

Intertemporal choice and consumption mobility*

Tullio Jappelli[†] and Luigi Pistaferri[‡]

August 2000

Abstract

The theory of intertemporal consumption choice makes sharp predictions about the evolution of the entire distribution of household consumption, not just about its conditional mean. In a first step, we study the empirical transition matrix of consumption using a panel drawn from the Bank of Italy Survey of Household Income and Wealth. In a second step, we simulate the transition matrix of the consumption distribution using parameters for the income process estimated on the same dataset. Comparison between the actual and the simulated transition matrix for consumption is favorable to the permanent income hypothesis once we allow for measurement error in consumption and a moderate degree of excess sensitivity to income shocks. The theory of consumption insurance is strongly rejected.

Keywords: Consumption dynamics, mobility

JEL Classification: D52; D91; I30.

*We thank Chris Carroll and Raffaele Miniaci for comments. The Training and Mobility of Researchers Network Program (TMR), the Italian National Research Council (CNR), the Italian Ministry for Universities and Scientific Research (MURST) and the “Taube Faculty Research Fund” at the Stanford Institute for Economic and Policy Research provided financial support.

[†]CSEF, University of Salerno, and CEPR.

[‡]Stanford University, SIEPR and CEPR.

1 Introduction

Household consumption is widely used for the purpose of measuring poverty and, more generally, for welfare analyses. There are also important policy implications related to the degree of consumption mobility a society experiences. At the two opposite extremes, societies could be completely polarized or completely mobile. In the first case, poverty is a permanent status (at least in relative terms) and any escape from the poverty trap does not last long; in the second case, poverty is a transitory state, but the high level of consumption volatility induced by a strong mobility up and down the socio-economic ladder may have the undesirable effect of increasing individual uncertainty.

To our knowledge, there is no empirical work attempting to relate changes in the distribution of income or in the demographic structure of the population with consumption mobility. Cutler and Katz (1992) examine consumption and income data in the Consumer Expenditure Survey, but their analysis, as that of Johnson and Shipp (1997), is descriptive and concerned with consumption inequality rather than consumption mobility. In contrast, there is a long tradition of studies of earnings and income mobility. Existing contributions can be divided into two broad groups. In a first group, the focus is on analyzing transition probabilities across quantiles of the earnings distribution by Markov-chain models of earnings mobility (Shorrocks, 1976). In a second group of studies, a process for the conditional mean of earnings is specified and estimated. To gauge the ability of the model to fit the existing data, transition probabilities are estimated using regression analysis and compared to the actual transition probabilities (Lillard and Willis, 1978).

In this paper we attempt to understand which model of intertemporal consumption choice is capable of explaining the amount of consumption mobility we observe in the data. We focus on the theory of consumption insurance, the rule-of-thumb model, and the permanent income hypothesis (PIH). We calculate the theoretical degree of consumption mobility stemming from these models and compare them statistically with the actual amount of mobility estimated in the data. The exercise is performed constructing a transition matrix for consumption and applying non-parametric statistical tools to test different hypotheses concerning consumption dynamics. Since to measure consumption mobility one needs to follow households over time, the empirical analysis is conducted on a panel drawn from the Bank of Italy Survey of Household Income and Wealth for the years 1987 to 1995. The survey we use is representative of the Italian population, spans nine years of data, contains a measure of total non durable consumption and has good quality income data. Since there are virtually no panel datasets with broad consumption measures, a by-product of this paper is to bring the dataset to the attention

of empirical macroeconomists.

To see how the theory of intertemporal choice delivers implications for consumption mobility, consider first the extreme case of full consumption insurance. According to this theory, the cross-sectional distribution of consumption of any group of households is constant over time. 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 over time. Consumption insurance makes therefore strong predictions about the entire consumption distribution, not just its mean or variance. In particular, consumption insurance implies the total absence of consumption mobility between any two time periods, regardless of the individual income shocks. It follows that if one observes people 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.

A second extreme case is the rule-of-thumb model which predicts that households set consumption equal to income in each period. Given that any change in current income translates into an equivalent change in consumption, one should expect a relatively high degree of consumption mobility if shocks are not correlated with the rank position in the initial distribution of consumption.

In more realistic models with incomplete insurance, individuals use saving as a self-insurance device and are able to smooth away at least some of the income variability. The case we consider is one in which income shifts over time because of transitory (e.g., mean reverting) and permanent (e.g., persistent or non-mean reverting) shocks. If people behave according to the PIH, consumption reacts mostly to permanent unanticipated income shocks, but is almost insensitive to transitory ones. Households will therefore move up and down in the consumption distribution only in response to permanent shocks. Thus one should expect a degree of mobility that is intermediate between the level predicted by the consumption insurance hypothesis and the rule-of-thumb model. As we shall see, versions of this model with quadratic utility or precautionary saving make similar (although not identical) predictions about the impact of income shocks on consumption mobility.

Our work is related to Deaton and Paxson (1994), who focus on consumption inequality - the cross-sectional dispersion in consumption - rather than consumption mobility - the individual transitions across the consumption distribution. They derive implications for consumption inequality in mod-

els of intertemporal choice and show that the PIH implies that the cross-sectional dispersion in consumption of any given generation increases over time. Consumption insurance instead implies that the cross-sectional variance of consumption of the same generations should be constant over time. As in their paper, we also derive unexplored theoretical predictions and confront them with the data. Our test requires panel data, while their analysis can be performed with repeated cross-sectional data.

The rest of the paper is organized as follows. Section 2 presents the mobility index and the non-parametric test of consumption mobility. The data and the empirical results are presented in Section 3. In Section 4 we review the implications for consumption dynamics of the theories of intertemporal consumption choice and consider how to account for measurement error in consumption. In Section 5 we calibrate the simulations. In section 6 we present the simulation results and confront the theoretical predictions with the empirical transition probabilities. We strongly reject full consumption insurance; the PIH is the model that fits the data best once we allow for realistic values of measurement error in consumption and a moderate amount of excess sensitivity to income shocks. The simulations are also able to reproduce remarkably well the difference in consumption mobility that we observe in samples stratified according to education and year of birth. Section 7 summarizes our results.

2 The mobility index

To summarize the transition matrix for consumption through an appropriate index of mobility, we 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 being the number of quantiles in the distribution. For notational simplicity we consider transition probabilities from period t to period $t + 1$; extending the argument to transition probabilities in periods $t + 2$, $t + 3$, and so on, is straightforward. The generic element of \mathbf{P} 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 \mathbf{P} . 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:¹

¹In 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.

$$S(\mathbf{P}) = \frac{q - \text{trace}(\mathbf{P})}{q} \quad (1)$$

If the probability of being in quantile i in period t is independent of that 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 , and $\text{trace}(\mathbf{P}) = 1$. If the probability of being in quantile i in period t equals that of being in quantile j in period $t + 1$, the typical entry of the transition matrix is $p_{ij} = 1$ for all $i = j$ and 0 otherwise. In this case $\text{trace}(\mathbf{P}) = q$ and $S(\mathbf{P}) = 0$. Since $0 \leq \text{trace}(\mathbf{P}) \leq q$, the mobility index satisfies the condition $0 \leq S(\mathbf{P}) \leq 1$.² $S(\mathbf{P})$ can be interpreted as the proportion of households moving across the consumption distribution between t and $t + 1$.

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

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

and one can test the null hypothesis that $S(\mathbf{P})$ assumes a given value V using the statistic:

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

To assess whether consumption mobility differs statistically over time or between population groups we will find it useful to construct a test of differential mobility between groups or time periods, based on the statistic:

$$Z_2 = \frac{S(\hat{\mathbf{P}}_d) - S(\hat{\mathbf{P}}_k)}{\sqrt{s.e.(S(\hat{\mathbf{P}}_d))^2 + s.e.(S(\hat{\mathbf{P}}_k))^2}} \sim N(0, 1) \quad (3)$$

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 (3) is also asymptotically distributed as a standard normal.

The Shorrocks index of mobility is based on a trace of a matrix, and therefore the same index can be produced by very different transition matrices.

²The upper bound is a case in which all households move to a different quantile so that $\text{trace}(\mathbf{P}) = 0$ and $S(\mathbf{P}) = 1$.

We therefore supplement the analysis by looking at the pattern of the entire transition matrix for consumption. An important advantage of studying transition probabilities is that they are not affected by 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. It is just for a matter of convenience that in the empirical section and in the simulations we study the transition matrix of log consumption.

3 Measuring consumption mobility

From the previous section it is clear that mobility can only be computed with longitudinal data on consumption. For this purpose we use the 1987-1995 panel of the Italian Survey of Household Income and Wealth (SHIW). This 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 intertemporal consumption choice.

The SHIW is conducted by the Bank of Italy which surveys a representative sample of the Italian resident population. Sampling is in two stages, first municipalities and then households. Municipalities are divided into 51 strata defined by 17 regions and 3 classes of population size (more than 40,000, 20,000 to 40,000, less than 20,000). Households are randomly selected from registry office records. From 1987 through 1995 the survey was conducted every other year and covered about 8,000 households, defined as groups of individuals related by blood, marriage or adoption and sharing the same dwelling. Starting in 1989, each SHIW has re-interviewed some households from the previous surveys. The panel component has increased over time: 15 percent of the sample was re-interviewed in 1989, 27 percent in 1991, 43 percent in 1993, and 45 percent in 1995.³ The net response rate (ratio of responses to contacted households net of ineligible units) was 64 percent in 1987, 38 percent in 1989, 33 percent in 1991, 58 percent in 1993, and 57 percent in 1995. Ample details on sampling, response rates, processing of results and comparison of survey data with macroeconomic data are provided

³In the panel component, the sampling procedure is also determined in two stages: (i) selection of municipalities (among those sampled in the previous survey); (ii) selection of households reinterviewed. This implies that there is a fixed component in the panel (for instance, households interviewed 5 times between 1987 to 1995, or 4 times from 1991 to 1995) and a new component every survey (for instance, households reinterviewed only in 1989).

by Brandolini and Cannari (1994).⁴

To minimize measurement error we 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. Self-employment slightly falls over time. Income strongly declines in 1993, a recession year, and consequently dispersion increases. In all years, household disposable income is more variable than consumption: the variance of log income is between 0.10 and 0.20 higher than the variance of log consumption. Note the stability of the variance of log consumption as opposed to the wide fluctuations in the variance of log income. The pattern of the Gini coefficients for consumption and income confirms that the income distribution is less equal than the consumption distribution (34 percent vs. 28 percent). Interestingly, the 1993 recession boosts income inequality while leaving consumption inequality unaffected. These descriptive statistics are consistent with models in which households are able to smooth away at least some of the income shocks.

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

Table 2 reports the transition matrix of (log) per capita consumption from 1987-89 to 1993-95. Recall that the SHIW is conducted every two years, so we observe transitions from period $t-2$ to period t . The elements of the main diagonal report the proportion of households that did not change quartile. For instance, the entry in the top left cell of the 1993-95 panel indicates that

⁴In the panel section, the net response rate was 25 percent in 1989, 54 percent in 1991, 71 percent in 1993, and 78 percent in 1995 (Brandolini, 1999). Response rates increase in 1991 because in that year households included in the panel were chosen among those that had previously expressed their willingness to being re-interviewed.

68 percent of the households in the first quartile in 1993 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 1993 to the second quartile in 1995. The transition matrices for other years are similar, displaying substantial amount of consumption mobility. Note also that all transition matrices are symmetric. The symmetry of a 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 for all years, and does not reject the null hypothesis of symmetry.

The mobility index $S(\hat{\mathbf{P}})$ corresponding to each of the transition matrices in Table 2 is reported in Table 3 together with the associated standard error and the number of transitions. On average, about half of the population moves up or down in the distribution every two years. Consumption mobility ranges from 44 percent in 1993-95 to 50 percent in 1991-93, and is precisely estimated in each year. The swings in mobility that we observe after 1991 are likely to be associated to the deep 1991-93 recession and to the subsequent expansionary period of 1993-95.

An objection to computing mobility using the distribution of per capita consumption is that household expenditures are affected by demographic variables that change over time. One of the most important demographic variables that can affect preferences is certainly household composition. For instance, the arrival of children alters family needs, hence consumption allocations. If household expenditures are characterized by economies of scale, one would observe consumption mobility in consumption per capita even if the distribution of consumption per adult equivalent is constant over time. We thus compute mobility defining transitions in terms of consumption per adult equivalent.⁵ The resulting index shows that using consumption per adult equivalent does not affect the pattern of either the transition matrix or the mobility index. As a further check, we restrict attention to households whose demographic structure did not change over the sample period and find, again, similar consumption transitions.⁶ In the remaining of the paper we thus focus exclusively on per-capita consumption.

As it stands, the mobility index in Table 3 is just a summary measure of the transition matrix. What is most interesting is to derive from theory

⁵The definition of adult equivalent assigns a weight of 1 to the first adult, 0.8 to any additional adult and 0.25 to each household member less than 18 years old. Thus adult equivalent are defined as: $1 + 0.8(\text{Number of adults} - 1) + 0.25(\text{Number of children})$.

⁶For instance, excluding households with changes in family composition results in a mobility index of 0.432 in 1993-95.

meaningful null hypotheses against which data can be confronted. For this purpose we will therefore calculate the theoretically based mobility index under different scenarios (full consumption insurance, rule-of-thumb behavior and the PIH) and test the validity of each of them.

4 Intertemporal choice and mobility

To explore the relation between the consumption and the income distributions, it is useful to start by presenting a fairly general characterization of the income process. Consider the following decomposition of (log) income for household h in period t :

$$\ln y_{h,t} = \beta X_{h,t} + p_{h,t} + e_{h,t} \quad (4)$$

where $X_{h,t}$ is a set of deterministic variables such as age and region of residence, $p_{h,t}$ and $e_{h,t}$ permanent and transitory components, respectively. The latter is the sum of an idiosyncratic ($\varepsilon_{h,t}$) and an aggregate component (ε_t); both are assumed to be serially uncorrelated. Since the permanent component of income changes very slowly, the standard assumption is to model it as a random walk process of the form:

$$p_{h,t} = p_{h,t-1} + z_{h,t} \quad (5)$$

where $z_{h,t}$ is the permanent innovation, which is again the sum of an idiosyncratic ($\zeta_{h,t}$) and an aggregate shock (ζ_t); both components are serially uncorrelated.

The original decomposition of income shocks into transitory and permanent components dates back to Friedman (1957) and starting with Hall and Mishkin (1982) it has become quite standard in panel data studies of income and consumption dynamics. Some of the income shocks are transitory (mean reverting) and their effect does not last long. Examples include fluctuations in overtime labor supply, bonuses, lottery prizes, and bequests. On the other hand, some of the innovations to earnings are highly persistent (non-mean reverting) and their effect cumulates over time. Examples of permanent innovations are generally associated with job mobility, promotions and lay-off.

Thus income growth can be written as:

$$\Delta \ln y_{h,t} = \beta \Delta X_{h,t} + z_{h,t} + \Delta e_{h,t}$$

As we shall see below, the specified income process delivers different implications for consumption mobility in models of intertemporal choice. We also consider how these implications change in the presence of measurement error in consumption.

4.1 The Permanent Income Hypothesis

Under a series of well known hypothesis concerning preferences and technology (infinitely lived households maximizing expected additive and separable utility, perfect credit markets, quadratic utility, and $r = \delta$), one obtains the classical prediction that consumption is a martingale, i.e. that the change in consumption is an innovation, and therefore orthogonal to all past information available to household h :

$$\Delta c_{h,t} = \psi_{h,t}$$

as in Hall (1978). In general, the consumption innovation depends on the sources of uncertainty of the model. If income is the only source of uncertainty, one can derive a closed form solution for the innovation term ψ :

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

Even though the theory delivers a relation between changes in consumption and changes in income, in the empirical analysis we consider a log-linear approximation to the optimal rule

$$\Delta \ln c_{h,t} = \psi'_{h,t}$$

where:

$$\psi'_{h,t} = \frac{r}{1+r} \sum_{\tau=0}^{\infty} \frac{(E_t - E_{t-1}) \ln y_{h,t+\tau}}{(1+r)^\tau}$$

The approximation is warranted because the process for log-income fits much better the data than a process for income levels, and because only log-linearity allows us to nest the three models in a unified framework. Since the mobility index depends only on the relative position of each household over time and not on whether consumption is measured in levels or logs, the approximation should not affect the implied mobility index under the PIH. If income follows the process (4)-(5), and if the transitory and the permanent shocks are mutually uncorrelated at all leads and lags, the optimal rule is to respond one-to-one to permanent shocks. In case of transitory shocks the optimal rule is instead to revise consumption only by the annuity value of the income innovation. In fact, substituting the income process (4)-(5) in $\psi'_{h,t}$:

$$\Delta \ln c_{h,t} = m_t^{PIH} + \frac{r}{1+r} \varepsilon_{h,t} + \zeta_{h,t} \quad (6)$$

where $m_t^{PIH} = \frac{r}{1+r}\varepsilon_t + \zeta_t$ is the effect of the aggregate permanent and transitory shocks on consumption.

Suppose now that we observe a given cross-sectional distribution of consumption at time $t - 1$ and that the income shocks are not perfectly correlated with the consumption rank of each household in the cross-section. Since aggregate shocks are by definition identical for all households, they do not change each consumer's rank in the consumption distribution and therefore they will not induce any consumption mobility: if they were the sole source of consumption fluctuations the mobility index would be zero. However, other shocks are idiosyncratic, and will move people up and down in the consumption distribution, to an extent that depends on the variance of the two shocks. But since the impact of transitory shocks is scaled down by the factor $\frac{r}{1+r}$, we expect the variance of the permanent shocks to have the greatest impact on mobility. The purpose of the simulations in the next section will be precisely to assess the amount of mobility that one should expect in the permanent income model for given parameters of the income process.

The model in this section has been derived assuming quadratic utility. Recent literature has stressed the importance of precautionary saving, i.e. that households respond to income risk (Kimball, 1990). These models do not deliver a closed-form solution. Recent simulation results produced by Carroll (2000) show that with constant relative risk aversion and an income process similar to the one we use, the implication of the permanent income model that transitory income shocks have a negligible impact on consumption still holds true. Permanent shocks, however, have a somewhat lower impact. In fact, with a precautionary saving motive, a permanent income shock reduces the ratio of wealth to permanent income, thus increasing also precautionary saving. Under a wide range of parameter values, Carroll shows that the marginal propensity to consume of a permanent income shock is about 0.9, not too different from that of the quadratic utility model. Therefore, empirically it is difficult to distinguish the certainty equivalence version of the model from a model with precautionary saving on the basis of the marginal propensity to consume out of permanent income shocks. But the main intuition is the same: transitory income shocks should have a negligible impact on consumption.

4.2 The rule-of-thumb model

Let's assume that consumption equals income in each period, i.e.:

$$\ln c_{h,t} = \ln y_{h,t}$$

This model has been often proposed as a simple, yet extreme alternative to the PIH. A way to rationalize the model is by appealing to the presence of binding liquidity constraints in each period. Alternatively, it can be interpreted as a particular case of the keynesian consumption function, i.e. as an upper bound for the sensitivity of consumption to income shocks. We prefer to term this model the rule-of-thumb model because liquidity constrained consumers cannot borrow but can save, and may react differently to positive and negative expected income growth.

Using the income process above the dynamic of consumption is given by:

$$\Delta \ln c_{h,t} = m_t^K + \varepsilon_{h,t} - \varepsilon_{h,t-1} + \zeta_{h,t} \quad (7)$$

where $m_t^K = \Delta \varepsilon_t + \zeta_t$ is the effect of the aggregate shocks on consumption in the rule-of-thumb model. According to the rule-of-thumb model the growth rate of consumption is therefore equally affected by current and lagged transitory shocks and by permanent shocks. The main difference with the PIH is that in the rule-of-thumb model transitory shocks impact one-to-one on consumption. It is precisely for this reason that in the rule-of-thumb model one should expect more consumption mobility than under the permanent income rule.

4.3 Consumption insurance

To illustrate the implications of the theory of intertemporal choice with complete insurance markets, let us proceed on the assumption that households have preferences of the CRRA type, $u(c) = (1 - \gamma)^{-1} c^{1-\gamma}$, where γ^{-1} is the intertemporal rate of substitution. The implications of the model are identical for any power utility function. As shown by Deaton (1997), the optimal transition law for consumption with complete markets can be obtained by assuming that there is a social planner who maximizes a weighted sum of individual households' utilities. The Lagrangian of this problem can be written as:

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

where h , s and t are subscripts for 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 the realization of state s in period t , and $C_{s,t}$ aggregate consumption in state s and period t .

The first order condition can be written in logarithms as:

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

To obtain the growth rate of consumption, subtract side-by-side from the same expression at time $t - 1$:

$$\Delta \ln c_{h,t} = -\gamma^{-1} \Delta \ln \mu_t + \gamma^{-1} \Delta \ln \pi_t \equiv m_t^{CI} \quad (8)$$

where we drop the subscript s because only one state is realized in each period. The two terms on the right-hand-side of equation (8) represent genuine aggregate effects. The first term 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 (μ and π in equation 8 are not indexed by h).

Equation (8) states that the growth rate of consumption of each household is the same. This implies that the initial cross-sectional distribution of consumption is a sufficient statistic to describe all future distributions. Since all households experience the same consumption growth rate, their rank in the consumption distribution 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.

This approach to test for consumption insurance can be contrasted with the standard procedure of regression analysis. Cochrane (1991) and Townsend (1994) 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.). The implication of the theory is that none of these shocks should impact household consumption growth, as in equation (8). Focussing instead on the prediction that consumption insurance implies absence of consumption mobility has the advantages that we need not identify any of these shocks, and that we need not assume that they are uncorrelated with unobservable or omitted preference shocks, including household fixed effects.

To summarize, consumption mobility is zero in the consumption insurance model, intermediate according to the permanent income model, and highest according to the rule-of-thumb model. This proposition is formally proved in the Appendix.

4.4 Measurement error

The consumption transition law is derived assuming that there is no measurement error in consumption. In practice the index could potentially be upward biased by reporting errors. If respondents report their consumption with errors, one will find units moving up and down even if their true rank in the consumption distribution is unchanged; hence, in the presence of measurement error affect consumption dynamics and the mobility index in Section 2 will tend to report higher mobility. Here we derive the consumption dynamics of the three models considered above in the presence of measurement error in consumption.

Suppose that true consumption is measured with a multiplicative error:

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

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

where $\ln c^*$ is measured consumption and v is an independently and identically normally distributed measurement error. The consumption dynamics in the three models changes in the following way:

$$\Delta \ln c_{h,t}^* = m_t^{PIH} + \frac{r}{1+r} \varepsilon_{h,t} + \zeta_{h,t} + v_{h,t} - v_{h,t-1} \quad (11)$$

$$\Delta \ln c_{h,t}^* = m_t^K + \varepsilon_{h,t} - \varepsilon_{h,t-1} + \zeta_{h,t} + v_{h,t} - v_{h,t-1} \quad (12)$$

$$\Delta \ln c_{h,t}^* = m_t^{CI} + v_{h,t} - v_{h,t-1} \quad (13)$$

The three equations show that measurement error induces a further reason for consumption to vary. Clearly, not only consumption dynamics changes, but the implied consumption mobility as well.⁷ The correct approach then becomes that of calibrating mobility indexes in the three models allowing for realistic impact of measurement error. Define α as the fraction of the cross-sectional standard deviation of measured log consumption that is contaminated by measurement error. This fraction ranges from 0 in the absence of measurement error to 1 when the variability of measured consumption is entirely explained by measurement error. To get a feeling for how measurement error affects consumption mobility in the three models, in the simulations we experiment with a wide range of realistic values for α .

⁷The clearest case in which this happens is in the model with consumption insurance: in the absence of measurement error there is absolutely no mobility in the consumption distribution.

5 Calibration

We now calibrate and simulate the mobility indexes for the three models at issue. One complication with the panel we use is that while income and consumption refer to calendar years, data are collected every other year from 1987 to 1995. The simulated transition laws for consumption must therefore be slightly modified to tackle this problem. In particular, in the case in which there is reporting error we obtain:

$$\begin{aligned} \ln c_{h,t}^* &= \ln c_{h,t-2}^* + \frac{r}{1+r} (\varepsilon_{h,t} + \varepsilon_{h,t-1}) \\ &\quad + (\zeta_{h,t} + \zeta_{h,t-1}) + (v_{h,t} - v_{h,t-2}) \end{aligned} \quad (14)$$

$$\ln c_{h,t}^* = \ln c_{h,t-2}^* + (\varepsilon_{h,t} - \varepsilon_{h,t-2}) + (\zeta_{h,t} + \zeta_{h,t-1}) + (v_{h,t} - v_{h,t-2}) \quad (15)$$

$$\ln c_{h,t}^* = \ln c_{h,t-2}^* + (v_{h,t} - v_{h,t-2}) \quad (16)$$

Since aggregate shocks do not affect consumption mobility, for notational simplicity the above equations omit the aggregate component m_t^j . However, in the estimation of the income process we control for aggregate shocks by introducing time dummies in the regression.

The distinction between the three models is useful but perhaps too stylized for empirical applications. Consumption insurance is no less unrealistic than assuming that all income is consumed in each period, or that all households follow exactly the PIH. In the simulations we therefore nest the three models as:

$$\begin{aligned} \ln c_{h,t}^* &= \ln c_{h,t-2}^* + \frac{\lambda+r}{1+r} \varepsilon_{h,t} + \frac{(1-\lambda)r}{1+r} \varepsilon_{h,t-1} - \lambda \varepsilon_{h,t-2} \\ &\quad + \zeta_{h,t} + \zeta_{h,t-1} + v_{h,t} - v_{h,t-2} \end{aligned} \quad (17)$$

The parameter λ provides enough flexibility as to allow departures from the stylized models of intertemporal choice. Under consumption insurance income shocks play no role, and the consumption dynamics is driven only by measurement error (see equation 16). Note also that if $\lambda = 0$ the expression reduces to the PIH (equation 14). If instead $\lambda = 1$ one obtains the extreme rule-of-thumb case in which consumption equals income each period (equation 15). But if $0 < \lambda < 1$, we obtain intermediate cases between the PIH and the rule-of-thumb models. The reason is that λ represents the extent to which consumption responds to current income over and above the amount that is warranted by the PIH, i.e. the degree of excess sensitivity of consumption to current and past income shocks. One way to interpret this parameter is that each household sets consumption equal to income with probability λ

(perhaps because of binding liquidity constraints) and follows the PIH with probability $(1 - \lambda)$. According to this interpretation, λ measures the fraction of rule-of-thumb consumers in our simulations.

The required parameters for the calibration exercise are the variances of the permanent and transitory income shocks, the amount of measurement error in consumption, the degree of excess sensitivity and the real interest rate. As for the interest rate, we assume a value of 2 percent throughout. As far as excess sensitivity is concerned, we experiment with values of λ spanning from 0 to 1. The simulation also requires assumptions about the value of α , the fraction of the cross-sectional standard deviation of measured log consumption that is contaminated by measurement error. In the absence of validation studies, we experiment with a wide range of values for α , from the highly implausible case of no measurement error ($\alpha = 0$) to an upper bound of $\alpha = 0.20$, i.e. assuming that 20 percent of the variability in consumption is pure noise.

As explained above, we specify the income process as $\ln y_{h,t} = d_t + \beta X_{h,t} + p_{h,t} + \varepsilon_{h,t}$, where $y_{h,t}$ is per-capita family disposable income and d_t a set of time dummies. Using the 1987-95 panel, we regress $\ln y_{h,t}$ on a set of demographic variables (North, South, a dummy for gender, a fourth-order age polynomial, and education dummies) and time dummies, so to remove the deterministic component of income. We save the residuals $u_{h,t} = p_{h,t} + \varepsilon_{h,t}$ and carefully examine their covariance properties. We estimate covariances using equally weighted minimum distance methods, as suggested by Altonji and Segal (1997).⁸

We find that the estimated covariances are consistent with the income process in equations (4) and (5), i.e. that there is a random-walk permanent component and a serially uncorrelated transitory shock. Recall that because of the sample design of the SHIW we can only construct the covariance matrix for two years apart income residuals, $u_{h,t} - u_{h,t-2} = \zeta_{h,t} + \zeta_{h,t-1} + \varepsilon_{h,t} - \varepsilon_{h,t-2}$. To check the consistency of the estimated income process with the model in equations (4) and (5), note that the income process implies the following testable restrictions on the covariance matrix of the first difference of the income residuals:

$$\begin{aligned} E \left[(u_{h,\tau} - u_{h,\tau-2})^2 \right] &= 2\sigma_{\zeta}^2 + 2\sigma_{\varepsilon}^2 \\ E \left[(u_{h,\tau} - u_{h,\tau-2})(u_{h,\tau-2} - u_{h,\tau-4}) \right] &= -\sigma_{\varepsilon}^2 \end{aligned}$$

⁸Covariances can be estimated by equally weighted minimum distance or optimal minimum distance. As shown by Altonji and Segal (1997), the latter can produce inconsistent estimates in small samples, so we adopt the former.

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

Provided that the restrictions are met in the data, one can estimate the variance of the transitory shock σ_ε^2 from the first order autocovariance of income residuals and the variance of the permanent shock σ_ζ^2 combining information on the variance and the first-order autocovariance of the residuals. We find that the estimated autocovariance at the second order is very small (-0.0056) and not statistically different from zero (a t -statistic of -1.1); the autocovariance at the third order is again small (-0.0178) and not statistically different from zero (a t -statistic of -1.1). In contrast, the first order autocovariance (which is an estimate of $-\sigma_\varepsilon^2$) is precisely estimated (a t -statistics of 6.4) at -0.0794. The estimate of the overall variance ($2\sigma_\zeta^2 + 2\sigma_\varepsilon^2$) is 0.2122 (with a t -statistics of 19.4), so we infer that $\sigma_\zeta^2 = 0.0267$ and $\sigma_\varepsilon^2 = 0.0794$.⁹ These parameter estimates are broadly consistent with the evidence available for the US, where researchers have found variances of similar magnitude.¹⁰

Next, we simulate consumption mobility. In each year we choose a sample size identical to the number of actual sample transitions (for instance, it is 2,982 in 1991-93 and 3,211 in 1993-95). We then generate normal distributions for the transitory and the permanent shocks. The shocks are assumed to be normally distributed with mean zero and variances $\sigma_\varepsilon^2 = 0.0794$ and $\sigma_\zeta^2 = 0.0267$, respectively. Measurement errors at times t and $t-2$ are drawn from a normal distribution with mean zero and standard deviation α times the standard deviation of measured log consumption at t and $t-2$.¹¹ True consumption $\ln c_{h,t-2}$ is drawn from a normal distribution with mean equal to the mean of measured consumption and standard deviation equal to $(1-\alpha)$ times the standard deviation of measured consumption at $t-2$. Finally, we generate $\ln c_{h,t}^*$ for each household using equation (17) assuming different values for α and λ and calculate the associated mobility index. This exercise is replicated 100 times for each sample period.

6 Simulating consumption mobility

The simulated mobility index is reported in Table 4 for different values of the α and λ parameters. The qualitative pattern of results is similar in all

⁹Unfortunately, with data collected every two years we cannot distinguish between this income process and one where the transitory component is an MA(1) process.

¹⁰For instance, Carroll and Samwick (1997) using the PSID, estimate $\sigma_\zeta^2 = 0.0217$ and $\sigma_\varepsilon^2 = 0.0440$.

¹¹By construction, the normality of the income shocks and of measurement error generates a symmetric transition matrix for consumption. This feature of the simulations is consistent with the symmetry of the empirical matrix documented in Table 2.

periods, suggesting that the simulated index is only marginally affected by the initial distribution of consumption (the income process and the associated variances of the shocks are in fact assumed to be the same across the different samples). For brevity, we report the results only for the most recent period (1993-95), that also features the largest number of transitions. The stars in the table denote that the simulated mobility index is statistically different from the actual index (0.4429 in 1993-95, see Table 3) at the 1 percent level.

To understand the simulation results it is useful to start with $\alpha = 0$ (no measurement error). Consider first the three benchmark cases of the PIH ($\lambda = 0$), rule-of-thumb ($\lambda = 1$) and consumption insurance (no effect of transitory or permanent shocks). Income shocks have the largest impact on the consumption distribution in the rule-of-thumb model (mobility is 55 percent), intermediate in the case of the permanent income model (40 percent) and no impact under consumption insurance. The ranking of the models in the terms of predicted mobility agrees with intuition because idiosyncratic income shocks translate into consumption changes entirely in the rule-of-thumb model, partially in the PIH via intertemporal smoothing, and are fully insured in the risk sharing model. It is encouraging to note that even in the benchmark models the index simulated under the PIH is closer to the empirical index than any of the two other models. However, from a statistical point of view, it is clear that none of the models is fully able to replicate the actual degree of mobility we observe in the data. Each model rejects the hypothesis that the simulated index equals the empirical index at the 1 percent significance level.

One way to explain the difference between simulated and empirical mobility is to introduce measurement error in consumption ($\alpha > 0$). Measurement error always increases simulated mobility, regardless of the model considered. In the PIH and rule-of-thumb models mobility increases from 40 to 48 percent, and from 55 to 60 percent, respectively. With consumption insurance, the highest value of α is able to generate at most 28 percent mobility. These results indicate that under consumption insurance even very large and implausible values of measurement error cannot reproduce actual mobility. The rule-of-thumb model already overestimates mobility when $\alpha = 0$, so that introducing measurement error will only shift the simulated index further away from the empirical one.¹² Once we allow for measurement error, the permanent income model can instead replicate quite well the amount of mobility we have in the data. The t -statistic associated with the test that the

¹²Consumption insurance can be made consistent with the data only assuming highly unlikely values of $\alpha = 0.35$. The rule-of-thumb model is always rejected, regardless of the size of measurement error.

actual mobility index is different from that simulated under the permanent income model does not reject the null hypothesis for values of $\alpha = 0.10$ and $\alpha = 0.15$.¹³

An alternative way to reconcile the data with the simulation results is to depart from the PIH by allowing a higher response to income shocks than predicted by the PIH. Raising the excess sensitivity parameter λ increases consumption mobility, regardless of the size of the measurement error. For instance, setting $\alpha = 0.05$, the simulated index of mobility is not statistically different from actual mobility for values of λ between 0.2 and 0.3. Glancing through Table 4, one can see that there is a region of parameter values α and λ for which the null hypothesis of equality between the simulated and actual index is not rejected (in italics). Furthermore, the simulated index is only marginally affected when we set the marginal propensity to consume out of permanent shocks at 0.9, as predicted by versions of the PIH with precautionary saving.

Since on statistical grounds we cannot exclude any value in this region, we face an apparently difficult identification problem. Unfortunately, we cannot rely on independent evidence to gather realistic values for the fraction of rule-of-thumb consumers and for the amount of measurement error in consumption. As for the latter, we already noted the absence of validation studies. The excess sensitivity parameter has sometimes been inferred from the income growth coefficient in Euler equations estimates. In previous work we perform such test in the same panel dataset we use in this study. We find an excess sensitivity coefficient of 0.32 with a t-statistic of 5 (Jappelli and Pistaferri, 2000). However, the result depends heavily on the choice of the instrument set, and vanishes if one controls for the structure of the forecast error of the Euler equation. In short, there is no way to benchmark the values of α and λ tightly with outside evidence. In our view a reasonable and balanced interpretation of the data is that they are generated by a model with intermediate values (within the “acceptance” region of Table 4) of measurement error ($\alpha = 0.10$) and excess sensitivity of consumption to income shocks ($\lambda = 0.20$).

The model simulated with $\alpha = 0.10$ and $\lambda = 0.20$ predicts almost per-

¹³Results for other years are similar with the exception of 1991-93. In that period actual mobility increases to 50 percent, a fact that is not captured by our simulations. One possible explanation is that the variance of the permanent shock, which is assumed to be time stationary, changed in 1993 due to the unprecedented strong recession. However, we cannot rule out that in 1993 the amount of measurement error is greater than in the other two years. Another possibility is that the 1993 recession impacted unevenly on households, a particular form of non-stationarity that we neglect in our simulation exercise.

fectly the actual transition matrix, not just the aggregate mobility index. In Table 5 we report the simulated transition probabilities and (in parenthesis) the actual transition probabilities, the same reported for 1993-95 in Table 2. The comparison between the two sets of numbers is striking: regardless of cell, the difference between the actual and simulated values is at most 2 percent. A battery of formal tests of the hypothesis that the simulated transition probability is equal to the actual probability is not rejected for any of the 16 values in Table 5. Other parameter values in the “acceptance” region of Table 4 yield qualitatively similar results.

6.1 Group heterogeneity

As a further test of the ability of the simulations in explaining consumption transitions, we check if the differences in mobility across specific population groups can be replicated by the simulations. The PIH suggests that if different population groups are systematically exposed to different idiosyncratic shocks (and therefore face different income processes), consumption mobility should differ across groups in a predictable way. We choose to focus on two education groups (compulsory schooling or less and high school or college degree), and two generations (born before and after 1940). Education and year of birth are exogenous characteristics by which one can partition the sample, and there is wide evidence that different education groups and generations face different earnings opportunities and uncertainties. As we shall see below, also in our sample the income generating process differs by education and year of birth, thus providing an ideal setting to test the validity of models of intertemporal choice. One further reason to focus on education and year of birth is that they are likely to be correlated with variables affecting preferences and therefore with different consumption behavior.

The upper part of Table 6 compares the income process and consumption mobility by education. We run the income regressions separately for households headed by individuals with high and low education. We then estimate the autocovariance matrix as explained in Section 5, and find $\sigma_\zeta^2 = 0.0296$ and $\sigma_\varepsilon^2 = 0.0754$ for the less well educated, and $\sigma_\zeta^2 = 0.0198$ and $\sigma_\varepsilon^2 = 0.0895$ for those with at least a high school degree. Overall, the estimated variances signal that the less well educated face a higher variance of permanent income shocks, a pattern also uncovered by Carroll and Samwick (1997) with US data.

Computing the transition matrix for consumption shows that the less well educated are more mobile than the high educated, the difference between the two groups being 3.8 percentage points. We apply the statistic on difference of means outlined in Section 2, and reject the hypothesis that the two in-

dexes are equal at the 1 percent level. The simulated index of mobility with $\alpha = 0.10$ and $\lambda = 0.20$ is almost identical to actual mobility in each group, implying that the model replicates quite well also the difference in actual and simulated mobility between the two groups.

In contrast, the rule-of-thumb model (or models with high values of excess sensitivity) does not predict any difference between education groups. The lower variance of the transitory shock for the less well educated is compensated by a higher variance of the permanent shock, resulting in the same mobility rates in the two groups. Similarly, the model with consumption insurance, is by assumption unable to explain differences in consumption dynamics emerging from idiosyncratic shocks (whether transitory or permanent).

The lower part of Table 6 refers to those born before and after 1940. Also in this case, we find wide differences in the parameters of the income process. The old generation exhibits a higher variance of the permanent shock and a lower variance of the transitory shock ($\sigma_{\zeta}^2 = 0.0305$ and $\sigma_{\varepsilon}^2 = 0.0507$) than the young generation ($\sigma_{\zeta}^2 = 0.0217$ and $\sigma_{\varepsilon}^2 = 0.1125$). This is consistent with the largely held view that the young generation experiences higher labor market uncertainties than the old generation, as a reflection of transitory shocks. On the other hand, permanent shocks impact more people towards the end of their career (because of permanent layoffs, quits, or health-related problems). The difference in consumption mobility between the two groups is 4.5 percentage points, again statistically different from zero at the 1 percent level. The simulated mobility index is not statically different from the actual index for the two groups, and reproduces the pattern of observed mobility quite well (3.0 percentage points against the actual value of 4.5).

In contrast to the simulated mobility rates by year of birth, the rule-of-thumb model (or models with high values of λ) predicts lower mobility for the old (0.55 against 0.59), a pattern opposite to what we find in the data (0.46 against 0.42). The reason is in the rule-of-thumb model transitory shocks have the largest impact on consumption, and σ_{ε}^2 is twice as high for the young. The model with consumption insurance, once again, predicts no difference between the two groups. Overall, we take the results of these experiments as strong evidence that the permanent income model with realistic values of measurement error and a moderate amount of excess sensitivity can replicate remarkably well the pattern of actual mobility not only in the aggregate, but also by selected population groups that differ in the income generating process.

6.2 Measurement error in income

It is important to consider the robustness of our conclusions in the presence of measurement error in income. This error inflates the variance of the transitory shock but does not affect the variance of the permanent shock. To see this point, assume that true income is measured with a multiplicative error: $\ln y_{h,t}^* = \ln y_{h,t} + \omega_{h,t}$, where $\omega_{h,t}$ is an independently and identically normally distributed measurement error with mean zero and variance σ_ω^2 . Using the income process (4)-(5): $\ln y_{h,t}^* = \beta X_{h,t} + p_{h,t} + \varepsilon_{h,t} + \omega_{h,t}$, the two years apart income residual is now: $u_{h,t} - u_{h,t-2} = \zeta_{h,t} + \zeta_{h,t-1} + \varepsilon_{h,t} - \varepsilon_{h,t-2} + \omega_{h,t} - \omega_{h,t-2}$. The covariance matrix of the first difference of the income residuals depends now on the variance of the measurement error:

$$\begin{aligned} E \left[(u_{h,\tau} - u_{h,\tau-2})^2 \right] &= 2\sigma_\zeta^2 + 2\sigma_\varepsilon^2 + 2\sigma_\omega^2 \\ E \left[(u_{h,\tau} - u_{h,\tau-2})(u_{h,\tau-2} - u_{h,\tau-4}) \right] &= -\sigma_\varepsilon^2 - \sigma_\omega^2 \\ E \left[(u_{h,\tau} - u_{h,\tau-2})(u_{h,\tau-j} - u_{h,\tau-j-2}) \right] &= 0 \text{ for all } j \geq 4 \end{aligned}$$

However, it can be checked that measurement error inflates the estimated variance of the transitory shock by σ_ω^2 , but not the variance of the permanent shock σ_ζ^2 , which is still identified by the difference between the variance and (minus twice) the first-order autocovariance. The conclusion is that even though the estimate of the variance of the permanent shock is unaffected by serially uncorrelated measurement error, the estimate of the variance of the transitory shock is not.

This implies that in the model with full consumption insurance, idiosyncratic income shocks play no role regardless of measurement error in income. In the permanent income model, the impact of measurement error in income is bound to be small, because transitory shocks play a very limited role. In contrast, measurement error may have a large impact in the rule-of-thumb model. Since we cannot identify σ_ω^2 from the data, we repeat our simulation: (a) dropping the self-employed from the sample on which we estimate the income process,¹⁴ and (b) downsizing the variance of the transitory shock, i.e. assuming that one third or one half of the estimated first-order autocovariance reflects measurement error.

The results of these experiments are very similar to the simulations reported in Tables 4, 5 and 6 and are not reported for brevity. Excluding the self employed, in all years the rule-of-thumb model and the model with consumption insurance yield mobility indexes similar to the simulations presented in

¹⁴Brandolini and Cannari (1994) note that in the SHIW income from self-employment is less well estimated than wages or salaries.

Table 4, and the hypothesis that they are equal to the actual mobility index is soundly rejected. On the other hand, simulated mobility with $\lambda = 0.2$ and $\alpha = 0.10$ is quite close to actual mobility and we cannot reject the hypothesis that it is equal to the actual mobility index. Similarly, assuming that measurement error in income is one third or half of the first-order autocovariance has negligible effects on simulated mobility.

7 Conclusions

The implications of the theories of intertemporal consumption choice for consumption mobility are as yet unexplored. In this paper we study transition probabilities for total non-durable consumption using the 1987-95 panel contained in the Bank of Italy Survey of Household Income and Wealth. We then summarize the transition matrix of consumption by appropriate mobility indexes and find that there is substantial consumption mobility: in any year, about 50 percent of the households moves up or down in the consumption distribution.

In the remainder of the paper we attempt to understand which model of intertemporal consumption choice is capable of explaining the amount of consumption mobility we observe in the data. We consider three popular alternative models of intertemporal consumption choice: the model with full consumption insurance, the rule-of-thumb model, and the PIH. These models provide the clearest implications for consumption mobility.

The model with complete insurance markets, where all idiosyncratic income shocks are insured, implies that in any time period the initial cross-sectional distribution of consumption is a sufficient statistic for all future distributions, and therefore no consumption mobility in the absence of measurement error in consumption. The rule-of-thumb model is one where income shocks have the greatest impact on consumption and therefore generates substantial consumption mobility. The permanent income model is one in which households react only to permanent income shocks. Thus, the degree of mobility predicted by the model is intermediate between the two other models.

We carefully parametrize an income process to distinguish between transitory and permanent shocks and use the estimated parameters to simulate theoretically the degree of mobility stemming from each of the three consumption models. We then compare them statistically with the actual amount of mobility estimated in the data. Overall, the simulations provide strong support in favor of a slightly modified version of the permanent income

model that allows for realistic values of measurement error in consumption and a moderate amount of excess sensitivity to income shocks. First of all, the aggregate mobility index generated under this model is quite close to the actual one, contrary to the other two models; second, the permanent income hypothesis is the only model that is able to match the actual transition matrix cell by cell; finally, and most importantly, the permanent income hypothesis captures the different patterns of consumption mobility across education and year of birth groups, while other models predict either no differences or an opposite pattern. The results are robust with respect to different definitions of consumption (in per capita or per adult equivalent terms) and to the presence of measurement error in income.

References

- [1] Altonji, Joseph G., and Lewis M. Segal (1996), “Small-sample bias in GMM estimation of covariance structures”, *Journal of Business and Economic Statistics* **14**, 353-66.
- [2] Atkinson, Anthony B., Francois Bourguignon and Christian Morrisson (1992), *Empirical studies of earnings mobility*. Philadelphia: Harwood Academic.
- [3] Attanasio, Orazio P., and Steve Davis (1996), “Relative wage movements and the distribution of consumption”, *Journal of Political Economy* **104**, 1227-62.
- [4] Banerjee, Abhijit V. and Andrew F. Newman (1991), “Risk-bearing and the theory of income distribution”, *Review of Economic Studies* **58**, 211-35.
- [5] Brandolini, Andrea (1999), “The distribution of personal income in post-war Italy: source description, data quality, and the time pattern of income inequality”, *Tem di Discussione n. 350*. Rome: Bank of Italy.
- [6] 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.
- [7] Campbell, John Y., and Gregory N. Mankiw (1989), “Consumption, income, and interest rates: reinterpreting the time-series evidence,” in *NBER Macroeconomics Annual 1989*, ed. by Olivier J. Blanchard and Stanley Fischer, 185-216. Cambridge: MIT Press.
- [8] Carroll, Christopher D. (2000), “Precautionary saving and the marginal propensity to consume out of permanent income,” mimeo, John Hopkins University.
- [9] Carroll, Christopher D., and Andrew A. Samwick (1997), “The nature of precautionary wealth”, *Journal of Monetary Economics* **40**, 41-72.

- [10] Cochrane, John (1991), “A simple test of consumption insurance”, *Journal of Political Economy* **99**, 957-976.
- [11] Cutler, David M., and Lawrence F. Katz (1992), “Rising inequality? Changes in the distribution of income and consumption in the 1980s”, *American Economic Review* **82**, 546-51.
- [12] Deaton, Angus (1997), *The analysis of household surveys*. Baltimore: The Johns Hopkins University Press.
- [13] Deaton, Angus, and Christina Paxson (1994), “Intertemporal choice and inequality”, *Journal of Political Economy* **102**, 384-94.
- [14] Hall, Robert E. (1978), “Stochastic implications of the life-cycle permanent income hypothesis: theory and evidence”, *Journal of Political Economy* **96**, 971-87.
- [15] Hall, Robert E., and Frederic Mishkin (1982), “The sensitivity of consumption to transitory income: evidence from PSID households,” *Econometrica* **50**, 461-81.
- [16] Jappelli, Tullio, and Luigi Pistaferri (2000), “Using subjective income expectations to test for the excess sensitivity of consumption to expected income growth,” *European Economic Review* **44**, 337-58.
- [17] Johnson, David, and Stephanie Shipp (1997), “Trends in inequality using Consumption Expenditures: the US from 1960 to 1993”, *Review of Income and Wealth* **43**, 133-151.
- [18] Kimball, Miles S. (1990), “Precautionary saving in the small and in the large”, *Econometrica* **58**, 53-73.
- [19] Schluter, Christian (1998), “Statistical inference for mobility indexes”, *Economic Letters* **59**, 157-62.
- [20] Shorrocks, Anthony F. (1978), “The measurement of mobility”, *Econometrica* **46**, 1013-24.
- [21] Silverman, B.W. (1986), *Density estimation for statistics and data analysis*, Chapman and Hall.
- [22] Townsend, Robert (1994), “Risk and insurance in village India”, *Econometrica* **62**, 539-91.

A Appendix

Recall the three distribution laws for log consumption:

$$\begin{aligned}\ln c_{h,t+1} &= \ln c_{h,t} + \zeta_{h,t+1} + \frac{r}{1+r} \varepsilon_{h,t+1} \\ \ln c_{h,t+1} &= \ln c_{h,t} + \zeta_{h,t+1} + \varepsilon_{h,t+1} - \varepsilon_{h,t} \\ \ln c_{h,t+1} &= \ln c_{h,t}\end{aligned}$$

respectively in three cases of the PIH, the rule-of-thumb model, and consumption insurance. Recall that ζ is the permanent shock and ε the transitory shock. Without loss of generality, we set aggregate consumption growth to zero.

Divide the distribution in quantiles. Denote with q_{j-1} and q_j two successive quantiles of the distribution ($q_{j-1} < q_j$), and assume that:

$$\begin{pmatrix} \zeta_{h,t+1} \\ \varepsilon_{h,t+1} \\ \varepsilon_{h,t} \end{pmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\zeta^2 & 0 & 0 \\ 0 & \sigma_\varepsilon^2 & 0 \\ 0 & 0 & \sigma_\varepsilon^2 \end{pmatrix} \right]$$

Proposition: Consumption mobility is highest in the rule-of-thumb model, intermediate in the PIH and zero in the consumption insurance model.

Proof: Start from the consumption insurance model and note that:

$$\begin{aligned}\Pr(q_{j-1} < \ln c_{h,t+1} < q_j \mid q_{j-1} < \ln c_{h,t} < q_j, \ln c_{h,t+1} = \ln c_{h,t}) \\ = \Pr(q_{j-1} < \ln c_{h,t} < q_j \mid q_{j-1} < \ln c_{h,t} < q_j) = 1\end{aligned}$$

In the rule-of-thumb model:

$$\begin{aligned}\Pr(q_{j-1} < \ln c_{h,t+1} < q_j \mid q_{j-1} < \ln c_{h,t} < q_j, \ln c_{h,t+1} = \ln c_{h,t} + \zeta_{h,t+1} + \varepsilon_{h,t+1} - \varepsilon_{h,t}) \\ = \Pr(q_{j-1} < \ln c_{h,t} + \zeta_{h,t+1} + \varepsilon_{h,t+1} - \varepsilon_{h,t} < q_j \mid q_{j-1} < \ln c_{h,t} < q_j) \\ = \Pr(q_{j-1} - \ln c_{h,t} < \zeta_{h,t+1} + \varepsilon_{h,t+1} - \varepsilon_{h,t} < q_j - \ln c_{h,t} \mid q_{j-1} - \ln c_{h,t} < 0 < q_j - \ln c_{h,t}) \\ = \Pr\left(\frac{q_{j-1} - \ln c_{h,t}}{\sigma_{ROT}} < \frac{\zeta_{h,t+1} + \varepsilon_{h,t+1} - \varepsilon_{h,t}}{\sigma_{ROT}} < \frac{q_j - \ln c_{h,t}}{\sigma_{ROT}} \mid q_{j-1} - \ln c_{h,t} < 0 < q_j - \ln c_{h,t}\right) \\ = \Phi\left(\frac{q_j - \ln c_{h,t}}{\sigma_{ROT}}\right) - \Phi\left(\frac{q_{j-1} - \ln c_{h,t}}{\sigma_{ROT}}\right) > 0\end{aligned}$$

where $\sigma_{ROT} = \sqrt{\sigma_\zeta^2 + 2\sigma_\varepsilon^2}$, and $\Phi(\cdot)$ is the c.d.f. of the $N(0, 1)$ distribution. The last inequality holds because $q_j - \ln c_{h,t} > 0$, $q_{j-1} - \ln c_{h,t} < 0$, and $\sigma_{ROT} > 0$.

In the PIH:

$$\begin{aligned}
& \Pr \left(q_{j-1} < \ln c_{h,t+1} < q_j \mid q_{j-1} < \ln c_{h,t} < q_j, \ln c_{h,t+1} = \ln c_{h,t} + \zeta_{h,t+1} + \frac{r}{1+r} \varepsilon_{h,t+1} \right) \\
&= \Pr \left(q_{j-1} < \ln c_{h,t} + \zeta_{h,t+1} + \frac{r}{1+r} \varepsilon_{h,t+1} < q_j \mid q_{j-1} < \ln c_{h,t} < q_j \right) \\
&= \Pr \left(q_{j-1} - \ln c_{h,t} < \zeta_{h,t+1} + \frac{r}{1+r} \varepsilon_{h,t+1} < q_j - \ln c_{h,t} \mid q_{j-1} - \ln c_{h,t} < 0 < q_j - \ln c_{h,t} \right) \\
&= \Pr \left(\frac{q_{j-1} - \ln c_{h,t}}{\sigma_{PIH}} < \frac{\zeta_{h,t+1} + \frac{r}{1+r} \varepsilon_{h,t+1}}{\sigma_{PIH}} < \frac{q_j - \ln c_{h,t}}{\sigma_{PIH}} \mid q_{j-1} - \ln c_{h,t} < 0 < q_j - \ln c_{h,t} \right) \\
&= \Pr \left(\frac{q_{j-1} - \ln c_{h,t}}{\sigma_{PIH}} < \frac{\zeta_{h,t+1} + \frac{r}{1+r} \varepsilon_{h,t+1}}{\sigma_{PIH}} < \frac{q_j - \ln c_{h,t}}{\sigma_{PIH}} \mid q_{j-1} - \ln c_{h,t} < 0 < q_j - \ln c_{h,t} \right) \\
&= \Phi \left(\frac{q_j - \ln c_{h,t}}{\sigma_{PIH}} \right) - \Phi \left(\frac{q_{j-1} - \ln c_{h,t}}{\sigma_{PIH}} \right) > 0
\end{aligned}$$

where $\sigma_{PIH} = \sqrt{\sigma_{\zeta}^2 + \left(\frac{r}{1+r}\right)^2 \sigma_{\varepsilon}^2}$. The last inequality holds because $q_j - \ln c_{h,t} > 0$, $q_{j-1} - \ln c_{h,t} < 0$, and $\sigma_{PIH} > 0$. Notice now that $\sigma_{ROT} > \sigma_{PIH}$, so that:

$$\frac{q_{j-1} - \ln c_{h,t}}{\sigma_{PIH}} < \frac{q_{j-1} - \ln c_{h,t}}{\sigma_{ROT}} < 0 < \frac{q_j - \ln c_{h,t}}{\sigma_{ROT}} < \frac{q_j - \ln c_{h,t}}{\sigma_{PIH}}$$

and finally:

$$\Phi \left(\frac{q_j - \ln c_{h,t}}{\sigma_{PIH}} \right) - \Phi \left(\frac{q_{j-1} - \ln c_{h,t}}{\sigma_{PIH}} \right) > \Phi \left(\frac{q_j - \ln c_{h,t}}{\sigma_{ROT}} \right) - \Phi \left(\frac{q_{j-1} - \ln c_{h,t}}{\sigma_{ROT}} \right)$$

The last inequality proves that the probability of remaining in the same quantile of the consumption distribution is greater under the PIH than under the rule-of-thumb model, i.e. that mobility is higher in the latter case. This completes the proof.

Table 1
Descriptive statistics

Cross-sectional means and variances are computed using sample weights. The variables c_t and y_t denote household non-durable consumption and disposable income, respectively. Demographic characteristics refer to the household head.

	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
Gini coefficient of c_t	0.28	0.27	0.29	0.29	0.27	0.28
$\ln y_t$	10.25	10.40	10.36	10.27	10.27	10.32
$var(\ln y_t)$	0.39	0.37	0.37	0.57	0.47	0.45
Gini coefficient of y_t	0.35	0.32	0.32	0.36	0.35	0.34
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
Years of schooling	7.38	7.97	8.19	8.03	8.10	8.03
Less well educated	0.78	0.72	0.70	0.72	0.70	0.72
More educated	0.22	0.28	0.30	0.28	0.30	0.28
Age	52.00	52.52	52.78	53.05	55.03	53.22
Born \leq 1940	0.60	0.58	0.54	0.49	0.49	0.53
Born $>$ 1940	0.40	0.42	0.46	0.51	0.51	0.47
Income recipients	1.63	1.72	1.72	1.74	1.78	1.73
Number of obs.	1,097	2,717	4,036	4,006	3,211	15,067

Table 2
The transition matrix of consumption

The table reports consumption transitions from period $t - 2$ to period t . The generic element of this table is \hat{p}_{ij} , the estimated probability of moving from quartile i in period $t - 2$ to quartile j in period t . Define n_{ij} as the number of households that move from quartile i in period $t - 2$ to quartile j in period t and $n_i = \sum_j n_{ij}$ as the total number of observations in each row i of the transition matrix. The maximum likelihood estimator of the first-order Markov transition probabilities is then: $\hat{p}_{ij} = \frac{n_{ij}}{n_i}$.

		1987-89			
		1989 quartile			
1987 quartile		1 st	2 nd	3 rd	4 th
	1 st	0.71	0.20	0.07	0.02
	2 nd	0.23	0.42	0.27	0.08
	3 rd	0.08	0.29	0.40	0.23
	4 th	0.03	0.09	0.29	0.60

		1989-91			
		1991 quartile			
1989 quartile		1 st	2 nd	3 rd	4 th
	1 st	0.66	0.25	0.07	0.01
	2 nd	0.25	0.41	0.27	0.06
	3 rd	0.10	0.27	0.41	0.25
	4 th	0.01	0.07	0.25	0.68

		1991-93			
		1993 quartile			
1991 quartile		1 st	2 nd	3 rd	4 th
	1 st	0.63	0.26	0.08	0.02
	2 nd	0.23	0.38	0.29	0.09
	3 rd	0.11	0.28	0.37	0.25
	4 th	0.04	0.10	0.26	0.60

		1993-95			
		1995 quartile			
1993 quartile		1 st	2 nd	3 rd	4 th
	1 st	0.68	0.25	0.07	0.01
	2 nd	0.24	0.43	0.26	0.07
	3 rd	0.07	0.27	0.44	0.23
	4 th	0.02	0.06	0.23	0.69

Table 3
Actual mobility index

The table reports the Shorrocks mobility index and the associated standard error and number of transitions for separate sample periods. The index is calculated as: $S(\hat{P}) = \frac{q - \text{trace}(\hat{P})}{q}$, where $q = 4$, $\text{trace}(\hat{P}) = \sum_i \hat{p}_{ii}$, and \hat{p}_{ii} the estimated probability of being in quartile i in both $t - 2$ and t . The standard error is $\sqrt{\frac{1}{q^2} \sum_i \frac{\hat{p}_{ii}(1 - \hat{p}_{ii})}{n_i}}$.

Sample period	$S(\hat{P})$	$s.e.(S(\hat{P}))$	Number of transitions
1987-1989	0.4702	0.0146	1,097
1989-1991	0.4705	0.0110	1,914
1991-1993	0.5029	0.0089	2,982
1993-1995	0.4429	0.0085	3,211

Table 4
Simulated mobility index

The table reports the simulated Shorrocks indexes of mobility for different values of measurement error in consumption and the degree of excess sensitivity to income shocks. The last row reports the simulated mobility index of the consumption insurance model for different values of measurement error in consumption. A star indicates that the simulated index is statistically different from the actual mobility index at the 1 percent level. Values in italics do not reject the hypothesis that the simulated index equals the actual index at the 1 percent level.

Degree of excess sensitivity (λ)	Fraction of measurement error (α)				
	0.00	0.05	0.10	0.15	0.20
0.0	0.4004*	0.4143*	0.4337	0.4568	0.4810*
0.1	0.4031*	0.4189*	0.4370	0.4590	0.4849*
0.2	0.4154*	0.4282	0.4467	0.4689*	0.4936*
0.3	0.4305	0.4432	0.4625*	0.4804*	0.5042*
0.4	0.4478	0.4608*	0.4765*	0.4964*	0.5168*
1	0.5494*	0.5597*	0.5714*	0.5840*	0.5965*
Consumption insurance	0.0000*	0.0610*	0.1286*	0.2014*	0.2780*

Table 5
Simulated and actual
transition matrix of consumption

The table reports the simulated consumption transitions between 1993 and 1995 and, in parenthesis, the actual consumption transitions. The generic element of this table is \hat{p}_{ij} , the estimated probability of moving from quartile i in period $t - 2$ to quartile j in period t . The simulated transitions are computed assuming that measurement error explains 10 percent of the overall variability in measured consumption and that the degree of excess sensitivity is 20 percent.

1993 quartile	1995 quartile			
	1 st	2 nd	3 rd	4 th
1 st	0.69 (0.68)	0.25 (0.25)	0.06 (0.07)	0.01 (0.01)
2 nd	0.25 (0.24)	0.42 (0.43)	0.27 (0.26)	0.06 (0.07)
3 rd	0.06 (0.07)	0.27 (0.27)	0.42 (0.44)	0.25 (0.23)
4 th	0.01 (0.02)	0.06 (0.06)	0.25 (0.23)	0.69 (0.69)

Table 6
Actual and simulated mobility for different groups

The table reports the parameters of the income process, and the actual and simulated Shorrocks index of consumption mobility for two education groups (compulsory schooling or less, and high school or college) and two generations (those born before and after 1940, respectively). The transition matrix refers to the 1993-95 period. The simulations assume that measurement error in consumption is 10 percent of the overall variability in measured consumption and that the degree of excess sensitivity is 20 percent.

	Low education	High education	Difference
Variance of permanent shock	0.0296	0.0198	0.0098
Variance of transitory shock	0.0754	0.0895	-0.0141
Actual mobility index	0.4631	0.4249	0.0381
Simulated mobility index	0.4676	0.4260	0.0416
	Born \leq 1940	Born $>$ 1940	Difference
Variance of permanent shock	0.0305	0.0217	0.0088
Variance of transitory shock	0.0507	0.1125	-0.0618
Actual mobility index	0.4592	0.4139	0.0453
Simulated mobility index	0.4635	0.4335	0.0300