or the probability of future unemployment are now becoming available also for the US (in particular, the Survey of Economic Expectations and the Health and Retirement Survey), and have been used, among others, by Dominitz and Manski (1998) and Barski, Juster, Kimball, and Shapiro (1997). This is an interesting avenue for future empirical research.

6 Conclusions

I began this chapter with the consideration that for models of partial or general equilibrium is important to measure microeconomic uncertainty. With this motivation in the background, I have presented an empirical strategy that allows to model the variance of income shocks as a stochastic autoregressive process; while there can be various ways to measure individual risk, I chose to focus on the variance of earnings innovation both for comparison with previous studies and because the variance is likely to proxy more accurately for uncertainty than other statistics of the distribution of earnings. I have showed that the autoregressive process for the conditional variance of earnings innovation is identifiable by combining information on the variance and the autocovariances at various length of earnings residuals in first difference. To discriminate amongst the various error structure presented in the literature, I construct a simple statistical test. In particular, I distinguish three traditional error structures: no permanent component, a time-invariant permanent component, and a stochastically evolving permanent component (i.e., a martingale in levels). To account for the possibility that individuals with different education levels face different income processes, both at the mean and the variance level, I estimate separate income processes for three education groups I can distinguish in the population.

The results of my empirical analysis, conducted using generalized method of moments estimators of the type discussed in Arellano and Bond (1991a), are generally consistent with the previous empirical literature: I find that simple models where permanent components are absent or fixed over time are rejected, and that the permanent component is more likely to follow a martingale process. On the other hand, transitory unobservable changes in earnings, due to measurement error or genuine shocks, follow a serially correlated process with low persistence. My preferred specification for the rate of growth of individual earnings is $ARMA(1,1)$, a finding that agrees with previous empirical studies conducted on the PSID.

Although I account for stochastic risk, it is fair to say that there are at least two problems that remain unsolved by my empirical strategy. The first is that transitory shocks are indistinguishable from genuine measurement errors;

the second is that even in the presence of perfectly measured income part of the variability I measure is not necessarily uncertainty as perceived by the household. This happens because households condition on a richer information set than the econometrician (for instance, households can exploit insurance opportunities that are ignored by the econometrician). Nonetheless, I find that while changes in the variance of the transitory component do not seem to behave according to an *ARCH* process, the variance of the permanent shock displays significant state dependence. Moreover, I find substantial heterogeneity when evaluating risk evolution across individuals. Such evidence weakens the conventional wisdom of treating microeconomic uncertainty as a parameter to estimate; my empirical analysis demonstrates that individual risk is more likely to follow a distinct stochastic process whose properties ought to be carefully accounted for.

I find striking differences across education groups. Transitory risk is mainly explained by long-run differences across individuals, with deviations around individual trends being due to counter-cyclical economy-wide fluctuations. Interestingly, the latter are not important for the most educated, pointing out to unemployment risk being particularly important for individuals with little education. Permanent risk for the high-school graduate (the largest group in my sample) does not display idiosyncratic dynamics, while for the groups at the two extremes of the distribution of schooling the estimate of the *ARCH* coefficient displays a negative sign. In other words, I find that for these groups permanent shocks in any given time period produce a decline in the conditional variance of tomorrow's permanent shock. One interpretation I have considered is that since permanent shocks are likely to occur at extremely low frequencies over a worker's labour market career, they generate a zero-conditional *mean preserving compression* in the distribution of income. To put it more simply, the likelihood of being promoted or fired drops immediately after either event has occurred. I interpret this finding on the ground of uncertainty resolution.

Of course, I cannot rule out that the same result can be explained by model misspecification. I am particularly aware of the fact that my empirical findings may signal model misspecification. In fact, there is no theoretical background

for microeconomic uncertainty. Lacking a well defined theoretical structure, a linear *ARCH* specification is just a choice available amongst many, and has been adopted here essentially because of its tractability. It is perhaps possible that such structure is just a local approximation for a highly non-linear specification, where risk depends on features of the distribution of earnings that I ignore or am unable to characterize.

Table 1A
Descriptive statistics

Year	Number of yearly observations	Number of years	Number of individuals	Number of individuals		
				High-school dropout	High-school graduate	College graduate
1967	675	11	192	52	100	40
1968	717	12	157	37	92	28
1969	760	13	155	44	84	27
1970	809	14	146	34	90	22
1971	851	15	149	38	78	33
1972	933	16	129	34	68	27
1973	1006	17	119	42	47	30
1974	1089	18	102	22	50	30
1975	1157	19	100	23	52	25
1976	1242	20	107	25	55	27
1977	1329	21	68	11	37	20
1978	1352	22	56	17	25	14
1979	1398	23	53	14	25	14
1980	1422	24	47	13	24	10
1981	1460	25	100	29	47	24
1982	1400					
1983	1337					
1984	1270					
1985	1227					
1986	1177					
1987	1128					
1988	1076					
1989	1031					
1990	981					
1991	933					

Table 1B¹²⁰

Descriptive statistics: log annual earnings

Year	High-school dropout	High-school graduate	College graduate
1967	9.87 (0.59)	10.34 (0.37)	10.60 (0.60)
1968	9.94 (0.60)	10.37 (0.38)	10.69 (0.52)
1969	9.96 (0.59)	10.41 (0.36)	10.68 (0.55)
1970	9.95 (0.60)	10.41 (0.36)	10.68 (0.54)
1971	9.98 (0.55)	10.39 (0.43)	10.62 (0.66)
1972	10.00 (0.62)	10.42 (0.45)	10.68 (0.55)
1973	10.05 (0.58)	10.43 (0.43)	10.66 (0.55)
1974	9.89 (0.61)	10.37 (0.46)	10.62 (0.54)
1975	9.89 (0.69)	10.30 (0.57)	10.60 (0.51)
1976	9.97 (0.67)	10.34 (0.56)	10.57 (0.68)
1977	10.00 (0.60)	10.36 (0.48)	10.63 (0.52)
1978	10.01 (0.66)	10.36 (0.48)	10.67 (0.48)
1979	10.00 (0.65)	10.33 (0.52)	10.63 (0.53)
1980	9.88 (0.69)	10.28 (0.52)	10.59 (0.55)
1981	9.83 (0.70)	10.26 (0.50)	10.59 (0.57)
1982	9.69 (0.83)	10.18 (0.63)	10.61 (0.52)
1983	9.78 (0.78)	10.20 (0.62)	10.61 (0.53)
1984	9.79 (0.77)	10.28 (0.55)	10.64 (0.52)
1985	9.83 (0.73)	10.28 (0.58)	10.68 (0.52)
1986	9.80 (0.78)	10.33 (0.52)	10.75 (0.49)
1987	9.85 (0.63)	10.35 (0.52)	10.80 (0.48)
1988	9.82 (0.69)	10.34 (0.54)	10.84 (0.47)
1989	9.80 (0.67)	10.33 (0.60)	10.84 (0.50)
1990	9.77 (0.75)	10.34 (0.53)	10.83 (0.51)
1991	9.77 (0.65)	10.32 (0.57)	10.83 (0.53)

¹²⁰Standard deviation in parenthesis.

Table 2
The conditional mean of transitory shocks
(Model A)¹²¹

	High-school dropout	High-school graduate	College graduate
Autoregr. coeff.	0.1919 (0.0212)	0.2887 (0.0213)	0.2806 (0.0229)
Age	0.0346 (0.0092)	0.0705 (0.0051)	0.1352 (0.0076)
Age ²	-0.0003 (0.0001)	-0.0007 (0.0051)	-0.0014 (0.0001)
Sargan (p-value)	0.147	0.445	0.249
m ₂ (p-value)	0.810	0.600	0.072
Variance	0.1501 (0.0033)	0.0948 (0.0018)	0.0716 (0.0022)
Number of obs.	7,217	14,082	6,174

¹²¹Asymptotic standard errors are reported in round brackets under the coefficient estimate. The Sargan test is robust to conditional heteroscedasticity. The m_2 is a test for the null hypothesis of no second-order correlation in the residuals (Arellano and Bond, 1991). In table 2 the variance is calculated using the formula reported in note 22.

Table 3
The conditional variance of transitory shocks
(Model A)¹²¹

	High-school dropout		High-school graduate		College graduate	
<i>ARCH</i> coeff.	0.0342 (0.0047)	0.0038 (0.0073)	0.0651 (0.0045)	0.0154 (0.0118)	0.0443 (0.0012)	0.0243 (0.0022)
Asymmetry		-0.0843 (0.0164)		-0.1635 (0.0135)		-0.1239 (0.0053)
Age	0.0024 (0.0005)	0.0025 (0.0006)	0.0007 (0.0006)	0.0018 (0.0005)	0.0010 (0.0002)	0.0012 (0.0002)
Sargan (p-value)	0.083	0.078	0.338	0.166	0.251	0.002
m ₂ (p-value)	0.569	0.632	0.323	0.594	0.287	0.510
Number of obs.	6,661	6,661	12,966	12,966	5,723	5,723

Table 4
The conditional mean of earnings¹²²

	GMM first differences			OLS levels		
	High school dropout	High school graduate	College graduate	High school dropout	High school graduate	College graduate
First lag	0.1740 (0.0237)	0.2684 (0.0247)	0.2637 (0.0260)	0.7877 (0.0120)	0.7835 (0.0054)	0.8091 (0.0072)
Age	-0.0018 (0.0003)	-0.0011 (0.0002)	-0.0019 (0.0002)	0.0030 (0.0066)	0.0057 (0.0037)	0.0113 (0.0049)
Age ²				-0.0000 (0.0001)	-0.0001 (0.0000)	-0.0001 (0.0001)
Wald test (d.o.f.) (p-value)	200.82 (23) (0.0000)	229.19 (23) (0.0000)	102.72 (23) (0.0000)	127.95 (23) (0.0000)	180.42 (23) (0.0000)	55.96 (23) (0.0000)
Instruments		y_{t-2}, y_{t-3}				
Adj. R^2	0.1355	0.0955	0.0612			
P-value reduced form	0.0000	0.0000	0.0000			
Sargan test (d.o.f.) (p-value)	44.78 (44) [0.439]	47.38 (44) [0.337]	46.72 (44) [0.361]			
m_2 test (p-value)	0.647	0.423	0.103	0.042	0.277	0.000
# of obs.	7,217	14,082	6,174	7,773	15,168	6,625

¹²² Asymptotic standard errors are reported in round brackets under the coefficient estimate. The Wald statistic tests for the null hypothesis that time dummies are not jointly significant. The Sargan test is robust to conditional heteroscedasticity. The p -value reduced form reports the p -value of Wald-test that the excluded instruments in the reduced form are not jointly significant. The m_2 is a test for the null hypothesis of no second-order correlation in the residuals (Arellano and Bond, 1991).

Table 5
The unconditional variance of income shocks¹²³

	Whole sample	High-school dropout	High-school graduate	College graduate
Transitory shock	0.0685 (0.0037)	0.0998 (0.0080)	0.0645 (0.0051)	0.0413 (0.0072)
Permanent shock	0.0117 (0.0023)	0.0192 (0.0062)	0.0102 (0.0028)	0.0066 (0.0030)

¹²³Standard errors in parenthesis.

Table 6
The identification of the ARCH coefficients¹²⁴

	Transitory shock			Permanent shock		
	High school dropout	High school graduate	College graduate	High school dropout	High school graduate	College graduate
<i>P</i> -value reduced form	0.0748	0.9038	0.6747	0.0000	0.1709	0.0000
Adj. <i>R</i> ² excluded instruments	0.0007	-0.0002	-0.0003	0.0068	0.0002	0.0086
	Estimated autocovariances of $\Delta v_{it+1}v_{it}$			Estimated autocovariances of $\Delta(v_{it}^2 + v_{it+1}v_{it} + v_{it}v_{it-1})$		
Lag order						
1	-0.3632 (0.0888)	-0.2795 (0.0990)	-0.2173 (0.1730)	-0.2802 (0.0581)	-0.1013 (0.0190)	-0.0814 (0.0338)
2	-0.0104 (0.0356)	-0.0474 (0.0308)	-0.0451 (0.0600)	0.1532 (0.0428)	0.0417 (0.0143)	0.0490 (0.0297)
3	0.0074 (0.0202)	0.0013 (0.0068)	0.0069 (0.0068)	-0.1158 (0.0412)	-0.0140 (0.0074)	-0.0309 (0.0201)
4	-0.0012 (0.0115)	0.0082 (0.0071)	-0.0019 (0.0019)	0.0329 (0.0186)	-0.0037 (0.0035)	0.0093 (0.0084)
5	0.0064 (0.0113)	-0.0036 (0.0063)	-0.0017 (0.0025)	-0.0036 (0.0166)	0.0023 (0.0031)	-0.0005 (0.0011)

¹²⁴This table report the *p*-value of a test that the excluded instruments in the reduced form are not jointly significant and the adjusted *R*² of the regression of the lagged endogenous variable on the excluded instruments. In the bottom, I report autocovariance estimates (standard errors in parenthesis).

Table 7
The conditional variance of transitory shocks

	High-school dropout	High-school graduate	College graduate
Age	-0.0310 (0.0207)	-0.0360 (0.0155)	0.0352 (0.0281)
Age ²	0.0001 (0.0002)	0.0001 (0.0001)	0.0000 (0.0002)
<i>P</i> -value time dummies	0.0084	0.0070	0.6253
<i>P</i> -value person dummies	0.0000	0.0000	0.0170

Table 8
The conditional variance of permanent shocks¹²⁵

	GMM first differences			GMM levels		
	High school dropout	High school graduate	College graduate	High school dropout	High school graduate	College graduate
First lag	-0.3755 (0.0154)	0.1082 (0.0384)	-0.1027 (0.0030)	-0.0689 (0.0115)	-0.1547 (0.0518)	-0.0258 (0.0018)
Age	-0.0001 (0.0005)	-0.0000 (0.0002)	-0.0000 (0.0002)	-0.0046 (0.0071)	-0.0050 (0.0050)	-0.0082 (0.0028)
Age ²				0.0001 (0.0001)	0.0001 (0.0001)	0.0001 (0.0000)
Wald test (<i>p</i> -value)	0.027	0.227	0.036	0.199	0.171	0.000
Instruments	<i>pm_{t-4}, pm_{t-5}</i>			<i>pm_{t-3}, pm_{t-4}</i>		
Sargan test (d.o.f.) (<i>p</i> -value)	24.63 (30) (0.743)	33.76 (30) (0.290)	33.67 (30) (0.294)	35.31 (32) (0.315)	36.45 (32) (0.269)	25.73 (32) (0.775)
<i>m</i> ₂ test (<i>p</i> -value)	0.383	0.021	0.322	0.011	0.012	0.221
# of obs.	4,993	9,738	4,377	5,549	10,824	4,821

¹²⁵Asymptotic standard errors in parenthesis. The Sargan test is robust to conditional heteroscedasticity. The Wald statistic tests for the null hypothesis that time dummies are not jointly significant. The *m*₂ is a test for the null hypothesis of no second-order correlation in the residuals. The variable $pm_{t-j} = v_{t-j-1}^2 + v_{t-j}v_{t-j-1} + v_{t-j-1}v_{t-j-2}$.

6

Conclusions

This dissertation has presented four empirical analyses aimed at testing the validity of the permanent income hypothesis (chapters 2 and 3), the full consumption insurance hypothesis (chapter 4), and models for the conditional variance of income shocks (usually taken as a proxy for income uncertainty). To synthesise the results, I have found the following:

i) Because of lack of power, the Euler equation is perhaps not the appropriate theoretical *locus* where to find evidence for liquidity constraints; even in a sample of Italian households, notoriously affected by borrowing constraints and other imperfections in credit and insurance markets, the excess sensitivity test does not signal model misspecification once one controls for the stochastic structure of the error term using individual subjective expectations. Nonetheless, there is strong support for a precautionary motive for saving.

ii) When the data reveal individual superior information (for instance in the form of subjective expectations of future earnings), it is possible not only to construct estimates of one's unanticipated earnings changes, but also to distinguish between shocks of different nature (transitory and permanent). This is important because the theory predicts that the propensity to save out of an income shock will differ depending upon the nature of the shock experienced.

The findings I report in chapter 2 are consistent with an extended version of the permanent income hypothesis that takes into account both liquidity constraints and precautionary savings.

iii) Consumption is mobile; between 25 and 45 percent of households in Italy move across quartiles of the distribution of consumption. In chapter 3 I have argued that consumption mobility can be interpreted a symptom of incomplete consumption insurance; this is simply because full insurance eliminates all idiosyncrasies in consumption growth rates.

iv) Based on evidence drawn from the 1967-1991 PSID, a longitudinal dataset of US households, chapter 4 shows that transitory income risk fluctuates with aggregate shocks and is shaped by individual heterogeneity; permanent shocks occur very infrequently over the life cycle of an individual and once they do, the conditional variance collapses to the unconditional one. Also, income uncertainty declines with labour market experience, a fact that matches both intuition and casual observation.

In chapters 2 and 3 I have argued that there is a very rich research agenda based on the use of subjective expectations in structural models of consumption behaviour. In particular, one can think of using the identification strategy presented in chapter 3 to run a nested test of five different models of consumption: the keynesian model, the full consumption insurance hypothesis, the life-cycle hypothesis, the permanent income hypothesis and the precautionary saving model. If income is defined by the equations (11)-(12) of chapter 1, the five models can be jointly tested within the following framework:

$$\Delta c_{it} = \alpha_t + (\beta \ \gamma \ \delta \ \eta \ \theta) \begin{pmatrix} \varepsilon_{it} \\ \varepsilon_{it-1} \\ \zeta_{it} \\ \sigma_\varepsilon^2 \\ \sigma_\zeta^2 \end{pmatrix} + v_{it}$$

Based on the keynesian model, $\beta = -\gamma = \delta > 0$ and $\eta = \theta = 0$; for the full consumption insurance hypothesis, $\beta = \gamma = \delta = \eta = \theta = 0$; for the life cycle hypothesis, $\gamma = 0$, $\beta = \left(1 - \frac{1}{(1+r)^{T-t+1}}\right)^{-1} \frac{r}{1+r}$, $\delta = 1$, $\eta = \theta = 0$; for the

permanent income hypothesis: $\gamma = 0$, $\beta = \frac{r}{1+r}$, $\delta = 1$, $\eta = \theta = 0$; finally, for the precautionary saving model: $\gamma = 0$, $\beta = \frac{r}{1+r}$, $\delta = 1$, $\eta = \left(\frac{r}{1+r}\right)^2$ and $\theta = 1$.

These are the obvious restrictions on the parameters arising from the different theories, but other restrictions could be added to take into account some particular form of liquidity constraints, uncertainty on the date of death, etc. Up to my knowledge, nested tests of the five different models have never been attempted on microeconomic data. This is an interesting avenue of research that I intend to pursue in the future.

The results presented in chapter 5 also call for further investigation, both on the theoretical and on the empirical side. At present, the more important challenge to confront is probably to try to develop theoretical models where both the mean and the variance of earnings are characterized. In this sense, we need for the variance the same kind of theoretical development that is associated with the studies of Becker (1975) and Farber and Gibbons (1995) on the evolution of mean earnings. On the empirical side, it is probably worth assessing whether the between-group differences I found are capturing genuine features of the functioning of the US labour market, in particular those related to the distribution of risk in the population. Moreover, it can be interesting to consider alternative empirical models of risk evolution, i.e. examining various models for the variance of income innovations to match what MaCurdy (1981) did for the evolution of the first moments of wages and earnings.




Figure 1: Fraction of households changing quartile, 1987-1995

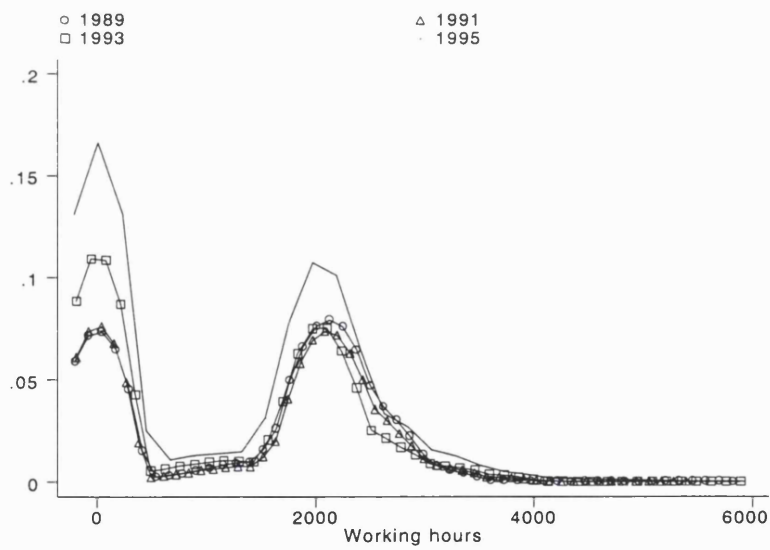


Figure 2: The distribution of working hour, 1989-1995

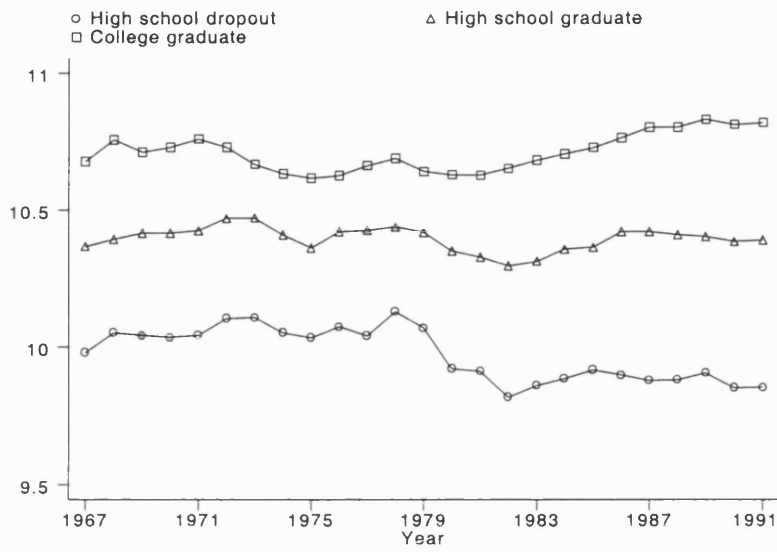


Figure 3: The median of log earnings, by education




Figure 4: The interquartile range of log earnings, by education

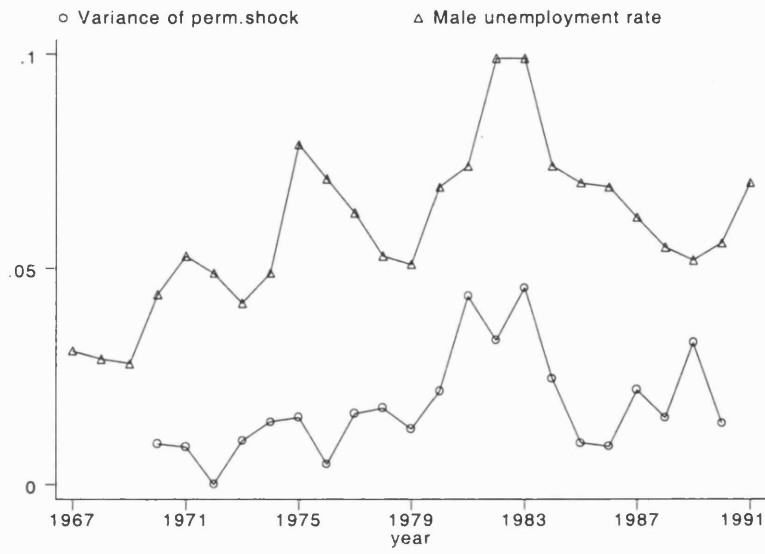


Figure 5: The variance of the permanent shock and the business cycle




Figure 6: The variance of the transitory shock and the business cycle

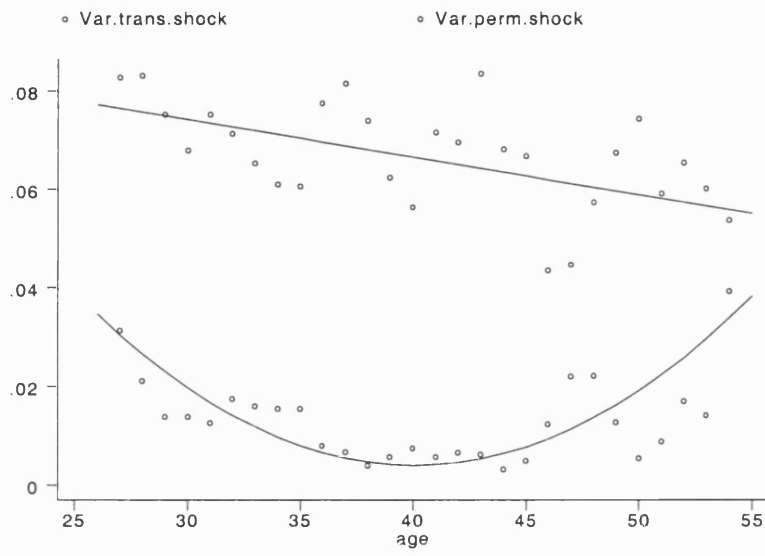


Figure 7: The variance of income shocks and the life-cycle

Bibliography

- [1] Abowd, John and David Card (1989), "On the covariance structure of earnings and hours changes", *Econometrica*, **57**, 411-45
- [2] Alessie, Rob, and Annamaria Lusardi (1997), "Saving and income smoothing: Evidence from panel data", *European Economic Review*, **41**, 7, 1251-79
- [3] Altonji, Joseph G., and Aloysious Siow (1987), "Testing the response of consumption to income with (noisy) panel data", *Quarterly Journal of Economics*, **102**, 293-328.
- [4] Alvarez, Fernando, and Urban J. Jermann (1998), "Asset pricing when risk sharing is limited by default", University of Chicago, *mimeo*.
- [5] Arellano, Manuel, and Stephen Bond (1991a), "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations", *Review of Economic Studies*, **58**, 277-97
- [6] Arellano, Manuel, and Stephen Bond (1991b), "Dynamic panel data estimation using DPD - A guide for users", *Institute for Fiscal Studies*, Working paper 88/15, London
- [7] Attanasio, Orazio P. (1998), "Consumption", forthcoming in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: North Holland.
- [8] Attanasio, Orazio P., and Steve Davis (1996), "Relative wage movements and the distribution of consumption", *Journal of Political Economy*, **104**, 1227-62.
- [9] Attanasio, Orazio P., and Guglielmo Weber (1995), "Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey", *Journal of Political Economy*, **103**, 1121-57.

- [10] Baker, Michael (1997), "Growth-rate heterogeneity and the covariance structure of life-cycle earnings", *Journal of Labour Economics*, **15**, 2, 338-375
- [11] Banerjee, Abhijit V. and Andrew F. Newman (1991), "Risk-bearing and the theory of income distribution", *Review of Economic Studies*, **58**, 211-35
- [12] Banks, James, Richard Blundell and Agar Brugiavini (1998), "Risk pooling and consumption growth", IFS, *mimeo*.
- [13] Barski R., Thomas Juster, Miles Kimball, and Mathew Shapiro (1997), "Preference parameters and behavioral heterogeneity: An experimental approach in the Health and Retirement Survey", *Quarterly Journal of Economics*, **12**, 537-80.
- [14] Becker, Gary S. (1975), "*Human Capital*", 2nd edition. Chicago: University of Chicago Press, 1975
- [15] Bekker, P.A. (1994), "Alternative approximations to the distribution of instrumental variables estimators", *Econometrica*, **62**, 657-82
- [16] Bertola, Giuseppe, and Andrea Ichino (1996), "Crossing the river: A comparative perspective of Italian employment dynamics, *Economic Policy*, 361-420
- [17] Bishop, Y.M.M., S.E. Fienberg, and P.W. Holland (1988), *Discrete multivariate analysis*, MIT Press.
- [18] Blundell, Richard W., and Ian Preston (1997), "Consumption inequality and income uncertainty", forthcoming in *Quarterly Journal of Economics*.
- [19] Browning, Martin J., Lars P. Hansen, and James J. Heckman (1998), "Micro data and general equilibrium models", forthcoming in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: North Holland.
- [20] Browning, Martin J., and Annamaria Lusardi (1996), "Household saving: mi-

cro theories and micro facts”, *Journal of Economic Literature*, **34**, 1797-1855.

- [21] Brandolini, Andrea (1999), “The personal distribution of income in post-war Italy: source description, data quality, and the time pattern of income inequality”, *Giornale degli Economisti* (forthcoming)
- [22] Brandolini, Andrea, and Luigi Cannari (1994), “Methodological Appendix: The Bank of Italy Survey of Household Income and Wealth”, in *Saving and the accumulation of wealth. Essays on Italian households and government behavior*, edited by Albert Ando, Luigi Guiso and Ignazio Visco. Cambridge: Cambridge University Press.
- [23] Caballero, Ricardo J. (1990), “Consumption puzzles and precautionary savings”, *Journal of Monetary Economics*, **25**, 113-36
- [24] Campbell, John Y. (1987), “Does saving anticipates declining labour income? An alternative test of the permanent income hypothesis”, *Econometrica*, **55**, 1249-73
- [25] Carroll, Chris D. (1992), “Buffer stock saving and the Life Cycle/Permanent Income Hypothesis”, *Brookings Papers on Economic Activity*, **2**, 61-156.
- [26] Carroll, Cristopher and Andrew Samwick (1997), “The nature of precautionary saving”, *Journal of Monetary Economics*, **40**, 41-71.
- [27] Chamberlain, Gary, (1984), “Panel data”, in *Handbook of Econometrics*, vol. 2, edited by Zvi Griliches and Michael D. Intriligator. Amsterdam: North-Holland.
- [28] Cochrane, John (1991), “A simple test of consumption insurance”, *Journal of Political Economy* **99**, 957-976
- [29] Cole, H. and Narayana Kocherlakota (1997), “Efficient allocations with hidden income and hidden storage”, Federal Reserve Bank of Minneapolis research

Department Staff Report 238.

- [30] Costain, James S. (1998), "A simple model of multiple equilibria based on risk", Universitat Pompeu Fabra, *mimeo*.
- [31] Das, Marcel, and Arthur Van Soest (1997), "Expected and realized income changes: evidence from the Dutch Socio-Economic Panel", *Journal of Economic Behavior and Organization*, **32**, 137-54.
- [32] Deaton, Angus (1991), "Saving and liquidity constraints", *Econometrica*, **59**, 1221-48.
- [33] Deaton, Angus (1992a), "Saving and income smoothing in the Côte d'Ivoire", *Journal of African Economics*, **1**, 1-24
- [34] Deaton, Angus (1992b), *Understanding consumption*. Oxford: Oxford University Press.
- [35] Deaton, Angus (1997), *The analysis of household surveys*. Baltimore: The Johns Hopkins University Press.
- [36] Deaton, Angus, and Christina Paxson (1994), "Intertemporal choice and inequality", *Journal of Political Economy*, **102**, 384-94.
- [37] Dominitz, Jeff, and Charles F. Manski (1997), "Using expectation data to study subjective income expectations", *Journal of the American Statistical Association*
- [38] Dynan, Karen E. (1993), "How prudent are consumers?", *Journal of Political Economy*, **101**, 1104-13.
- [39] Eberly, Janice (1994), "Adjustments of consumers'durables stocks: Evidence from automobile purchases", *Journal of Political Economy*, **102**, 403-36.
- [40] Engle, Richard F. (1982), "Autoregressive conditional heteroscedasticity with estimates of the variance of the United Kingdom inflation", *Econometrica*,

50, 987-1007.

- [41] Farber, Henry S. and Robert Gibbons (1995), "Learning and wage dynamics", *Quarterly Journal of Economics*, 1007-46
- [42] Flavin, Marjorie (1981), "The adjustment of consumption to changing expectations about future income", *Journal of Political Economy*, **89**, 974-1009
- [43] Flavin, Marjorie (1993), "The excess smoothness of consumption: Identification and interpretation", *Review of Economic Studies*, **60**, 651-66
- [44] Flavin, Marjorie (1994), "The joint consumption/asset demand decision: a case study in robust estimation". University of California - San Diego, *mimeo*.
- [45] Friedman, Milton (1957), *A theory of the consumption function*, Princeton: Princeton University Press
- [46] Garcia, René, Annamaria Lusardi, and Serena Ng (1997), "Excess sensitivity and asymmetries in consumption: an empirical investigation", *Journal of Money, Credit, and Banking*, **29**, 154-76.
- [47] Gottshalk, Peter and Robert Moffitt (1993), "The growth of earnings instability in the U.S. labour market", *Brooking Papers on Economic Activity*, 217-272
- [48] Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese (1992), "Earnings uncertainty and precautionary saving", *Journal of Monetary Economics*, **30**, 307-37.
- [49] Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese (1994), "Why is Italy's saving rate so high?", in *Saving and the accumulation of wealth. Essays on Italian households and government behavior*, edited by Albert Ando, Luigi Guiso and Ignazio Visco. Cambridge: Cambridge University Press.
- [50] Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese (1996), "Income risk, borrowing constraints and portfolio choice," *American Economic Review*, **86**,

- [51] Hall, Robert (1978), "Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence", *Journal of Political Economy*, **96**, 971-87
- [52] Hall, Robert, and Frederic Mishkin (1982), "The sensitivity of consumption to transitory income: Estimates from panel data of households", *Econometrica*, **50**, 461-77
- [53] Hassler, John A.A. (1996), "Variations in risk and fluctuations in demand: A theoretical model", *Journal of Economic Dynamics and Control*, **20**, 1115-43.
- [54] Hayashi, Fumio (1985), "The permanent income hypothesis and consumption durability: analysis based on Japanese panel data", *Quarterly Journal of Economics*, **100**, 1083-113.
- [55] Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff (1996) "Risk sharing between and within families", *Econometrica* **64**, 261-94.
- [56] Hill, Martha S. (1992), *The Panel Study of Income Dynamics: A user's guide*, Newbury Park, California: Sage Publications.
- [57] Hubbard, Glenn, Jonathan Skinner, and Stephen Zeldes (1994), "The importance of precautionary motives in explaining individual and aggregate saving", *Carnegie Rochester Series on Public Policy*, **40**, 59-125.
- [58] Kimball, Miles (1990), "Precautionary saving in the small and in the large", *Econometrica*, **58**, 53-73.
- [59] Kimball, Miles (1993), "Standard risk aversion", *Econometrica*, **61**, 589-611.
- [60] Krueger, Dirk and Fabrizio Perri (1999) "Risk sharing: Private insurance markets or redistributive taxes", Federal Reserve Bank of Minneapolis Research

Department, *mimeo*.

- [61] Levy, Frank, and Richard J. Murnane (1992), "Earnings levels and earnings inequalities: A review of recent trends and proposed explanations", *Journal of Economic Literature*, **30**, 1333-81
- [62] Lillard, Lee, and Robert Willis (1978), "Dynamic aspects of earnings mobility", *Econometrica*, **46**, 985-1012
- [63] Ludvigson, Sydney, and Christina Paxson (1997), "Approximation bias in linearized Euler equations", Princeton University, *mimeo*.
- [64] Lusardi, Annamaria (1996), "Permanent income, current income and consumption: evidence from two panel data sets", *Journal of Business and Economics Statistics*, **14**, 81-90.
- [65] Lusardi, Annamaria (1997), "Precautionary saving and subjective earnings variance", *Economic Letters*
- [66] MaCurdy, Thomas E. (1982), "The use of time series processes to model the error structure of earnings in a longitudinal data analysis", *Journal of Econometrics*, **18**, 82-114.
- [67] Manacorda, Marco (1997), "Mind the step. The evolution of male wage inequality in Italy and the escalator clause", Centre for Economic Performance, *mimeo*
- [68] Meghir, Costas, and Guglielmo Weber (1996), "Intertemporal non-separability or borrowing restrictions? A disaggregated analysis using a US consumption panel", *Econometrica*, **64**, 1151-82.
- [69] Miniaci, Raffaele, and Guglielmo Weber (1999), "The Italian recession of 1993: aggregate implications of microeconomic evidence", *Review of Economics and Statistics* (forthcoming)

- [70] Moffitt, Robert and Peter Gottschalk (1994), "Trends in the autocovariance structure of earnings in the U.S.: 1969-1987", Brown University, Department of Economics, *mimeo*.
- [71] Muth, J.F. (1960), "Optimal properties of exponentially weighted forecasts", *Journal of the American Statistical Association*, **55**
- [72] Quah, Danny T. (1990), "Permanent and transitory movements in labor income: An explanation for excess smoothness in consumption", *Journal of Political Economy*, **98**, 449-75
- [73] Schluter, Karl (1998), "Statistical inference for mobility indexes", *Economic Letters*
- [74] Sentana, Enrique (1995), "Quadratic ARCH models", *Review of Economic Studies*, **62**, 639-61.
- [75] Shea, John (1995), "Union contracts and the life-cycle permanent income hypothesis", *American Economic Review*, **85**, 186-200.
- [76] Shorrocks, Anthony F. (1976), "Income mobility and the Markov assumption", *Economic Journal*, **86**, 566-78
- [77] Silverman, B.W. (1986), *Density estimation for statistics and data analysis*, Chapman and Hall.
- [78] Staiger, D., and J.H. Stock (1994), "Instrumental variables regression with weak instruments", *Econometrica*, **65**, 557-86.
- [79] Storesletten, Kjetik, Chris Telmer, and Amir Yaron (1997), "Persistent idiosyncratic shocks and incomplete markets", Carnegie Mellon University, *mimeo*.
- [80] Townsend, Robert (1994), "Risk and insurance in village India", *Econometrica*, **62**, 539-91.

- [81] Zeldes, S. P. (1989a) "Consumption and liquidity constraints: an empirical investigation," *Journal of Political Economy*, **97**, 305-46.
- [82] Zeldes, Stephen (1989b), "Optimal consumption with stochastic income: Deviations from certainty equivalence", *Quarterly Journal of Economics*, **104**, 275-98.
- [83] Zeldes, Stephen (1992), "Comment" to C. Carroll, "Buffer stock saving and the Life Cycle/Permanent Income Hypothesis", *Brookings Papers on Economic Activity*, **2**, 61-156.
- [84] Ziliak, James P. (1997), "Efficient estimation with panel data when instruments are predetermined: An empirical comparison of moment-condition estimators", *Journal of Business Economics and Statistics*, **15**, 419-31