

CONSUMPTION INEQUALITY AND PARTIAL INSURANCE*

Richard Blundell

University College London and Institute for Fiscal Studies.

Luigi Pistaferri

Hoover Institution, Stanford University and CEPR.

Ian Preston

University College London and Institute for Fiscal Studies.

September 2003

Abstract

This paper uses panel data on household consumption and income to evaluate the degree of insurance to income shocks. Our aim is to describe the transmission of income inequality into consumption inequality by contrasting shifts in the cross-sectional distribution of income growth with shifts in the cross-sectional distribution of consumption growth. We combine panel data on income from the PSID with consumption data from repeated CEX cross-sections. The results point to some partial insurance but reject the complete market restrictions. We find a greater degree of insurance for transitory shocks and differences in the degree of insurance over time and across demographic groups. We also document the importance of durables and of taxes and transfers as a means of insurance.

Key words: Consumption, Insurance, Inequality.

JEL Classification: D52; D91; I30.

*This paper is a revised version of "Partial Insurance, Information and Consumption Dynamics". We thank Joe Altonji, Orazio Attanasio, Erich Battistin, David Johnson, Arie Kapteyn, John Kennan, Robert Lalonde, Hamish Low, Bruce Meyer, Ken West, seminar participants in Northwestern, Oxford, Wisconsin-Madison, Berkeley, Maryland, the Federal Reserve Board, UCLA, Carlos III, Chicago, the 2002 SED Conference in New York, and the NBER's 26th Annual International Seminar on Macroeconomics for useful comments, and Cristobal Huneus for able research assistance. The paper is part of the program of research of the ESRC Centre for the Microeconomic Analysis of Public Policy at IFS. Financial support from the ESRC, the Joint Center for Poverty Research/Department of Health and Human Services and the National Science Foundation (under grant SES-0214491) is gratefully acknowledged. All errors are ours.

1 Introduction

Under complete markets agents can sign contingent contracts providing full insurance against idiosyncratic shocks to income. Moral hazard and asymmetric information, however, make these contracts hard to implement, and in fact they are rarely observed in reality. Even a cursory look at consumption and income data reveals the weakness of the complete markets hypothesis. Thus volatility of individual consumption is much higher than the volatility of aggregate consumption, a fact against full insurance [Aiyagari, 1994]. Moreover, there is a substantial amount of mobility in consumption [Jappelli and Pistaferri, 2003]. Indeed, Blundell and Preston [1998] use the growth in consumption inequality over the 1980s in the U.K. to identify growth in permanent (uninsured) income inequality. Formal tests of the complete markets hypothesis [see Attanasio and Davis, 1996], find that the null hypothesis of full consumption insurance is soundly rejected.¹ Attempts to salvage the theory by allowing for risk sharing within the family and no risk sharing among unrelated families have also failed [Hayashi, Altonji and Kotlikoff, 1996].

In the textbook permanent income hypothesis the only mechanism available to agents to smooth income shocks is personal savings. The main idea is that people attempt to keep the expected marginal utility of consumption stable over time. Since insurance markets for income fluctuations are assumed to be absent, the marginal utility of consumption is not stabilized across states. If income is shifted by permanent and transitory shocks, self-insurance through borrowing and saving may allow intertemporal consumption smoothing against the latter but not against the former. This is simply because one cannot borrow to smooth out a permanent income decline without violating the budget constraint, so that permanent shocks to income will be permanent shocks to consumption.² In this paper we start from the premise of some, but not necessarily full, insurance and consider the importance of distinguishing between transitory and permanent shocks.

Models that feature a myriad of markets and those that allow for just personal savings as a smoothing mechanism are clearly extreme characterization of individual behavior and of the economic environment faced by the consumers. Deaton and Paxson [1994] notice this and envision

¹Notable exceptions are Altug and Miller [1990] and Mace [1991]. In these papers, however, failure to reject the null hypothesis of full consumption insurance is likely to be due to econometric and sample selection issues. See, e.g., Nelson [1994].

²Even with precautionary saving, permanent shocks to labour income will typically be almost fully transmitted into consumption (see below).

“the construction and testing of market models under partial insurance”, while Hayashi, Altonji and Kotlikoff [1996] call for future research to be “directed to estimating the extent of consumption insurance over and above self-insurance”. In this paper we address the issue of whether partial consumption insurance is available to agents and estimate the degree of insurance over and above self-insurance through savings. We do this by contrasting shifts in the distribution of income growth with shifts in the distribution of consumption growth, and analyze the way these two measures of household welfare correlate over time. Our research is related to other papers in the literature, particularly Hall and Mishkin [1982], Altonji, Martins and Siow [2002], Deaton and Paxson [1994], and Blundell and Preston [1998].

Using data from a combination of the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX), we document a number of key findings. We find a strong growth in permanent income shocks during the early 1980s. The variance of permanent shocks thereafter levels off. The variance of consumption mimics these trends very closely. We find strong evidence against full insurance for permanent income shocks but not for transitory income shocks, except for a low income subsample where transitory shocks seem, unsurprisingly, less insurable. Further there is evidence of some partial insurance of permanent income shocks (not uniform across demographic groups); the results point to much of this partial insurance occurring through the adjustment of durable expenditures. Finally we show that taxes and transfers provide an important insurance mechanism for permanent income shocks.

We use the term *partial insurance* to denote smoothing devices other than credit markets for borrowing and saving. There is scattered evidence on the role played by such devices on household consumption. Theoretical and empirical research have analyzed the role of extended family networks [Kotlikoff and Spivak, 1981; Attanasio and Rios-Rull, 2000], added worker effects [Stephens, 2002], the timing of durable purchases [Browning and Crossley, 2001], progressive income taxation [Mankiw and Kimball, 1992, Auerbach and Feenberg, 2001, and Kniesner and Ziliak, 2002], personal bankruptcy laws [Fay, Hurst and White, 2002], insurance within the firm [Guiso, Pistaferri and Schivardi, 2003], financial markets [Davis and Willen, 2000], mortgage refinancing [Hurst and Stafford, 2003], and the role of government public policy programs, such as unemployment insurance [Engen and Gruber, 2001], Medicaid [Gruber and Yelowitz, 1999], AFDC [Gruber, 2000], and food

stamps [Blundell and Pistaferri, 2003].

While we do not take a precise stand on the mechanisms (other than savings) that are available to smooth idiosyncratic shocks to income, we emphasize that our evidence can be used to uncover whether some of these mechanisms are actually at work, how important they are quantitatively, and how they differ across households and over time. Our approach of examining the relationship between consumption and income inequality follows the suggestion of Deaton [1995] that “although it is possible to examine the mechanisms [providing partial insurance against income shocks], their multiplicity makes it attractive to look directly at the magnitude that is supposed to be smoothed, namely consumption”.

The distinction between permanent and transitory shocks stressed in this paper is an important one, as we might expect to uncover less insurance for more persistent shocks. This point has been emphasized in the early work on the permanent income hypothesis and also in the recent analysis of limited commitment models. The literature on insurance under limited commitment [Kocherlakota, 1996, Kehoe and Levine, 2001, Alvarez and Jermann, 2000] explores the nature of income insurance schemes in economies where agents cannot be prevented from withdrawing participation if the loss from the accumulated income gains they are asked to forgo becomes greater than the gains from continuing participation. Such schemes, if feasible, allow individuals to keep some of the positive shocks to their income and therefore offer only partial income insurance. The proportion of income shocks which is insured will vary with the variance of the underlying shocks. As the variance increases the value of future participation increases, alleviating the participation constraint. Krueger and Perri [2001] investigate insurance of transitory shocks through analytic solution of simple models and simulation of more complex cases and demonstrate the possibility that consumption variance can actually fall with an increase in the variance of income shocks. The results in Alvarez and Jermann demonstrate that if income shocks are persistent enough and agents are infinitely lived, then participation constraints become so severe that no insurance scheme is feasible. This suggests that the degree of insurance should be allowed to differ between transitory and permanent shocks and should also be allowed to change over time and across different groups.

Uncovering the degree of partial insurance is likely to matter for a number of reasons. First, the presence of mechanisms that allow households to smooth idiosyncratic shocks has a bearing

on aggregation results [see Blundell and Stoker, 2002]. Second, it may help to understand the characteristics of the economic environment faced by the agents. This may prove crucial when evaluating the performance of macroeconomic models, especially those that explicitly account for agents' heterogeneity. Moreover, it is important to understand to what extent changes in social insurance systems affect smoothing abilities, and the consequences of this for private saving behavior. This is important as far as the efficient design and evaluation of social insurance policy is concerned. It is also particularly relevant in the US and the UK, where quantitatively large changes in the structure of relative prices (most notably, wages) have occurred over the last three decades. Much research exists about the rise in wage inequality, and we shall have very little to add about this. Less evidence exists on consumption inequality [Cutler and Katz, 1992; Dynarski and Gruber, 1997]. We show that consumption inequality follows closely the trends in permanent earnings inequality documented, among others, by Moffitt and Gottschalk [1994].

A study of this kind requires in principle good quality longitudinal data on household consumption and income. It is well known that the PSID contains longitudinal income data but the information on consumption is scanty (limited to food and few more items). Our strategy is to impute consumption to all PSID households combining PSID data with consumption data from repeated CEX cross-sections. Previous studies [Skinner, 1987] impute non-durable consumption data in the PSID using CEX regressions of non durable consumption on consumption items (food, housing, utilities) and demographics available in both the PSID and the CEX. Although related, our approach starts from a standard demand function for food at home (a consumption item available in both surveys); we make this depend on prices, total non durable expenditure, and a host of demographic and socio-economic characteristics of the household. Under monotonicity of food demands these functions can be inverted to obtain a measure of non durable consumption in the PSID. We review the conditions that make this procedure reliable and show that it is able to reproduce remarkably well the trends in the consumption distribution.

The paper continues with an illustration of the model we estimate and of the identification strategy we use (Section 2). In Section 3 we discuss data issues and the imputation procedure. Section 4 contains a discussion of the results and a critical analysis of our finding. Section 5 concludes. The Appendixes discuss technical details about estimation, the approximation of the

Euler equation used in the empirical section and the imputation procedure.

2 Income and Consumption dynamics

2.1 The income process

The unit of analysis is a household, comprising a couple and, possibly, their children. Our sample selection focuses on income risk and we do not model divorce, widowhood, and other household breaking-up factors. We recognize that these may be important omissions that limit the interpretation of our study. However, by focusing on stable households and the interaction of consumption and income we are able to develop a complete identification strategy.³ We also confine our analysis to links between labor income and consumption that become less important after retirement. Consequently we only select households during the working life of the husband.

We assume that the main source of uncertainty faced by the consumer is income (defined as the sum of labor income and transfers, such as welfare payments). We also assume that labor is supplied inelastically and make the assumption of preference separability between consumption and leisure. This means all insurance provided through, say, an added worker effect, will pass through disposable income. Similarly, it is possible that the wage component of family income may have already been smoothed out relative to productivity by implicit agreements within the firm. If this insurance is present, it will be reflected in the variability of income. The income process we consider is:

$$y_{i,a,t} = Z'_{i,a,t}\varphi_t + P_{i,a,t} + v_{i,a,t} \quad (1)$$

where a and t index age and time, respectively, $y = \log Y$ is the log of real income, and Z is a set of observable income characteristics. Note that we allow the effect of demographic characteristics to shift with calendar time. Equation (1) decomposes unexplained income into a permanent component $P_{i,a,t}$ and a transitory or mean-reverting component, $v_{i,a,t}$. By writing $y_{i,a,t}$ rather than $y_{i,t}$ we emphasize the importance of cohort effects in the evolution of earnings over the life-cycle and, more importantly, across generations entering the labor market in different time periods (and thus

³Whether stable families have access to more or less insurance than non-stable families is an issue that cannot be settled in principle. On the one hand, stable families have often more incomes and assets and therefore are less likely to be eligible for social insurance, which is typically mean-tested. On the other hand, they can plausibly be more successful in securing access to credit, family networks and other informal insurance devices, over and above self-insurance through saving.

facing different economic environments and opportunities). In keeping with this remark, we also study consumption decisions of different cohorts.

For consistency with previous empirical studies [MaCurdy, 1982; Abowd and Card, 1989; Moffitt and Gottschalk, 1994; Meghir and Pistaferri, 2003], we assume that the permanent component $P_{i,a,t}$ follows a martingale process of the form:

$$P_{i,a,t} = P_{i,a-1,t-1} + \zeta_{i,a,t} \quad (2)$$

where $\zeta_{i,a,t}$ is serially uncorrelated, and the transitory component $v_{i,a,t}$ follows an MA(q) process, where the order q is to be established empirically:

$$v_{i,a,t} = \sum_{j=0}^q \theta_j \varepsilon_{i,a-j,t-j}$$

with $\theta_0 \equiv 1$. It follows that income growth is:

$$\Delta y_{i,a,t} = \Delta Z'_{it} \varphi_t + \zeta_{i,a,t} + \Delta v_{i,a,t} \quad (3)$$

The covariance restrictions implied by (3) are explored in the next Section.

2.2 Self Insurance and Consumption Growth

Consider the optimization problem faced by household i . Suppose the objective is to:

$$\max E_{a,t} \sum_{j=0}^{T-a} \frac{1}{1 + D_{i,a+j,t+j}} u(C_{i,a+j,t+j}) \quad (4)$$

subject to the intertemporal budget constraints and the initial and terminal conditions on financial assets:

$$A_{i,a+j+1,t+j+1} = (1 + r_{t+j}) (A_{i,a+j,t+j} + Y_{i,a+j,t+j} - C_{i,a+j,t+j}) \quad (5)$$

$$A_{i,a,t} \text{ given} \quad (6)$$

$$A_{i,T,t+T-a} = 0 \quad (7)$$

where $D_{i,a+j,t+j}$ represents taste shifts, discount rate heterogeneity, etc. We set the end of the life-cycle at age T (and retirement at age L), and assume that there is no interest rate uncertainty or uncertainty about the date of death. If preferences are of the CRRA form ($u(C) = \frac{C^{1-\gamma}-1}{1-\gamma}$) and

credit markets are perfect, then one obtains the approximate Euler equation (see Appendix A.1 for more details on the approximation):

$$\Delta c_{i,a,t} \cong \Gamma_{b,t} + \xi_{i,a,t} + \Delta Z'_{i,a,t} \vartheta_t + \pi_{i,a,t} \zeta_{i,a,t} + \gamma_{a,L} \pi_{i,a,t} \varepsilon_{i,a,t} \quad (8)$$

where $c_{i,a,t} = \log C_{i,a,t}$ is the log of real consumption, $\Gamma_{b,t}$ a parameter that varies over time and by cohort (indexed by b), $\xi_{i,a,t}$ a random term, $\Delta Z'_{i,a,t}$ are changes in observable characteristics affecting tastes and impatience, $\gamma_{a,L}$ is a weight that is an increasing function of age, and $\pi_{i,a,t}$ the share of future labor income in the present value of lifetime wealth. In the empirical analysis we assume that $\gamma_{a,L}$ is a known constant rather than a parameter to estimate. The term $\Gamma_{b,t}$ is the slope of the consumption path for different year of birth cohorts, while $\xi_{i,a,t}$ can be interpreted as the individual deviation from the cohort-specific consumption gradient.⁴

For individuals a long time from the end of their life with the value of current financial assets small relative to remaining future labor income, $\pi_{i,a,t} \simeq 1$, and permanent shocks pass through more or less completely into consumption whereas transitory shocks are (almost) completely insured against through saving. Precautionary saving can provide effective insurance against permanent shocks only if the stock of assets built up is large relative to future labor income, which is to say $\pi_{i,a,t}$ is appreciably smaller than unity, in which case there will be some smoothing of permanent shocks through self insurance (see also Carroll, 2001, for numerical simulations). From here onwards, y and c should be interpreted as the income and consumption components after removing demographic characteristics and aggregate effects. The terms $Z_{i,a,t}$ and $\Gamma_{b,t}$ will thus be omitted from now on. The remainder of this section considers the case $\pi_{i,a,t} \simeq 1$ in which no part of permanent shocks is insured through precautionary saving. We defer a discussion of the consequences of removing this assumption to Section 4.4. To pre-empt, when we allow for partial insurance we are unable to separately identify how precautionary saving (through $\pi_{i,a,t}$) and partial insurance over and above saving smooth the impact of shocks on consumption. However, this will be practically of little importance. We will be identifying a parameter that combines self-insurance, partial insurance, and perhaps even the crowding out effect of public insurance on private insurance. In other words, we will still be able to pin down the degree of transmission of income shocks into consumption.

⁴Innovations to the conditional variance of consumption growth (precautionary savings) are captured by $\Gamma_{b,t}$.

We assume that $\zeta_{i,a,t}$, $\xi_{i,a,t}$ and $v_{i,a,t}$ are mutually uncorrelated processes. The only source of serial correlation, if present, arises from the transitory component $v_{i,a,t}$. Equation (3) can be used to derive the following covariance restrictions in panel data

$$\text{cov}(\Delta y_{a,t}, \Delta y_{a+s,t+s}) = \begin{cases} \text{var}(\zeta_{a,t}) + \text{var}(\Delta v_{a,t}) & \text{for } s = 0 \\ \text{cov}(\Delta v_{a,t}, \Delta v_{a+s,t+s}) & \text{for } s \neq 0 \end{cases} \quad (9)$$

where $\text{var}(\cdot)$ and $\text{cov}(\cdot, \cdot)$ denote cross-sectional variances and covariances, respectively (the index i is consequently omitted). These moments can be computed for the whole sample or for individuals belonging to a homogeneous group (i.e., born in the same year, with the same level of schooling, etc.). The covariance term $\text{cov}(\Delta v_{a,t}, \Delta v_{a+s,t+s})$ depends on the serial correlation properties of v . If v is an $MA(q)$ serially correlated process, then $\text{cov}(\Delta v_{a,t}, \Delta v_{a+s,t+s})$ is zero whenever $|s| > q + 1$. Note also that if v is serially uncorrelated ($v_{i,a,t} = \varepsilon_{i,a,t}$), then $\text{var}(\Delta v_{a,t}) = \text{var}(\varepsilon_{a,t}) + \text{var}(\varepsilon_{a-1,t-1})$. The moment restriction for $s = 0$ in (9) can also be recovered from repeated cross-section data, an observation underlying the analysis in Blundell and Preston [1998].

The panel data restrictions on consumption growth from (8) are as follows:

$$\text{cov}(\Delta c_{a,t}, \Delta c_{a+s,t+s}) = \text{var}(\xi_{a,t}) + \text{var}(\zeta_{a,t}) + \gamma_{a,L}^2 \text{var}(\varepsilon_{a,t}) \quad (10)$$

for $s = 0$ and zero otherwise (due to the consumption martingale assumption). The equivalent repeated cross-section moments are used in Blundell and Preston in conjunction with those relating to income to separate the growth in the variance of transitory shocks to log income from the variance of permanent shocks.

Finally, the covariance between income growth and consumption growth at various lags is:

$$\text{cov}(\Delta c_{a,t}, \Delta y_{a+s,t+s}) = \begin{cases} \text{var}(\zeta_{a,t}) + \gamma_{a,L} \text{var}(\varepsilon_{a,t}) \\ \gamma_{a,L} \text{cov}(\varepsilon_{a,t}, \Delta v_{a+s,t+s}) \end{cases} \quad (11)$$

for $s = 0$, and $s \neq 0$ respectively. Blundell and Preston note that the corresponding repeated cross section moments give overidentification to their empirical analysis.

In the context of identifying sources of variation in household income and consumption, it is worth stressing that the availability of panel data presents several advantages over a repeated cross-sections analysis. In the latter case identification requires assuming that shocks are cross-sectionally orthogonal to past consumption and income [see also Deaton and Paxson, 1994]. This allows one to replace the unobservable $\text{var}(\Delta c_{a,t})$ with $\Delta \text{var}(c_{a,t})$, $\text{var}(\Delta y_{a,t})$ with $\Delta \text{var}(y_{a,t})$, and

$\text{cov}(\Delta y_{a,t}, \Delta c_{a,t})$ with $\Delta \text{cov}(y_{a,t}, c_{a,t})$. This assumption will be violated if, say, knowledge of one's position in the income (or consumption) distribution conveys information about the distribution of future shocks to income. In panel data, identification of the various parameters does not require making such assumption and can allow for serial correlation in transitory shocks as well as measurement error in consumption and income data. Finally, panel data provide more overidentifying restrictions vis-à-vis repeated cross-section and thus can afford more flexibility in terms of model testing.

As an example of the advantages offered by panel data, note for instance that identification of the variances of shocks to income requires only panel data on income, not consumption. In the simple case of serially uncorrelated transitory shock, for example:⁵

$$\text{var}(\zeta_{a,t}) = \text{cov}(\Delta y_{a,t}, \Delta y_{a-1,t-1} + \Delta y_{a,t} + \Delta y_{a+1,t+1}) \quad (12)$$

$$\text{var}(\varepsilon_{a,t}) = -\text{cov}(\Delta y_{a,t}, \Delta y_{a+1,t+1}) \quad (13)$$

We refer the reader to Section 4.2 for a discussion of the estimates of these variances.

2.3 Partial insurance

We now consider the possibility of partial insurance and suppose there are mechanisms (that we do not model explicitly here but were discussed above) that allow insurance of a fraction $(1 - \phi_{b,t})$ and $(1 - \psi_{b,t})$ of permanent and transitory shocks, respectively. We might expect $\phi_{b,t}$ to be close to unity and $\psi_{b,t}$ close to zero. As noted above, precautionary saving might allow partial insurance of permanent shocks if assets were large enough relative to future labor income (i.e. $\pi_{i,a,t} < 1$), but interpersonal insurance mechanisms might also underlie this.⁶ For simplicity of notation, we confine ourselves to the case where the transitory shock is serially uncorrelated, consumption is measured without error and $\text{var}(\xi_{a,t}) = 0$. These extensions are taken up in the empirical analysis.

In the partial insurance case residual income and consumption growth can be written, respec-

⁵See Meghir and Pistaferri [2003] for a generalization to serially correlated transitory shocks and measurement error in income.

⁶If there are no interpersonal mechanisms or transfers of any sort, then $\phi = \psi = \pi = 1$.

tively, as:

$$\Delta y_{i,a,t} = \zeta_{i,a,t} + \Delta \varepsilon_{i,a,t} \quad (14)$$

$$\Delta c_{i,a,t} \cong \phi_{b,t} \zeta_{i,a,t} + \psi_{b,t} \gamma_{a,L} \varepsilon_{i,a,t} \quad (15)$$

The economic interpretation of the partial insurance parameter is such that it nests the two polar cases of full insurance of income shocks ($\phi_{b,t} = \psi_{b,t} = 0$), as contemplated by the complete markets hypothesis, and no insurance ($\phi_{b,t} = \psi_{b,t} = 1$), as predicted by the PIH with just self-insurance through savings. A value $0 < \phi_{b,t} < 1$ ($0 < \psi_{b,t} < 1$) is consistent with partial insurance with respect to permanent (transitory) shocks. The lower the coefficient, the higher the degree of insurance.

The relevant panel data moments are:

$$\begin{aligned} \text{var}(\Delta y_{a,t}) &= \text{var}(\zeta_{a,t}) + \text{var}(\varepsilon_{a,t}) + \text{var}(\varepsilon_{a-1,t-1}) \\ \text{cov}(\Delta y_{a,t}, \Delta y_{a-1,t-1}) &= -\text{var}(\varepsilon_{a-1,t-1}) \\ \text{cov}(\Delta y_{a+1,t+1}, \Delta y_{a,t}) &= -\text{var}(\varepsilon_{a,t}) \\ \text{var}(\Delta c_{a,t}) &= \phi_{b,t}^2 \text{var}(\zeta_{a,t}) + \psi_{b,t}^2 \gamma_{a,L}^2 \text{var}(\varepsilon_{a,t}) \\ \text{cov}(\Delta c_{a,t}, \Delta y_{a,t}) &= \phi_{b,t} \text{var}(\zeta_{a,t}) + \psi_{b,t} \gamma_{a,L} \text{var}(\varepsilon_{a,t}) \\ \text{cov}(\Delta c_{a,t}, \Delta y_{a+1,t+1}) &= -\psi_{b,t} \gamma_{a,L} \text{var}(\varepsilon_{a,t}) \end{aligned}$$

Since $\text{var}(\zeta)$ and $\text{var}(\varepsilon)$ can still be identified from panel data on income (the first three moments above), there are only two parameters left to identify: $\phi_{b,t}$ and $\psi_{b,t}$. Take first the simple case $L - a \rightarrow \infty$ and $r \rightarrow 0$ ($\gamma_{a,L} \rightarrow 0$). Then either:

$$\phi_{b,t} = \frac{\text{var}(\Delta c_{a,t})}{\text{cov}(\Delta c_{a,t}, \Delta y_{a,t})} \quad (16)$$

or:

$$\phi_{b,t} = \frac{\text{cov}(\Delta c_{a,t}, \Delta y_{a,t})}{\text{cov}(\Delta y_{a,t}, \Delta y_{a-1,t-1} + \Delta y_{a,t} + \Delta y_{a+1,t+1})}. \quad (17)$$

identify the extent of insurance against permanent shocks ($\psi_{b,t}$ is obviously not identified). Under these assumptions $\phi_{b,t}$ would be overidentified and the overidentifying restrictions could be tested using standard methods. Relaxing, as we do in our empirical work, the assumption that $\text{var}(\xi_{a,t}) = 0$ would undermine (16) leaving identification reliant on (17).

As a matter of interpretation, note that the numerator of (16) captures the variance of shifts in consumption. In a model with no transitory shock effects on consumption, the volatility of consumption growth depends only on the arrival of permanent shocks to income and the availability of insurance mechanisms above self-insurance. The denominator of (16) measures the association between consumption growth and income growth. In a model with no transitory shocks, consumption growth tracks income growth only through its long run component. In the absence of partial insurance mechanisms, the numerator and the denominator will be measuring exactly the same (permanent) variability in income. Recall that in the self-insurance, infinite-horizon case any permanent shift in the variance of the distribution of income is paralleled by an equivalent permanent shift in the variance of the distribution of consumption. With partial insurance, however, the latter is attenuated by the fact that permanent income shocks translate less than one-for-one into consumption; the amount of attenuation (given by the ratio in (16)) is exactly measured by the parameter $\phi_{b,t}$.

In the more general case of finite horizon and $r \neq 0$, more complicated expressions are available to identify the coefficients of interest. For instance,

$$\begin{aligned}\psi_{b,t} &= \gamma_{a,L}^{-1} \frac{\text{cov}(\Delta c_{a,t}, \Delta y_{a+1,t+1})}{\text{cov}(\Delta y_{a+1,t+1}, \Delta y_{a,t})} \\ \phi_{b,t} &= \frac{\text{cov}(\Delta c_{a,t}, \Delta y_{a,t} + \Delta y_{a+1,t+1})}{\text{cov}(\Delta y_{a,t}, \Delta y_{a-1,t-1} + \Delta y_{a,t} + \Delta y_{a+1,t+1})}\end{aligned}$$

and $\psi_{b,t}$ and $\phi_{b,t}$ are generally overidentified even in this simple model.

Finally, it is very likely that *measurement error* will contaminate the observed income and consumption data. Assume that both consumption and income are measured with multiplicative (independent and quasi-classical) error,⁷ e.g.,

$$y_{i,a,t}^* = y_{i,a,t} + u_{i,a,t}^y \quad (18)$$

and

$$c_{i,a,t}^* = c_{i,a,t} + u_{i,a,t}^c \quad (19)$$

⁷Quasi-classical in the sense that the variance of the measurement error is not (necessarily) constant over time.

where x^* denote a measured variable, x its true, unobservable value, and u the measurement error. In Appendix A.2 we show that the partial insurance parameter $\phi_{b,t}$ remains identified under measurement error.

2.4 Information

In the analysis presented so far we have assumed that in the innovation process for income (14) the random variables $\zeta_{i,a,t}$ and $\varepsilon_{i,a,t}$ represent the arrival of new information to the agent i of age a in period t . If part of this random term was known to the agent then the consumption model would argue that it should already be incorporated into consumption plans and would not directly effect consumption growth (15). Suppose a proportion κ_ζ of the permanent shock was **not** known in advance to the consumer. Then the consumption growth relationship (15) would become

$$\Delta c_{i,a,t} \cong \tilde{\phi}_{b,t} \kappa_\zeta \zeta_{i,a,t} + \psi_{b,t} \gamma_{a,L} \varepsilon_{i,a,t}. \quad (20)$$

In this case the estimated $\phi_{b,t}$ would overstate the extent of partial insurance by the information factor κ_ζ .

The econometrician will treat $\zeta_{i,a,t}$ as the permanent shock. Whereas the individual may have already adapted to this change. Consequently, although transmission of income inequality to consumption inequality is correctly identified, the estimated $\phi_{b,t}$ has to be interpreted as reflecting a combination of insurance and information. This is discussed further in section 4.4 where we interpret our empirical results.

3 The data

Our empirical analysis is conducted on two microeconomic data sources: the 1978-1992 PSID and the 1980-1992 CEX. We describe their main features and our sample selection procedures in turn.

3.1 The PSID

Since the PSID has been widely used for microeconomic research, we shall only sketch the description of its structure in this section.⁸

⁸See Hill [1992] for more details about the PSID.

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 Survey of Economic Opportunities, or SEO sample). Thereafter, both the original families and their split-offs (children of the original family forming a family of their own) have been followed.

The PSID includes a variety of socio-economic characteristics of the household, including age, education, labor supply, and income of household members. Questions referring to income are retrospective; thus, those asked in 1993, say, refer to the 1992 calendar year. In contrast, many researchers have argued that the timing of the survey questions on food expenditure is much less clear [Hall and Mishkin, 1982; Altonji and Siow, 1987]. Typically, the PSID asks how much is spent on food in an average week. Since interviews are usually conducted around March, it has been argued that people report their food expenditure for an average week around that period, rather than for the previous calendar year as is the case for family income. We assume that food expenditure reported in survey year t refers to the previous calendar year, but check the effect of alternative assumptions.

Households in the PSID report their taxable family income (which includes transfers and financial income). The measure of income used in the baseline analysis below excludes income from financial assets, subtracts taxes and deflates the corresponding value by the CPI. We obtain an after-tax measure of income subtracting federal taxes paid. Before 1991, these are computed by PSID researchers and added into the data set using information on filing status, adjusted gross income, whether the respondent itemizes or takes the standard deduction, and other household characteristics that make them qualify for extra deductions, exemptions, and tax credits. Federal taxes are not computed in 1992 and 1993. We impute taxes for the last two years using regression analysis for the years where taxes are available (results not reported but available on request).

Education level is computed using the PSID variable “grades of school finished”. Individuals who changed their education level during the sample period are allocated to the highest grade achieved. We consider two education groups: with and without college education (corresponding to 13 grades or more and 12 grades or less, respectively).

Since CEX data are available on a consistent basis since 1980, we construct an unbalanced

PSID panel using data from 1978 to 1992 (the first two years are retained for initial conditions purposes). Due to attrition, changes in family composition, and various other reasons, household heads in the 1978-1992 PSID may be present from a minimum of one year to a maximum of fifteen years. We thus create unbalanced panel data sets of various length. The longest panel includes individuals present from 1978 to 1992; the shortest, individuals present for two consecutive years only (1978-79, 1979-80, up to 1991-92).

The objective of our sample selection is to focus on a sample of continuously married couples headed by a male (with or without children). The step-by-step selection of our PSID sample is illustrated in Table I. We eliminate households facing some dramatic family composition change over the sample period. In particular, we keep only those with no change, and those experiencing changes in members other than the head or the wife (children leaving parental home, say). We next eliminate households headed by a female. We also eliminate households with missing report on education and region,⁹ and those with topcoded income. We keep continuously married couples and drop some income outliers.¹⁰ We then drop those born before 1920 or after 1959.

As noted above, the initial 1967 PSID contains two groups of households. The first is representative of the US population (61 percent of the original sample); the second is a supplementary low income subsample (also known as SEO subsample, representing 39 percent of the original 1967 sample). To account for the changing demographic structure of the US population, starting in 1990 a representative national sample of 2,000 Latino households has been added to the PSID database. For the most part we exclude both Latino and SEO households and their split-offs. However, we do consider the robustness of our results in the low income SEO subsample.

Finally, we drop those aged less than 30 or more than 65. This is to avoid problems related to changes in family composition and education, in the first case, and retirement, in the second. The final sample used in the minimum distance exercise below is composed of 17,788 observations and 1,788 households.

We use information on age and the survey year to allocate individuals in our sample to four cohorts defined on the basis of the year of birth of the household head: born in the 1920s, 1930s,

⁹When possible, we impute values for education and region of residence using adjacent records on these variables.

¹⁰An income outlier is defined as a household with an income growth above 500 percent, below -80 percent, or with a level of income below \$100 a year or below the amount spent on food.

1940s, and 1950s. Years where cell size is less than 100 are discarded.¹¹

3.2 The CEX

The Consumer Expenditure Survey provides a continuous and comprehensive flow of data on the buying habits of American consumers. The data are collected by the Bureau of Labor Statistics and used primarily for revising the CPI. Consumer units are defined as members of a household related by blood, marriage, adoption, or other legal arrangement, single person living alone or sharing a household with others, or two or more persons living together who are financially dependent. The definition of the head of the household in the CEX is the person or one of the persons who owns or rents the unit; this definition is slightly different from the one adopted in the PSID, where the head is always the husband in a couple. We make the two definitions compatible.

The CEX is based on two components, the Diary, or record keeping survey and the Interview survey. The Diary sample interviews households for two consecutive weeks, and it is designed to obtain detailed expenditures data on small and frequently purchased items, such as food, personal care, and household supplies. The Interview sample follows survey households for a maximum of 5 quarters, although only inventory and basic sample data are collected in the first quarter. The data base covers about 95% of all expenditure, with the exclusion of expenditures for housekeeping supplies, personal care products, and non-prescription drugs. Following most previous research, our analysis below uses only the Interview sample.¹²

The CEX collects information on a variety of socio-demographic variables, including characteristics of members, characteristics of housing unit, geographic information, inventory of household appliances, work experience and earnings of members, unearned income, taxes, and other receipts of consumer unit, credit balances, assets and liabilities, occupational expenses and cash contributions of consumer unit. Expenditure is reported in each quarter and refers to the previous quarter; income is reported in the second and fifth interview (with some exceptions), and refers to the previous twelve months. For consistency with the timing of consumption, fifth-quarter income data

¹¹Median (average) cell sizes are 249 (219), 245 (246), 413 (407), and 398 (363), respectively for those born in the 1920s, 1930s, 1940s, and 1950s.

¹²There is some evidence that trends in consumption inequality measured in the two CEX surveys have diverged in the 1990s [Battistin, 2003]. While research on the reasons for this divergence is clearly warranted, our analysis, which uses data up to 1992, will only be marginally affected.

are used.

We select a CEX sample that can be made comparable, to the extent that this is possible, to the PSID sample. Our initial 1980-1998 CEX sample includes 1,249,329 monthly observations, corresponding to 141,289 households. We drop those with missing record on food and/or zero total nondurable expenditure, and those who completed less than 12 month interviews. This is to obtain a sample where a measure of annual consumption can be obtained. A problem is that many households report their consumption for overlapping years, i.e. there are people interviewed partly in year t and partly in year $t + 1$. Pragmatically, we assume that if the household is interviewed for at least 6 months at $t + 1$, then the reference year is $t + 1$, and it is t otherwise. Prices are adjusted accordingly. We then sum food at home, food away from home and other nondurable expenditure over the 12 interview months. This gives annual expenditures. For consistency with the timing of the PSID data, we drop households interviewed after 1992. We also drop those with zero before-tax income, those with missing region or education records, single households and those with changes in family composition. Finally, we eliminate households where the head is born before 1920 or after 1959, those aged less than 30 or more than 65, those with outlying income (defined as a level of income below the amount spent on food), and households entering the CEX after 1992. Our final sample contains 15,380 households. Table II details the sample selection process in the CEX.

The definition of total non durable consumption is similar to Attanasio and Weber [1995]. It includes food (at home and away from home), alcoholic beverages and tobacco, services, heating fuel, transports (including gasoline), personal care, clothing and footwear, and rents. It excludes expenditure on various durables, housing (furniture, appliances, etc.), health, and education. In our empirical results we assess the sensitivity of our results to the inclusion of durables and other non-durable items.

3.3 Comparing the two data sets

How similar are the two data sets in terms of average demographic and socio-economic characteristics? Mean comparisons are reported in Table III for selected years: 1980, 1983, 1986, 1989, and 1992.

PSID respondents are slightly younger than their CEX counterparts; there is, however, little

difference in terms of family size and composition. The percentage of whites is slightly higher in the PSID. The distribution of the sample by schooling levels is quite similar, while the PSID tends to under-represent the proportion of people living in the West. Due to slight differences in the definition of family income, PSID figures are higher than those in the CEX. It is possible that the definition of family income in the PSID is more comprehensive than that in the CEX, so resulting in the underestimation of income in the CEX that appears in the Table.

Trends in food expenditure are initially quite similar, but tend to diverge afterwards. One explanation for this is that the wording of the food expenditure question in the CEX changed several times over this period, resulting in PSID figures being higher than the corresponding CEX figures. In particular, the CEX survey question asks about average monthly expenditure over the last quarter in 1980-81, expenditure in an average week of the last quarter in 1982-87, and reverts to the monthly expenditure question starting in 1988. Some researchers [Garner et al., 1998] have noted that the change in the survey question induces a dramatic underestimation of food expenditure in the CEX in the 1982-87 period. As we shall see, accounting for differences in the way the food question is asked in the two data sets is crucial when replicating trends in average consumption. Food away from home is, in contrast, under-reported in the PSID *vis-à-vis* the CEX.

The last two rows of Table III compare the labor market activity of the household head and of the spouse. Both male and female participation rates in the PSID are comparable to those in the CEX.

3.4 The imputation procedure

In deriving the theoretical restrictions above, we have assumed that a researcher has access to panel data on household income and total non-durable consumption. However, this is a very strong data requirement. In the US, panel data typically lack household data on total non-durable consumption; and those surveys, such as the CEX, that contains good quality data on consumption, lack a panel feature. We may however combine the two data sets to impute non durable consumption to PSID households. This of course requires making some assumptions detailed below.

The PSID collects data on few consumption items, mainly food at home and food away from home. Moreover, food data are not available in 1987 and 1988. Our strategy is to write a demand

equation for food as a function of prices, demographics, labor supply variables, food away from home and total non-durable expenditure. Within each cohort, variability of food consumption is then explained by variability in those components plus unobserved heterogeneity, measurement error, etc. We can invert this relationship to obtain a measure of the variability in total non-durable consumption, one of the main objects of interest from the previous section. This inversion operation requires consistent estimation of the parameters of the demand function for food and monotonicity of the underlying demand function.

Our paper is not the first to combine CEX data with PSID data to impute a measure of non durable consumption in the PSID [examples include Skinner, 1987; Bernheim, Skinner and Weinberg, 2001; Dynan, 2000]. We adopt a demand function approach¹³ which we show is quite successful in replicating the trends in the variance of consumption that are the main object of our empirical analysis. To make matching of the two data sets feasible, the demand function for food must be invertible with respect to total expenditure. This requires the demand function to be monotonic in total expenditure.¹⁴ After some experimentation, we selected a loglinear functional form. The main advantage of the loglinear demand function is that it provides “ready-to-use” predictions for total nondurable expenditure, avoiding, for instance, the problem of negative predicted values faced when using the linear expenditure demand function. The loglinear demand function has also a series of shortcomings, however. In particular, it cannot capture zero expenditures, it does not satisfy adding up if applied to all goods in a demand system, and it does not capture apparent non-linearities in Engel curve relationships. Nevertheless, these shortcomings do not appear particularly relevant here. There are no zeros in food spending, the specification below is applied to just one good, and the Engel curve for food is not far from being log linear.

Formally, we write the following demand equation for food at home in the CEX:

$$f_{i,a,t} = W'_{i,a,t}\mu + \beta(Z_{i,a,t})c_{i,a,t} + e_{i,a,t} \quad (21)$$

where f is the log of food expenditure (which is available in both surveys), W contains prices, food away from home and a set of demographics and labor supply variables (also available in both data

¹³Skinner [1987] regresses non durable consumption on all the consumption items available in both surveys (food at home, food away from home, utilities, rents, etc.).

¹⁴While theoretically more appealing, flexible functional forms, such as variants of the AIDS model of Deaton and Muellbauer [1980], or demand functions that are non-linear in total nondurable expenditure or contain interactions [Banks, Blundell, and Lewbel, 1999], may violate this requirement.

sets), c is the log of total non-durable expenditure (available only in the CEX), and e captures unobserved heterogeneity in the demand for food and measurement error in food expenditure. We allow for the elasticity $\beta(\cdot)$ to vary with time and with observable household characteristics.

We pool all the CEX data from 1980 to 1992. Our specification includes the log of the price of food at home,¹⁵ the interaction of this with region of residence, the log of total nondurable expenditure and its interaction with year dummies, education dummies and indicators for number of children (no children, one child, two or three children, four children or more). We further include indicators for male and female labor market participation and their interactions with the number of children, food away from home,¹⁶ and a vector of demographics (an age spline, dummies for education and region of residence, year of birth dummies, indicators for number of children as above, family size and its interaction with region dummies, and a dummy for whites). We use the price of food at national level. Measurement error in nondurable expenditure biases the expenditure elasticity towards zero; we thus instrument this variable with the log of before-tax income (and interactions with demographics and year dummies).¹⁷ Standard errors are corrected for time clustering. The estimation results are reported in Table IV.

We estimate an average expenditure elasticity of 0.66; this declines with education and increases with the number of children. It also varies significantly with calendar time. The main effect is tightly estimated, and so are most of the interactions.¹⁸ The estimate of the price elasticity is -0.9 , and is marginally significant. Interaction of the price variable with region dummies are insignificant. Other demographics have the expected sign. Labor market participation reduces expenditure on food at home (more strongly so for females); the effect of male participation is insignificant but becomes stronger in the presence of more children.

Armed with the estimated demand parameters, we invert the demand equation for food and obtain a measure of total nondurable expenditure in the PSID matching on observable characteristics that are common to the two data sets. As explained in Appendix A.4, a good inversion procedure should have two defining properties: (a) average (imputed) consumption in the PSID should coin-

¹⁵We omit prices of other commodities due to multicollinearity problems.

¹⁶The inclusion of labor market participation variables and food away from home is justified by a conditional demand approach.

¹⁷This is important as far as replicating trends in consumption variance is concerned (see Appendix A.4).

¹⁸The average expenditure elasticity is lower in the OLS case (0.61).

cide with average consumption in the CEX, and (b) the variance of (imputed) consumption in the PSID should exceed the variance of consumption in the CEX by an additive factor (the variance of the error term of the demand equation scaled by the square of the expenditure elasticity). If this factor is constant over time the *trends* in the two variances should be identical.

As for the latter, trends in the variance of consumption are indeed remarkably similar in the two data sets, as Figure 1 shows. Between 1980 and 1986 the variance of PSID imputed consumption and the variance of CEX raw consumption are growing at a similar rate. Afterwards, they are flat. The levels differ by a common factor as expected if the imputation procedure is reliable, see Appendix A.4.

As for average consumption, our procedure appears to do less well. As shown in Panel A of Figure 2, imputed average log consumption in the PSID tend to systematically exceed CEX average log consumption. One reason for this divergence is, as noted above, that the food question is asked differently in the two data sets, which may result in average food consumption (the main input variable of our imputation procedure) being different as well. Panel B of Figure 2 shows that this is indeed the case: in the 1982-1987 period CEX average food consumption (in nominal terms) is dramatically different from PSID average food consumption. As shown in Appendix A.4, average imputed consumption in the PSID may differ from average CEX consumption simply because average food expenditure differs in the two data sets, not because the imputation procedure is unreliable. Indeed, the overestimation of non-durable consumption in the PSID (Panel A) is the mirror image of the underestimation of food expenditure in the CEX (Panel B). This is shown in Panel C, where we plot the difference between PSID and CEX non-durable consumption and between PSID and CEX food expenditure (scaled by the consumption elasticity). The two basically coincide. Once we correct for this discrepancy, average consumption in the two data sets line up remarkably well (see Panel D). Note that this is not “mechanically” true; using a biased estimate of the consumption elasticity, for example, will not eliminate the discrepancy.

The evidence discussed in this section thus provides confidence in our use of imputed data to estimate the parameters of interest discussed in Section 2. We now turn to the results of our empirical analysis.

4 The results

We organize the empirical analysis in three parts: evidence on consumption inequality (section 4.1), unrestricted consumption-income autocovariance estimation from the PSID (section 4.2), and minimum distance estimation using longitudinal data on household income and predicted consumption (section 4.3). We then discuss our findings and a variety of experiments (Section 4.4).

4.1 Consumption Inequality and Income Uncertainty - Evidence from the CEX

Table V reports estimates of the variance of log consumption, the variance of log income, and their covariance for all years and for four cohorts (born in the 1920s, 1930s, 1940s and 1950s).¹⁹ We use data from 1980 to 1992 from the CEX. Figure 3 graphs income and consumption variances and their covariance over the life cycle (we smooth trends using a rolling MA(3)). As said in the previous Section, these trends are replicated by our imputed measure of consumption in the PSID. For the two middle cohorts the variance of consumption increases throughout the sample period (both variance grow very moderately in the second half of the 1980s). For the youngest cohort the variance of income is flat while the variance of consumption actually declines. Finally, for the oldest cohort both variances increase in the early 1980s and decline afterward. Trends in covariances resemble those for the consumption variance (the level is higher). Two things are worth noting: income inequality is higher and it grows more rapidly than consumption inequality. The difference in the slopes of the two profiles is prima facie evidence against a simple model where variances of shocks are stable and no insurance, apart from savings, is available. This simple model would predict that income and consumption inequality grow at the same rate (the variance of permanent shocks). Such a simple model is neither borne out visually by the data nor, as we shall see, confirmed by more formal analysis below.

4.2 Autocovariance Estimates of Consumption and Income: Longitudinal Evidence from the Matched PSID

The PSID data set contains longitudinal records on income and imputed consumption. We remove the effect of deterministic effects on log income and (imputed) consumption by separate regressions

¹⁹These are the variances of deviations of consumption and income per household member from the cohort-specific cross-sectional means. The covariance is the covariance of such residuals.

of these variables on year and year of birth dummies, and on a set of observable family characteristics (dummies for education, race, family size, number of children, region, employment status, residence in a large city, outside dependent, and presence of income recipients other than husband and wife). We allow for the effect of these characteristics to vary with calendar time. These variables reflect deterministic growth in consumption and income (information). We then work with the residuals of these regressions, $c_{i,a,t}$ and $y_{i,a,t}$.

To pave the way to the formal analysis of Section 4.3, Table VI reports unrestricted minimum distance estimates of several moments of interest for the whole sample: the variance of unexplained income growth, $\text{var}(\Delta y_{a,t})$, the first-order autocovariances ($\text{cov}(\Delta y_{a+1,t+1}, \Delta y_{a,t})$), and the second-order autocovariances ($\text{cov}(\Delta y_{a+2,t+2}, \Delta y_{a,t})$). Estimates are reported for each year. Table VII repeats the exercise for our measure of consumption. Finally, Table VIII reports minimum distance estimates of contemporaneous and lagged consumption-income covariances.

Looking through Table VI, one can notice the strong increase in the variance of income growth, especially in the early 1980s. Also notice the strong blip in the final year (in 1992 the PSID converted the questionnaire to electronic form and imputations of income done by machine). The absolute value of the first-order autocovariance also increases in the early 1980s and then is stable or even declines after 1986. Second- and higher order autocovariances are small and only in few cases statistically significant. At least at face value, this evidence seems to tally quite well with a canonical MA(1) process, as implied by a traditional income process given by the sum of a martingale permanent component and a serially uncorrelated transitory component:²⁰

$$\Delta y_{i,a,t}^* = \zeta_{i,a,t} + \Delta v_{i,a,t} \quad (22)$$

(with the permanent shock ζ being serially uncorrelated and the transitory shock v possibly MA(1), i.e., $v_{i,a,t} = \varepsilon_{i,a,t} + \theta \varepsilon_{i,a-1,t-1}$). Most attention has been paid to estimating the variance of the two components ζ and v [Gottschalk and Moffitt, 1994; Blundell and Preston, 1998; Meghir and Pistaferri, 2003]. In particular, their time trends can be used to understand whether the recent rise in inequality has long- or short-run (instability) characteristics. This can be achieved by imposing

²⁰Since evidence on second-order autocovariances is mixed, in estimation we allow for MA(1) serial correlation in the transitory component.

the theoretical restrictions of equation (22) on the autocovariances of income at various lags as reported in Table VI.

In Figures 4 and 5 we plot equally weighted minimum distance estimates of $\text{var}(\zeta_{a,t})$ and $\text{var}(\varepsilon_{a,t})$, respectively, against time.²¹ Since estimates in the last two years are statistically imprecise, we focus only on the 1980s. Our results show a strong growth in permanent income shocks during the early 1980s. The variance of permanent shocks levels off thereafter. The variance of the transitory shock is basically flat in the period where the variance of permanent shock is increasing, and it increases only when the variance of permanent shock slows down. This evidence is similar to that reported by Moffitt and Gottschalk [1994] using PSID earnings data. It is also worth noting that from trough to peak the variance of the permanent shock doubles, but the variance of the transitory shock only goes up by about 50%.

While income moments are informative about shifts in the income distribution (and on the temporary or persistent nature of such shifts), they cannot be used to make conclusive inference about shifts in the consumption distribution. For this purpose, one needs to complement the analysis of income moments with that of consumption moments and of the joint income-consumption moments. This is done in Tables VII and VIII.

Table VII shows that the variance of imputed consumption growth increases quite strongly in the early 1980s, peaks in 1984 and then it is essentially flat afterwards. Note the high value of the variance which is clearly the result of our imputation procedure. The variance of consumption growth captures in fact the genuine association with shocks to income, but also the contribution of slope heterogeneity and measurement error.²² The absolute value of the first-order autocovariance of consumption growth should be a good estimate of the variance of the imputation error. This is in fact quite high and approximately stable over time. Second-order consumption growth autocovariances are mostly statistically insignificant and economically small.

Table VIII looks at the association, at various lags, of unexplained income and consumption growth. The contemporaneous covariance should be informative about the effect of income shocks

²¹Since we cannot separately identify measurement error and transitory shock, the estimated variance of the transitory shock will also reflect the contribution of income data noise. In the figures we smooth the estimates using simple 3-years moving averages.

²²To a first approximation, the variance of consumption growth that is not contaminated by error can be obtained by subtracting twice the (absolute value of) first order autocovariance $\text{cov}(\Delta c_{t+1}, \Delta c_t)$ from the variance $\text{var}(\Delta c_t)$.

on consumption growth if measurement errors in consumption are orthogonal to measurement errors in income. This covariance increases in the early 1980s and then is flat or even declining afterwards.

The covariance between current consumption growth and future income growth $\text{cov}(\Delta y_{a+1,t+1}, \Delta c_{a,t})$ should reflect the extent of insurance with respect to transitory shocks. Note that in the pure self-insurance case and with infinite horizon, the impact of transitory shocks on consumption growth is the annuity value $\frac{r}{1+r}$. With a small interest rate, this will be indistinguishable from zero, at least statistically. The addition of partial insurance $\psi_{b,t} < 1$ makes this even more likely. In fact, this covariance is hardly statistically significant and economically close to zero. As we shall see, the formal analysis below will confirm this.

The covariance between current consumption growth and past income growth $\text{cov}(\Delta c_{a+1,t+1}, \Delta y_{a,t})$ plays no role in the PIH model with perfect capital markets, but may be important in alternative models where liquidity constraints are present. The estimates of this covariance in Table VIII are close to zero. We should note, however, that for the low income sample examined further in the empirical results below we do find some sensitivity to transitory shocks. To sum up, there is weak evidence that transitory shocks impact consumption growth or that liquidity constraints are empirically important. In the sensitivity results reported below we note that there is more evidence of responsiveness to transitory shocks for the low income poverty sample of the PSID. We now turn to more formal minimum distance estimation, where we impose the theoretical restrictions outlined in Section 2.3 on the unrestricted income and consumption moments of Table VI, VII, and VIII.

4.3 Partial Insurance

Here we focus on the results of a non-stationary model where the parameters vary across cohorts or education groups (depending on the specification adopted) and time: in particular, we assume that they shift at some point in the mid-1980s (consistent with most of the Figures or results discussed above). We estimate the parameters that characterize the income and consumption process by equally weighted minimum distance.²³ Technical details are in Appendix A.3.

There are several parameters to estimate: the variances of the permanent and transitory income

²³In general, the choice is between a non-linear least squares procedure (equally weighted minimum distance, or EWMD) and a non-linear generalized least squares procedure (optimal minimum distance, or OMD). Altonji and Segal [1996] show that EWMD dominates OMD even for moderately large sample sizes. OMD estimates are qualitatively similar and available on request.

shocks (σ_ζ^2 and σ_ε^2 , respectively), the MA coefficient θ of the transitory shock, the variance of individual consumption gradients (σ_ξ^2) and imputation error (σ_u^2), the partial insurance coefficient for the permanent shock (ϕ) and for the transitory shock (ψ). We assume $L - a \rightarrow \infty$ and thus the annuitization factor $\gamma_{a,L} = \frac{r(1+r-\theta)}{(1+r)^2}$, where θ is the MA(1) parameter of the transitory income component. We set $r = 0.05$.

Table IX reports the results of the model for the whole sample, two representative cohorts (born in the 1940s and in the 1920s), and two education groups (with and without college education).²⁴

Starting with income growth parameters, note that both the variance of the permanent shock and the variance of the transitory shock are generally higher in 1985-92 than in 1979-84, which is expected from the evidence given above. The MA parameter for the transitory shock is small and generally not well measured. Turning to consumption parameters, note that the imputation error absorbs a large amount of the cross-sectional variability in consumption in the PSID, anything between 0.11 and 0.15. The variance of the imputation error σ_u^2 is always precisely measured. The variance of heterogeneity in the consumption slope is also sizable; it tends to be measured with little precision when we stratify our sample by education or decade of birth.

In the whole sample the estimate of ϕ , the partial insurance coefficient for the permanent shock, provides evidence in favor of partial insurance. In contrast, the evidence on ψ accords with a simple PIH model with infinite horizon. In no case do we reject the null that there is full smoothing with respect to transitory shocks ($\psi = 0$). The estimate of ϕ should be compared with the conventional belief, typical of simple consumption models, that permanent shocks to income are permanent shocks to consumption. As said above, prudent individuals will attempt to smooth this kind of fluctuation by a greater extent than predicted by a model with quadratic preferences; moreover, interpersonal insurance mechanisms or even public insurance will allow more smoothing than predicted by simple models where markets for insuring shocks are all shut down. The scope of the next subsections is to corroborate this evidence.²⁵

Finally note that the insurance coefficient ϕ decreases between the early 1980s and the late

²⁴Results for other cohorts are available on request. We could not achieve convergence for the cohort born in the 1950s for the 1979-84 period due to the small number of observations.

²⁵If we assume that food in the PSID reported in survey year t refers to that year rather than to the previous calendar year, we obtain similar results. The estimate of ϕ is slightly higher, but the qualitative pattern of results (and sensitivity checks) is unchanged.

1980s-early 1990s, suggesting that the degree of insurance has increased over this period. This may also, however, reflect the nature of the permanent shocks that occurred over this period rather than a change in the insurance mechanisms themselves.²⁶ Moreover, the trend is not uniform across groups: the less well educated and those born in the 1940s, for instance, face an increase in ϕ . Finally, it is not clear that a formal test will reject the null of no change.

When the sample is stratified by year of birth or education, we find qualitatively similar results: there is evidence for partial insurance with respect to the permanent shocks, and full insurance with respect to transitory shocks. The evidence across cohorts does not reveal any economically interesting heterogeneity in the degree of insurance available to consumers, perhaps because of small cell sizes that inflate standard errors; those with college education appear to be more able to smooth consumption in the face of permanent shocks to income. For individuals without college education and for the oldest cohort in our sample there is no statistically significant evidence for insurance against permanent income shocks. For these groups, permanent shocks to income prompt full consumption adjustment as in the traditional permanent income hypothesis model.

Finally, we note that the χ^2 goodness of fit statistics reveal some support for our model specification despite its simplicity for certain of the periods and household types.

4.4 Discussion: Insurance and Measurement

In this section we provide further interpretation of our results and discuss potential sources of bias in our estimates. Table X reports the results of various sensitivity checks where we change our definition of consumption and income, and examine the effect of extending our sample to the families of the SEO (the low-income subsample in the PSID), and of including young households in our baseline sample. We use whole sample moments throughout.

4.4.1 Insuring Income Shocks through Durables

Consider that the PIH could hold with respect to total consumption rather than non-durable consumption, the measure we use. The main consequence of this is that the Euler equation contains

²⁶Suppose for example that permanent earnings shocks are the combination of permanent wage and employment shocks, and that only the latter can be (partially) insured. Suppose further that permanent employment shocks become relatively more important than permanent wage shocks in the second subperiod. Then we might record an increase in insurance opportunities even if they are, in fact, stable over time or even declining.

an omitted variable, the growth in the fraction of total consumption that is devoted to non-durable expenditure. In this case $\phi_{b,t}$ will reflect, at least in part, the sensitivity of durable expenditure to income shocks. It is possible to correct for this by extending our analysis to a measure of consumption that includes durable expenditure. This is available in the CEX and our imputation procedure can easily handle such extension.²⁷ In columns (3) and (4) of Table X, we use a comprehensive measure of consumption that includes durables and nondurables.²⁸ This has a rather dramatic effect on our results, as we now find no evidence for partial insurance with respect to permanent shocks and again evidence for full insurance as far as transitory shocks are concerned. Indeed, we cannot reject the restriction of ϕ equal unity and imposing this leaves the remaining results very similar. This also suggests that it is unlikely that the information story of section 2.4 is quantitatively important, or else the estimate of ϕ would be less than one regardless of the consumption measure used.

These results suggest that much of the insurance we estimate in the baseline model for non-durables arises from optimal durable choice and timing, as argued, among others by Browning and Crossley [2001]. One might expect the ϕ coefficient to rise simply because durables are more income elastic than non durables. There are various other arguments in favor of including durables in our measure of consumption and why that may also explain the lack of evidence for insurance. In these arguments durable expenditure serves as an implicit insurance mechanism for non durable consumption. The stock of durables can be upgraded or downgraded in response to permanent shocks to income. Consider a permanent negative shock. In the absence of the durable hedge, one should reduce non durable consumption by the same amount of the shock. Downgrading one's house, car etc., and slowing the rate of replacement can help smoothing the non durable consumption effects of the permanent shock. A symmetric argument holds for a positive shock.

²⁷See Meyer and Sullivan [2001] for a detailed discussion of the measurement of durables in the CEX.

²⁸Total consumption includes food (at home and away from home), alcoholic beverages, tobacco, housing (utilities, fuels and public services, mortgage interests, property tax, maintenance and repairs, rents, other lodging, domestic services, textiles, furniture, floor coverings, appliances), clothing and footwear, transports (new and used cars, other vehicles, gasoline, vehicle finance charges and insurance, maintenance and repairs, rentals and leases, public transports), personal care, entertainment, health (insurance, prescription drugs, medical services), reading and education, cash contributions, and personal insurance (life insurance and retirement).

4.4.2 The Insurance Value of Taxes and Transfers

To see the impact of public insurance, suppose we exclude transfers (of any kind) from our measure of income. If taxes and transfers provide insurance for permanent income shocks, the insurance parameter in this specification should fall by an amount that reflects the degree of insurance. This happens because consumption still incorporates any insurance value of taxes and transfers but the new measure of income no longer does. The results of this experiment are reported in columns (5) and (6) of Table X. A comparison with the baseline results shows that the estimated insurance parameter declines from 0.6 to 0.52 in the early 1980s and from 0.48 to 0.26 in the second sub-period. That is, by excluding transfers the partial insurance coefficient drops on average by 30%, an estimate of the insurance provided by private and public transfers. This insurance can also be seen through the change in the estimated variance of permanent and transitory shocks. With taxes and transfers excluded, the variances of income shocks are indeed much higher.

4.4.3 Total Disposable Income

It has been suggested that people may use financial assets to hedge against labor market shocks, including permanent ones [Davis and Willen, 2001]. Another experiment we consider is to include income from assets in our definition of income. If portfolio choice is used to hedge against income risk, adding financial income back in would induce an increase in the estimate of ϕ , and such increase would reflect the amount of insurance provided by portfolio choice. The results (reported in columns (7) and (8)) are mixed, as in the original David and Willen's paper. There is some evidence in favor of this hypothesis, but the amount of insurance involved does not seem to be quantitatively important, probably due to low financial market participation.

4.4.4 Precautionary Asset Accumulation

The assumption $\pi_{i,a,t} \simeq 1$ made in Section 2.2 could be violated in the presence of precautionary asset accumulation. This could be particularly relevant for cohorts close to retirement, and it would signal partial insurance even when this is absent. However, focusing on young cohorts allows us to address this point directly, because young individuals are a long time from retirement and have very little precautionary assets to rely on. Nevertheless, one can notice that for the oldest cohort (where

the bias should be more severe) we find a very high partial insurance coefficient (see Table 9), while for young cohorts the insurance coefficient is similar to that estimated for the whole sample. Finally, while we cannot separately identify precautionary saving effects (π) from insurance effects (ϕ), we still pin down the degree of transmission of income shocks into consumption, as discussed in Section 2.2.

4.4.5 The Specification of Imputation Error

Our imputation procedure requires that the variance of predicted consumption differ from the variance of true consumption only by an additive term (which is estimated, see σ_u^2 in Table 9). But if in addition to this there is a scaling factor, the latter is generally not separately identifiable from the partial insurance parameter $\phi_{b,t}$. Thus the estimate of $\phi_{b,t}$ will be biased (upward, if the problem resembles that discussed in Appendix A.4). However, we have checked that this is not the case (see the discussion in the Appendix). As a further check, we use PSID food data directly without imputing. The estimates of the partial insurance coefficient reported in columns (9) and (10) of Table X (around 0.35) should be scaled by the food expenditure income elasticity (0.66, from Table IV) giving, e.g., $\frac{0.35}{0.66} = 0.53$, which is not far from the coefficient estimated with imputed data in the “Baseline” columns. This is indirect evidence that our imputation procedure is not responsible for the results.

4.4.6 Low income households

For a further sensitivity check, we include families from the poverty subsample of the PSID in our sample. The definition of consumption is the baseline one, i.e., non-durable. The results are in columns (11) and (12). Two pieces of evidence are worth mentioning: the estimate of ϕ is higher reflecting less insurance opportunities in this sample, and we would now reject full insurance with respect to transitory shocks, at least in the second sub-period (an estimate of ψ of 0.12, not far from the 0.2 benchmark found by other researchers, Hall and Mishkin, 1982). In fact, in the SEO subsample alone there is not much evidence for any insurance of permanent shocks (we do not reject the null that $\phi = 1$ regardless of the sub-period considered). For example, the estimate of ϕ in the second period is 1.05 and the estimate of ψ rises to 0.361, indicating an appreciable degree of sensitivity even to transitory shocks among the low income sample.

4.4.7 Young Households

In our final experiment (see columns (13) and (14)), we include young households (aged 20-29) in our sample. The results for ϕ are similar to those obtained for the baseline sample. The only relevant difference is that we find some evidence for no insurance against transitory shock, perhaps due to the fact that young households are more likely to be liquidity constrained.

5 Conclusions

The extensive research on the dynamics of income inequality in recent years has not been paralleled by comparable research on consumption inequality. Moreover, the limited research that is available for the US has reached contrasting conclusions (some papers find an increase in consumption inequality, while others find little or no change).²⁹ Finally, little attempt has been made to interpret the empirical findings using traditional consumption theory.³⁰

This paper used individual panel data on consumption and income to evaluate the degree of consumption insurance with respect to income shocks. Our framework allowed for self-insurance, in which consumers smooth idiosyncratic shocks through saving. It also considered the complete markets assumption in which all idiosyncratic shocks are insured. These two models sit amidst a wide range of missing insurance opportunities. We were able to assess the degree of insurance over and above self-insurance through savings. We did this by contrasting shifts in the cross-sectional distribution of income growth with shifts in the cross-sectional distribution of consumption growth, and analyzing the way these two measures of household welfare correlate over time. A major innovation of our study was to combine panel data on income from the PSID with consumption data from repeated CEX cross-sections in a structural way, i.e. using conventional demand analysis rather than reduced form imputation procedures. We also allowed a general form for heterogeneity.

Our results show a strong growth in permanent income shocks in the US during the early 1980s

²⁹Thus Cutler and Katz [1992] find evidence that consumption inequality increases in parallel with income inequality, while Slesnick [1994] shows that consumption inequality has historically fallen and raised only slightly in recent years.

³⁰Dynarski and Gruber [1997] regress consumption growth on income growth (instrumented to avoid the downward bias due to measurement error), and interpret the coefficient on income growth as measuring the extent of insurance against income shocks (i.e., as the discrepancy between income and consumption inequality). While their analysis is valuable, the interpretation of their estimates outside a structural consumption model is problematic. The modern theory of consumption makes a sharp distinction between the effect of anticipated and unanticipated income growth (and between transitory and permanent income growth), which is absent in their analysis.

(the variance of transitory shock also increases, but at a later stage). From trough to peak the variance of the permanent shock doubles, while the variance of the transitory shock only goes up by about 50%. The variance of permanent shocks levels off in the second half of the 1980s. The variance of the transitory shock is basically flat in the period where the variance of permanent shock is increasing, and it increases only when the variance of permanent shock slows down.

We find strong evidence against full insurance for permanent income shocks but not for transitory income shocks. Interestingly, this latter result needs adapting for the low income subsample where transitory shocks seem less insurable. Further there is evidence of partial insurance of permanent income shocks (albeit not uniform across groups); the results point to much of this partial insurance occurring through the adjustment of durable expenditures. We also find differences in the degree of insurance over time, although we are unable to distinguish changing insurance opportunities from changing nature of shocks. Finally we show that taxes and transfers provide an important insurance mechanism for permanent income shocks.

Our results have implications for both macroeconomics and labor economics. The macroeconomic literature has long been concerned with explaining why modern economies depart from the complete markets benchmark. Recent work has examined the role of asymmetric information, moral hazard, heterogeneity, etc., and asked whether the complete markets model can be amended to include some form of imperfect insurance. This issue has not been subject to a systematic empirical investigation. Insofar as lack of smoothing opportunities implies a greater vulnerability to income shocks, our research can be relevant to issues of the incidence and permanence of poverty studied in the labor economics literature. Studying how well families smooth income shocks, how this changes over time in response to changes in the economic environment confronted, and how different household types differ in their smoothing opportunities, is an important complement to understanding the effect of redistributive policies and anti-poverty strategies.

References

- [1] Abowd, J., and D. Card (1989), “On the covariance structure of earnings and hours changes”, *Econometrica*, **57**, 411-45.
- [2] Aiyagari, R. (1994), “Uninsured risk and aggregate saving”, *Quarterly Journal of Economics*, **109**, 659-84
- [3] Altonji, J., A.P. Martins, and A. Siow (2002), “Dynamic factor models of consumption, hours, and income”, *Research in Economics*, **56**, 3-59.
- [4] Altonji, J., and L. Segal (1996), “Small-sample bias in GMM estimation of covariance structures”, *Journal of Business and Economic Statistics*, **14**, 353-66.
- [5] Altonji, J., and A. Siow (1987), “Testing the response of consumption to income changes with (noisy) panel data”, *Quarterly Journal of Economics*, **102**, 293-28.
- [6] Altug, S., and R. Miller (1990), “Household choices in equilibrium”, *Econometrica*, **58**, 543-70.
- [7] Alvarez, F. and U. Jermann (2000), “Efficiency, equilibrium and asset pricing with risk of default”, *Econometrica*, **68**, 775-97.
- [8] Attanasio, O., and S. Davis (1996), “Relative wage movements and the distribution of consumption”, *Journal of Political Economy*, **104**, 1227-62.
- [9] Attanasio, O., and V. Rios Rull (2000), “Consumption smoothing in island economies: Can public insurance reduce welfare?”, *European Economic Review*, **44**, 1225-58.
- [10] Attanasio, O., and G. Weber (1995), “Is consumption growth consistent with intertemporal optimization: Evidence from the Consumer Expenditure Survey”, *Journal of Political Economy*, **103**, 1121-57.
- [11] Auerbach, A.J., and D. Feenberg (2000), “The significance of Federal Taxes as automatic stabilizers”, *Journal of Economic Perspectives*, **14** (Summer), 37-56.
- [12] Banks, J., Blundell, R. and A. Lewbel (1999), “Quadratic Engel curves and consumer demand”, *Review of Economics and Statistics*, **79**, 527-39.

- [13] Battistin, E. (2003), “Errors in survey reports of consumption expenditures”, IFS Working Paper 03/07.
- [14] Bernheim, D., J. Skinner, and S. Weinberg (2002), “What accounts for the variation in retirement wealth among U.S. households?”, *American Economic Review*, **91**, 832-57.
- [15] Blundell, R., and L. Pistaferri (2003), “Income volatility and household consumption: The impact of food assistance programs”, *Journal of Human Resources* (forthcoming).
- [16] Blundell, R., and I. Preston (1998), “Consumption inequality and income uncertainty”, *Quarterly Journal of Economics* **113**, 603-640.
- [17] Blundell, R., and T. Stoker (2002), “Aggregation in economic relationships: Heterogeneity and selection”, forthcoming in *Handbook of Econometrics*, vol. 6.
- [18] Browning, M., and T. Crossley (2001), “Shocks, stocks and socks: Consumption smoothing and the replacement of durables”, University of Copenhagen and McMaster University, mimeo..
- [19] Carroll, C. (2001), “Precautionary saving and the marginal propensity to consume out of permanent income”, NBER Working Paper 8233.
- [20] Chamberlain, G. (1984), “Panel data”, in *Handbook of Econometrics*, vol. 2, edited by Zvi Griliches and Michael D. Intriligator. Amsterdam: North-Holland.
- [21] Cutler, D., and L. Katz (1992), “Rising inequality? Changes in the distribution of income and consumption in the 1980s”, *American Economic Review*, **82**, 546-51.
- [22] Davis, S., and P. Willen (2001), “Using financial assets to hedge labor income risks”, University of Chicago, mimeo.
- [23] Deaton, A. (1995), *The Analysis of Household Surveys*. Baltimore: John Hopkins University Press.
- [24] Deaton, A. and J. Muellbauer (1980), “An Almost Ideal Demand System”, *American Economic Review*, 70, 312-26.

- [25] Deaton, A., and C. Paxson (1994), “Intertemporal choice and inequality”, *Journal of Political Economy*, **102**, 384-94.
- [26] Dynan, K. E. (2000), “Habit formation in consumer preferences: Evidence from panel data”, *American Economic Review*, **90**, 391-406.
- [27] Dynarski, S., and J. Gruber (1997), “Can families smooth variable earnings?”, *Brooking Papers on Economic Activity*, **1**, 229-305.
- [28] Engen, E., and J. Gruber (2001), “Unemployment insurance and precautionary savings”, *Journal of Monetary Economics*, **47**, 545-79.
- [29] Fay, S., E. Hurst, and M. White (2002), “The consumer bankruptcy decision”, *American Economic Review*, **92**, 706-18.
- [30] Garner T., K. Short, S. Shipp, C. Nelson, and G. Paulin (1998), “Experimental poverty measurement for the 1990s”, *Monthly Labor Review*, **121**, 39-68.
- [31] Gruber, J. (2000), “Cash welfare as a consumption smoothing mechanism for single mothers”, *Journal of Public Economics*, **75**, 157-82.
- [32] Gruber, J., and A. Yelowitz (1999), “Public health insurance and private savings”, *Journal of Political Economy*, **107**, 1249-74.
- [33] Guiso, L., L. Pistaferri, and F. Schivardi (2003), “Insurance within the firm”, University of Sassari, Stanford University and Bank of Italy, *mimeo*.
- [34] Hall, R., and F. Mishkin (1982), “The sensitivity of consumption to transitory income: Estimates from panel data of households”, *Econometrica*, **50**, 261-81.
- [35] Hayashi, F., J. Altonji, and L. Kotlikoff (1996), “Risk sharing between and within families”, *Econometrica*, **64**, 261-94.
- [36] Hill, M. (1992), *The Panel Study of Income Dynamics: A user’s guide*, Newbury Park, California: Sage Publications.

- [37] Hurst, E., and F. Stafford (2003), “Home is where the equity is: Liquidity constraints, refinancing and consumption”, *Journal of Money, Credit, and Banking*, forthcoming.
- [38] Jappelli, T., and L. Pistaferri (2003), “Intertemporal choice and consumption mobility”, University of Salerno and Stanford University, *mimeo*.
- [39] Kehoe, T., and D. Levine (2001), “Liquidity constrained markets versus debt constrained markets”, *Econometrica*, **69**, 575-98
- [40] Kniesner, T., and J. Ziliak (2002), “Tax reform and automatic stabilization”, *American Economic Review*, **92**, 590-612.
- [41] Kocherlakota, N. (1996), “Implications of efficient risk sharing without commitment”, *Review of Economic Studies*, **63**, 595-609
- [42] Kotlikoff, L., and A. Spivak (1981), “The family as an incomplete annuities market”, *Journal of Political Economy*, **89**, 372-91.
- [43] Krueger, D., and F. Perri (2001), “Risk sharing: Private insurance markets or redistributive taxes?”, University of Pennsylvania and New York University, *mimeo*.
- [44] Mace, B. (1991), “Full insurance in the presence of aggregate uncertainty”, *Journal of Political Economy*, **99**, 928-56.
- [45] MaCurdy, T. (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.
- [46] Mankiw, M., and M. Kimball (1992), “Precautionary saving and the timing of taxes”, *Journal of Political Economy*, **97**, 863-79.
- [47] Meghir, C., and L. Pistaferri (2003), “Income variance dynamics and heterogeneity”, *Econometrica* (forthcoming).
- [48] Meyer, B. and J. Sullivan (2001), “The living conditions of single mothers in the 1980s and 1990s”, NBER Working Paper 8298, May.

- [49] Moffitt, R., and P. Gottschalk (1994), "Trends in the autocovariance structure of earnings in the US: 1969-1987", Brown University, *mimeo*.
- [50] Nelson, J. (1994), "On testing for full insurance using Consumer Expenditure Survey data", *Journal of Political Economy* **102**, 384-394.
- [51] Skinner, J. (1987), "A superior measure of consumption from the Panel Study of Income Dynamics", *Economic Letters*, **23**, 213-16.
- [52] Slesnick, D. (1994), "Consumption, needs and inequality", *International Economic Review*, **35**, 677-703.
- [53] Stephens, M. Jr. (2002), "Worker displacement and the added worker effect", *Journal of Labor Economics*, **20**, 504-37.

Table I
Sample selection in the PSID

	# dropped	# remain
Initial sample (1968-1992)	0	145,940
Interviewed prior to 1978	52,408	93,532
Change in family composition	18,570	74,962
Female head	23,779	51,183
Missing values and topcoding	308	50,875
Change in marital status	5,882	44,993
Income outliers	2,407	42,586
Born before 1920 or after 1959	8,510	34,076
Poverty subsample	12,600	21,476
Aged less than 30 or more than 65	3,674	17,778

Table II
Sample selection in the CEX

	# dropped	# remain
Initial sample	0	141,289
Missing expenditure data	1,351	139,938
Present for less than 12 months	76,773	63,165
Observed after 1992	19,310	43,855
Zero before-tax income	1,308	42,547
Missing region or education	14,029	28,418
Marital status	5,848	22,570
Born before 1920 or after 1959	4,648	17,922
Aged less than 30 or more than 65	1,843	16,079
Income outliers	699	15,380

Table III
Comparison of means, PSID and CEX

	1980		1983		1986		1989		1992	
	PSID	CEX	PSID	CEX	PSID	CEX	PSID	CEX	PSID	CEX
Age	42.96	43.71	43.36	45.01	43.83	46.03	43.97	45.26	45.91	47.01
Family size	3.61	3.95	3.52	3.74	3.48	3.64	3.44	3.61	3.42	3.55
# of children	1.32	1.47	1.25	1.26	1.21	1.19	1.19	1.17	1.14	1.15
White	0.91	0.89	0.92	0.88	0.92	0.88	0.93	0.89	0.93	0.88
HS dropout	0.21	0.20	0.18	0.19	0.16	0.18	0.14	0.14	0.13	0.15
HS graduate	0.30	0.32	0.31	0.33	0.32	0.30	0.32	0.31	0.31	0.30
College dropout	0.49	0.48	0.51	0.48	0.53	0.52	0.54	0.56	0.56	0.55
Northeast	0.21	0.20	0.21	0.25	0.22	0.21	0.22	0.23	0.22	0.23
Midwest	0.33	0.28	0.31	0.26	0.30	0.27	0.30	0.28	0.30	0.29
South	0.31	0.29	0.31	0.28	0.30	0.27	0.31	0.27	0.30	0.26
West	0.15	0.24	0.17	0.21	0.18	0.25	0.18	0.23	0.18	0.22
Family income	32,759	29,078	37,907	35,587	45,035	43,473	52,919	50,690	61,911	55,956
Food at home	3,683	3,501	3,905	3,305	4,176	3,501	4,557	4,266	5,098	4,565
Food away	759	1,165	936	1,334	1,118	1,695	1,281	1,881	1,507	1,857
Husband's particip.	0.96	0.97	0.94	0.92	0.93	0.91	0.94	0.93	0.93	0.88
Wife's particip.	0.69	0.67	0.71	0.67	0.74	0.71	0.78	0.73	0.77	0.73

Table IV
The demand for food in the CEX

This table reports IV estimates of the demand equation for (the logarithm of) food at home in the CEX. We instrument the log of total nondurable expenditure (and its interaction with age, time education dummies) with the log of family before-tax income (and its interaction with age, time and education dummies). The estimates of the interaction of $\ln c$ with time dummies are omitted.

Variable	Estimate	Variable	Estimate
$\ln c$	0.6620 (0.0317)	Family size	0.0642 (0.0043)
$\ln c$ *One child	0.0885 (0.0280)	Family size*Northeast	0.0158 (0.0042)
$\ln c$ *Two children	-0.0573 (0.0372)	Family size*Midwest	-0.0053 (0.0053)
$\ln c$ *Three children+	0.1195 (0.0260)	Family size*South	0.0144 (0.0037)
$\ln c$ *HS dropout	-0.0124 (0.0331)	Born 1955-59	-0.1671 (0.0372)
$\ln c$ *HS graduate	0.0249 (0.0322)	Born 1950-54	-0.1093 (0.0324)
$\ln p$	-0.9003 (0.0630)	Born 1945-49	-0.0847 (0.0257)
$\ln p$ *Northeast	0.0390 (0.0740)	Born 1940-44	-0.0525 (0.0229)
$\ln p$ *Midwest	-0.0171 (0.0569)	Born 1935-39	-0.0389 (0.0158)
$\ln p$ *South	-0.1046 (0.0608)	Born 1930-34	-0.0131 (0.0137)
Aged 36-40	0.0266 (0.0134)	Born 1925-29	-0.0024 (0.0108)
Aged 41-45	0.0386 (0.0147)	Male participant	-0.0108 (0.0128)
Aged 46-50	0.0323 (0.0225)	Female participant	-0.0640 (0.0109)
Aged 51-55	0.0368 (0.00272)	Male part.*# of children	-0.0182 (0.0061)
Aged 56-60	0.0367 (0.0340)	Fem. part.*# of children	0.0060 (0.0056)
Aged 61-65	0.0067 (0.0405)	One child	-0.7692 (0.2739)
High school dropout	0.1956 (0.3174)	Two children	-0.4028 (0.3677)
High school graduate	-0.1900 (0.3122)	Three children+	-0.9920 (0.2646)
Northeast	-0.0477 (0.0183)	White	0.0784 (0.0129)
Midwest	-0.1243 (0.0155)	Food away/1000	-0.0613 (0.0042)
South	-0.0298 (0.0238)	Constant	1.5622 (0.3157)
F-test: $\ln c$ *year dummies (p-value)			<0.0001

Table V
Consumption and income variances and covariances,
CEX 1980-1992

	<i>Born 1920s</i>			<i>Born 1930s</i>		
	<i>var(c)</i>	<i>var(y)</i>	<i>cov(c, y)</i>	<i>var(c)</i>	<i>var(y)</i>	<i>cov(c, y)</i>
1980	0.2315	0.4860	0.4230	0.2184	0.4479	0.4009
1981	0.2500	0.5567	0.4789	0.2523	0.4618	0.4616
1982	0.2470	0.5874	0.4836	0.2344	0.4830	0.4591
1983	0.2678	0.5657	0.5358	0.2472	0.5424	0.5194
1984	0.2729	0.6053	0.5375	0.2712	0.6245	0.5412
1985	0.3060	0.7733	0.5836	0.2652	0.6013	0.4999
1986	0.2706	0.6963	0.5069	0.2970	0.6008	0.6067
1987	0.2849	0.7177	0.6205	0.2422	0.5404	0.4194
1988	0.2522	0.6562	0.4788	0.3014	0.7079	0.5498
1989	0.2765	0.4751	0.4691	0.3213	0.5681	0.5624
1990	0.2379	0.7406	0.4769	0.2840	0.6094	0.5154
1991	0.2757	0.4989	0.5591	0.2982	0.7166	0.6482
1992	0.3240	0.6334	0.5004	0.2747	0.6677	0.5030

	<i>Born 1940s</i>			<i>Born 1950s</i>		
	<i>var(c)</i>	<i>var(y)</i>	<i>cov(c, y)</i>	<i>var(c)</i>	<i>var(y)</i>	<i>cov(c, y)</i>
1980	0.2115	0.3646	0.3933	0.1155	0.2654	0.1562
1981	0.1921	0.4107	0.3609	0.2499	0.4438	0.3563
1982	0.2137	0.4904	0.4281	0.2291	0.4793	0.4194
1983	0.2030	0.4892	0.4060	0.3005	0.5715	0.4979
1984	0.2187	0.4915	0.4289	0.2438	0.5166	0.5005
1985	0.2371	0.4488	0.4198	0.2552	0.4765	0.4337
1986	0.3031	0.5443	0.5784	0.2678	0.5970	0.5374
1987	0.2263	0.5521	0.4638	0.2356	0.5250	0.4220
1988	0.2889	0.5664	0.4959	0.2262	0.5008	0.4373
1989	0.2461	0.5467	0.4791	0.2569	0.5514	0.4470
1990	0.2975	0.6343	0.5268	0.2252	0.5131	0.4243
1991	0.2931	0.5672	0.5028	0.2454	0.5756	0.4315
1992	0.2657	0.6230	0.4724	0.2286	0.5184	0.4204

Table VI
The autocovariance matrix of income growth

Year	$\text{var}(\Delta y_t)$	$\text{cov}(\Delta y_{t+1}, \Delta y_t)$	$\text{cov}(\Delta y_{t+2}, \Delta y_t)$
1980	0.0844 (0.0090)	-0.0225 (0.0042)	-0.0020 (0.0030)
1981	0.0821 (0.0093)	-0.0303 (0.0050)	-0.0039 (0.0035)
1982	0.0811 (0.0067)	-0.0242 (0.0040)	-0.0059 (0.0029)
1983	0.0881 (0.0094)	-0.0236 (0.0041)	-0.0092 (0.0053)
1984	0.0852 (0.0058)	-0.0313 (0.0039)	-0.0030 (0.0037)
1985	0.0917 (0.0069)	-0.0284 (0.0052)	-0.0012 (0.0042)
1986	0.1097 (0.0117)	-0.0427 (0.0093)	-0.0072 (0.0058)
1987	0.1138 (0.0112)	-0.0368 (0.0050)	0.0011 (0.0045)
1988	0.0893 (0.0083)	-0.0308 (0.0040)	-0.0014 (0.0031)
1989	0.0917 (0.0071)	-0.0292 (0.0076)	-0.0027 (0.0041)
1990	0.0977 (0.0137)	-0.0301 (0.0055)	-0.0055 (0.0045)
1991	0.1142 (0.0244)	-0.0417 (0.0142)	
1992	0.1309 (0.0153)		

Table VII
The autocovariance matrix of consumption growth

Year	$\text{var}(\Delta c_t)$	$\text{cov}(\Delta c_{t+1}, \Delta c_t)$	$\text{cov}(\Delta c_{t+2}, \Delta c_t)$
1980	0.2501 (0.0203)	-0.1345 (0.0177)	0.0029 (0.0109)
1981	0.2720 (0.0233)	-0.1256 (0.0163)	-0.0001 (0.0107)
1982	0.2841 (0.0233)	-0.1393 (0.0191)	0.0035 (0.0103)
1983	0.3122 (0.0276)	-0.1399 (0.0153)	-0.0062 (0.0104)
1984	0.3308 (0.0286)	-0.1455 (0.0250)	-0.0188 (0.0167)
1985	0.3255 (0.0292)	-0.1399 (0.0183)	NA
1986	0.3055 (0.0223)	NA	NA
1987	NA	NA	NA
1988	NA	NA	NA
1989	NA	NA	NA
1990	0.3182 (0.0315)	-0.1090 (0.0117)	-0.0065 (0.0102)
1991	0.3263 (0.0225)	-0.1614 (0.0197)	NA
1992	0.3065 (0.0232)	NA	NA

Table VIII
The consumption-income growth covariance matrix

Year	$\text{cov}(\Delta y_t, \Delta c_t)$	$\text{cov}(\Delta y_t, \Delta c_{t+1})$	$\text{cov}(\Delta y_{t+1}, \Delta c_t)$
1980	0.0037 (0.0058)	0.0072 (0.0056)	-0.0009 (0.0058)
1981	0.0152 (0.0058)	-0.0051 (0.0054)	-0.0048 (0.0050)
1982	0.0170 (0.0057)	-0.0003 (0.0062)	-0.0044 (0.0046)
1983	0.0268 (0.0066)	-0.0077 (0.0060)	-0.0101 (0.0066)
1984	0.0248 (0.0065)	-0.0090 (0.0051)	-0.0014 (0.0061)
1985	0.0127 (0.0056)	0.0017 (0.0061)	-0.0023 (0.0062)
1986	0.0226 (0.0058)	NA	-0.0009 (0.0075)
1987	NA	NA	NA
1988	NA	NA	NA
1989	NA	0.0089 (0.0062)	NA
1990	0.0041 (0.0067)	-0.0012 (0.0053)	0.0032 (0.0070)
1991	0.0111 (0.0062)	-0.0033 (0.0063)	-0.0015 (0.0065)
1992	0.0153 (0.0061)	NA	NA

Table IX
Minimum distance partial insurance and variance estimates

This table reports EWMD results of the parameters of interest: θ is the MA(1) coefficient of the transitory component of income, σ_u^2 the variance of the measurement error in consumption, σ_ξ^2 the variance of heterogeneity in the consumption slope, σ_ζ^2 the variance of permanent shocks to income, σ_ε^2 the variance of transitory shock to income, ϕ and ψ the partial insurance coefficients with respect to permanent and transitory income shocks, respectively. We assume $L - a \rightarrow \infty$, and an interest rate of 5 percent. Standard errors in parenthesis. For the goodness-of-fit statistic we report the p-value.

	<i>Whole sample</i>		<i>No College</i>		<i>College</i>		<i>Born 1940s</i>		<i>Born 1920s</i>	
	79-84	85-92	79-84	85-92	79-84	85-92	79-84	85-92	79-84	85-92
σ_u^2	0.1275 (0.0086)	0.1363 (0.0106)	0.1333 (0.0132)	0.1458 (0.0168)	0.1211 (0.0108)	0.1280 (0.0133)	0.1119 (0.0131)	0.1516 (0.0199)	0.1432 (0.0156)	0.1518 (0.0256)
σ_ξ^2	0.0915 (0.0447)	0.0681 (0.0416)	0.0171 (0.0166)	0.0226 (0.0230)	0.0144 (0.0116)	0.0445 (0.0173)	0.0156 (0.0139)	0.0238 (0.0301)	0.0030 (0.0221)	0.0232 (0.0448)
θ	0.0159 (0.0100)	0.0366 (0.0142)	0.1044 (0.0634)	0.0952 (0.0573)	0.0684 (0.0624)	0.0423 (0.0595)	0.0379 (0.1125)	0.0566 (0.0663)	0.0442 (0.0741)	0.1146 (0.0394)
σ_ζ^2	0.0223 (0.0033)	0.0309 (0.0047)	0.0249 (0.0054)	0.0267 (0.0063)	0.0200 (0.0039)	0.0341 (0.0066)	0.0195 (0.0046)	0.0262 (0.0074)	0.0327 (0.0075)	0.0231 (0.0233)
σ_ε^2	0.0333 (0.0033)	0.0395 (0.0044)	0.0389 (0.0053)	0.0420 (0.0057)	0.0272 (0.0039)	0.0374 (0.0064)	0.0236 (0.0046)	0.0361 (0.0057)	0.0421 (0.0076)	0.0975 (0.0356)
ϕ	0.5968 (0.1452)	0.4769 (0.1205)	0.7367 (0.2266)	0.7882 (0.2378)	0.4234 (0.1800)	0.2755 (0.1375)	0.4323 (0.2145)	0.5020 (0.2845)	0.7780 (0.2273)	0.9782 (1.0253)
ψ	0.0526 (0.0916)	-0.0400 (0.0798)	0.0894 (0.1183)	-0.1102 (0.1108)	0.0054 (0.1449)	0.0280 (0.1144)	0.2107 (0.1929)	0.0724 (0.1710)	-0.1532 (0.1623)	-0.1098 (0.1223)
χ^2	105.12 [0.0053]	132.83 [0.0005]	115.46 [0.0007]	110.29 [0.0288]	96.77 [0.0228]	118.83 [0.0074]	102.76 [0.0082]	143.45 [0.0000]	114.39 [0.0008]	189.57 [0.0000]

Table X
Sensitivity analysis

This table reports EWMD results of the parameters of interest: θ is the MA(1) coefficient of the transitory component of income, σ_u^2 the variance of the measurement error in consumption, σ_ξ^2 the variance of heterogeneity in the consumption slope, σ_ζ^2 the variance of permanent shocks to income, σ_ε^2 the variance of transitory shock to income, ϕ and ψ the partial insurance coefficients with respect to permanent and transitory income shocks, respectively. We assume $L - a \rightarrow \infty$, and an interest rate of 5 percent. Standard errors in parenthesis. For the goodness-of-fit statistic we report the p-value.

Consumption:	Non-durable		Total		Non-durable		Non-durable		Food		Non-durable		Non-durable	
Income:	Earn.+Transf.		Earn.+Transf.		Earnings		Total		Earn.+Transf.		Earn.+Transf.		Earn.+Transf.	
Sample:	Baseline		Baseline		Baseline		Baseline		Baseline		Baseline+SEO		Baseline+20s	
Period:	79-84	85-92	79-84	85-92	79-84	85-92	79-84	85-92	79-84	85-92	79-84	85-92	79-84	85-92
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
σ_u^2	0.1275 (0.0086)	0.1363 (0.0106)	0.3450 (0.0231)	0.3959 (0.0323)	0.1275 (0.0086)	0.1363 (0.0106)	0.1275 (0.0086)	0.1363 (0.0106)	0.0624 (0.0043)	0.0681 (0.0056)	0.1532 (0.0090)	0.1542 (0.0098)	0.1492 (0.0096)	0.1360 (0.0104)
σ_ξ^2	0.0915 (0.0447)	0.0681 (0.0416)	0.0895 (0.0475)	0.0666 (0.0417)	0.0659 (0.0442)	0.0218 (0.0651)	0.1005 (0.0529)	0.0615 (0.0432)	0.0908 (0.0442)	0.0680 (0.0417)	0.0806 (0.0407)	0.0846 (0.0290)	0.1203 (0.0420)	0.0884 (0.0372)
θ	0.0159 (0.0100)	0.0366 (0.0142)	0.0473 (0.0269)	0.1072 (0.0425)	0.0145 (0.0108)	0.0380 (0.0141)	0.0143 (0.0102)	0.0354 (0.0145)	0.0082 (0.0049)	0.0167 (0.0072)	0.0272 (0.0100)	0.0353 (0.0125)	0.0196 (0.0117)	0.0436 (0.0141)
σ_ζ^2	0.0223 (0.0033)	0.0309 (0.0047)	0.0225 (0.0035)	0.0310 (0.0046)	0.0345 (0.0057)	0.0810 (0.0126)	0.0212 (0.0027)	0.0226 (0.0026)	0.0224 (0.0033)	0.0309 (0.0047)	0.0233 (0.0031)	0.0269 (0.0034)	0.0211 (0.0030)	0.0299 (0.0044)
σ_ε^2	0.0333 (0.0033)	0.0395 (0.0044)	0.0332 (0.0035)	0.0394 (0.0044)	0.0539 (0.0066)	0.0684 (0.0100)	0.0270 (0.0025)	0.0334 (0.0027)	0.0333 (0.0033)	0.0395 (0.0044)	0.0397 (0.0031)	0.0450 (0.0034)	0.0350 (0.0032)	0.0399 (0.0041)
ϕ	0.5968 (0.1452)	0.4769 (0.1205)	0.9219 (0.2444)	0.8099 (0.2003)	0.5161 (0.1495)	0.2588 (0.0660)	0.6726 (0.1469)	0.6041 (0.1432)	0.3536 (0.0953)	0.3159 (0.0817)	0.8102 (0.1530)	0.6469 (0.1266)	0.5945 (0.1619)	0.4947 (0.1206)
ψ	0.0526 (0.0916)	-0.0400 (0.0798)	0.0889 (0.1516)	-0.0934 (0.1377)	0.0682 (0.0697)	0.0637 (0.0835)	0.0130 (0.1080)	0.0492 (0.0831)	0.0289 (0.0636)	-0.0373 (0.0561)	0.0761 (0.0765)	0.1184 (0.0640)	0.2138 (0.0921)	-0.0372 (0.0771)
χ^2	105.12 [0.0053]	132.83 [0.0005]	109.19 [0.0024]	131.26 [0.0008]	97.59 [0.0199]	102.41 [0.0840]	95.77 [0.0267]	141.87 [0.0001]	106.39 [0.0042]	131.72 [0.0007]	85.59 [0.1142]	131.40 [0.0007]	124.65 [0.0001]	137.49 [0.0002]




Figure 1: The Variance of Consumption, PSID and CEX.




Figure 2: The Mean of Consumption, PSID and CEX.

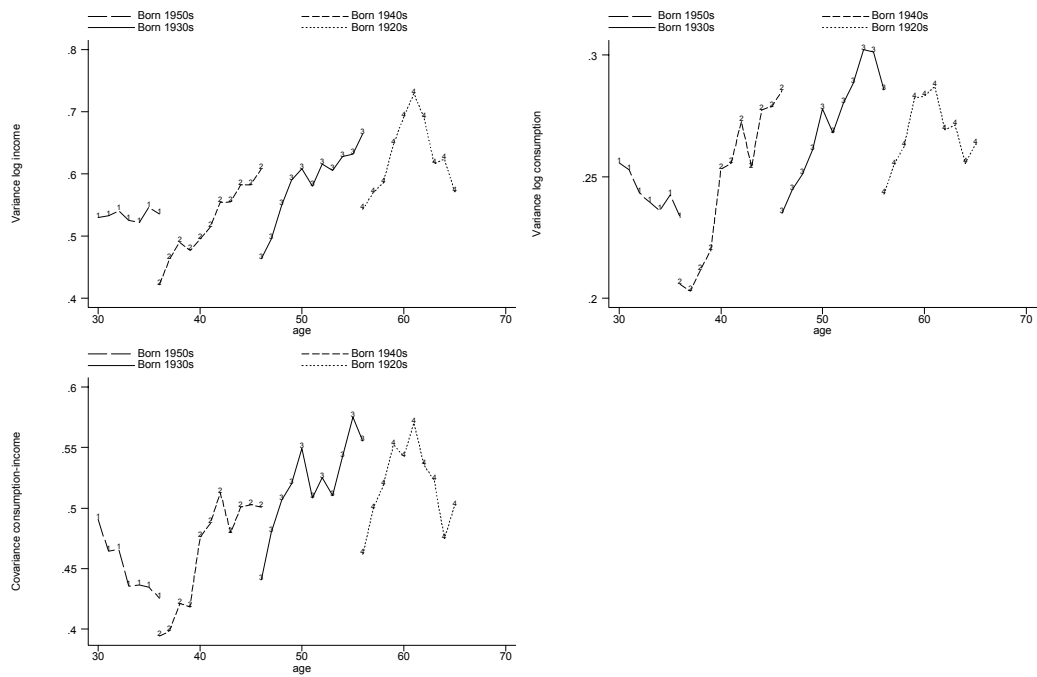


Figure 3: Cohort Profiles of the Variances of Log Income and Log Consumption and their Covariance, CEX 1980-1992.



Figure 4: The Variance of the Permanent Shock in the 1980s.



Figure 5: The Variance of the Transitory Shock in the 1980s

A.1 Appendix: The Euler Equation Approximation

Consider:

$$\ln \sum_{k=0}^{T-t} X_{t+k} = \ln X_t + \ln \left[1 + \sum_{k=1}^{T-t} \exp(\ln X_{t+k} - \ln X_t) \right]$$

Taking a Taylor expansion around $\ln X_{t+k} = \ln X_t + \sum_{i=0}^k \delta_{t+i}$, $k = 1, \dots, T-t$ with $\delta_t = 0$,

$$\begin{aligned} \ln \sum_{k=0}^{T-t} X_{t+k} &\simeq \ln X_t + \ln \left[1 + \sum_{k=1}^{T-t} \exp\left(\sum_{i=0}^k \delta_{t+i}\right) \right] \\ &\quad + \sum_{k=1}^{T-t} \frac{\exp\left(\sum_{i=0}^k \delta_{t+i}\right)}{\left[1 + \sum_{k=1}^{T-t} \exp\left(\sum_{i=0}^k \delta_{t+i}\right) \right]} (\ln X_{t+k} - \ln X_t - \sum_{i=0}^k \delta_{t+i}) \\ &\simeq \sum_{k=0}^{T-t} \alpha_{t+k,T}^{\delta} \ln X_{t+k} - \sum_{k=0}^{T-t} \alpha_{t+k,T}^{\delta} \ln \alpha_{t+k,T}^{\delta} \end{aligned}$$

where $\alpha_{t+k,T}^{\delta} = \exp\left(\sum_{i=0}^k \delta_{t+i}\right) / \left[1 + \sum_{k=1}^{T-t} \exp\left(\sum_{i=0}^k \delta_{t+i}\right) \right]$.

Take the consumption income model as in (4). With CRRA preferences, optimization implies the Euler equation

$$C_{i,a-1,t-1}^{-\gamma} = (1 + r_{t-1}) \frac{1 + D_{i,a-1,t-1}}{1 + D_{i,a,t}} E_{a-1,t-1} C_{i,a,t}^{-\gamma}$$

and therefore, approximately,

$$\Delta C_{i,a,t} = \eta_{i,a,t} + \omega_{i,a,t}$$

where $\eta_{i,a,t}$ is a consumption shock with $E_{a-1,t-1} \eta_{i,a,t} = 0$ and $\omega_{i,a,t}$ captures any slope in the consumption path due to interest rates, impatience, precautionary savings³¹ and so on.

From (3) we also have

$$\Delta y_{i,a+k,t+k} = \Delta Z'_{i,a+k,t+k} \varphi_t + \zeta_{i,a+k,t+k} + \sum_{j=0}^q \theta_j \varepsilon_{i,a+k-j,t+k-j}.$$

The intertemporal budget constraint is

$$\sum_{k=0}^{T-t} q_{t+k} C_{t+k} = \sum_{k=0}^{L-t} q_{t+k} Y_{t+k} + A_t$$

where T is death, L is retirement and q_{t+k} is appropriate discount factor $\prod_{i=1}^k (1 + r_{t+i})$, $k = 1, \dots, T-t$ (and $q_t = 1$). Using the above approximation

$$\begin{aligned} &\sum_{k=0}^{T-t} \alpha_{t+k,T}^{\omega-r} [\ln C_{t+k} - \ln q_{t+k} - \ln \alpha_{t+k,T}^{\omega-r}] \\ &\simeq \pi_{i,a,t} \sum_{k=0}^{L-t} \alpha_{t+k,L}^{\Delta Z \varphi - r} [\ln Y_{t+k} - \ln q_{t+k} - \ln \alpha_{t+k,L}^{\Delta Z \varphi - r}] \\ &\quad + (1 - \pi_{i,a,t}) \ln A_t - [(1 - \pi_{i,a,t}) \ln(1 - \pi_{i,a,t}) + \pi_{i,a,t} \ln \pi_{i,a,t}] \end{aligned}$$

where $\pi_{i,a,t} = \frac{\sum_{k=0}^{L-t} q_{t+k} Y_{t-k}}{\sum_{k=0}^{L-t} q_{t+k} Y_{t-k} + A_t}$ is the share of future labor income in current human and financial wealth.

³¹In particular, by the nature of the approximation, $\omega_{i,a+1,t+1}$ contains a term reflecting the conditional variance of $\eta_{i,a+1,t+1}$.

Taking differences in expectations and allowing for revisions to expectations of ω (because, say, of revised expectations regarding future consumption variability) gives

$$\begin{aligned}
\eta_{i,a,t} &\simeq \pi_{i,a,t} \left[\zeta_{i,a,t} + \left(\sum_{j=0}^q \alpha_{t+j,L}^{\Delta Z \varphi^{-r}} \theta_j \right) \varepsilon_{i,a,t} \right] \\
&\quad + \sum_{k=0}^{T-t} (E_t - E_{t-1}) (\alpha_{t+k,T}^{\omega^{-r}} \sum_{j=1}^k [\omega_{i,a+j,t+j} + r_{t+j}]) \\
&\quad - \sum_{k=0}^{T-t} (E_t - E_{t-1}) (\alpha_{t+k,T}^{\omega^{-r}} \ln \alpha_{t+k,T}^{\omega^{-r}}) \\
&= \pi_{i,a,t} [\zeta_{i,a,t} + \gamma_{t,L} \varepsilon_{i,a,t}] \\
&\quad + (E_t - E_{t-1}) \ln \left[1 + \sum_{j=1}^k [\omega_{i,a+j,t+j} + r_{t+j}] \right]
\end{aligned}$$

Suppose that any idiosyncratic gradient to the consumption path due to impatience, demographic change and so on, as captured in $\omega_{i,a,t}$ or revisions to future expected values of $\omega_{i,a,t}$, can be adequately picked up by a cohort/time-specific component $\Gamma_{b,t}$, an individual element $\xi_{i,a,t}$ and a component associated with observables $\Delta Z'_{i,a,t} \vartheta_t$

$$\omega_{i,a,t} + (E_t - E_{t-1}) \ln \left[1 + \sum_{j=1}^k [\omega_{i,a+j,t+j} + r_{t+j}] \right] \simeq \Gamma_{b,t} + \Delta Z'_{i,a,t} \vartheta_t + \xi_{i,a,t}.$$

Then

$$\Delta c_{i,a,t} \simeq \Gamma_{b,t} + \xi_{i,a,t} + \Delta Z'_{i,a,t} \vartheta_t + \pi_{i,a,t} \zeta_{i,a,t} + \gamma_{t,L} \pi_{i,a,t} \varepsilon_{i,a,t}.$$

For large $L-t$ we may be prepared to assume that $\gamma_{t,L} = \sum_{j=0}^q \alpha_{t+j,L}^{\Delta Z \varphi^{-r}} \theta_j$ is small enough to be ignored and $\pi_{i,a,t} \simeq 1$.

If $\Delta Z'_{i,t+k} \varphi_t = 0$ and $r_t = r$ is constant then $\alpha_{t+j,L}^{\Delta Z \varphi^{-r}} = \alpha_{t+j,L}^{-r} = \exp(-jr) / \sum_{k=0}^L \exp(-kr) \simeq r / (1+r)^k$ and $\gamma_{t,L} \simeq \frac{r}{1+r} [1 + \sum_{j=1}^q \theta_j / (1+r)^j]$.

A.2 Appendix: Measurement error

In the light of our imputation procedure, let's assume that both consumption and income are measured with multiplicative (independent and quasi-classical) error,³² e.g., $y_{i,a,t}^* = y_{i,a,t} + u_{i,a,t}^y$ and $c_{i,a,t}^* = c_{i,a,t} + u_{i,a,t}^c$, where x^* denote a measured variable, x its true, unobservable value, and u the measurement error. Equations (14) and (15) then rewrite as:

$$\begin{aligned}
\Delta y_{i,a,t}^* &= \zeta_{i,a,t} + \Delta \varepsilon_{i,a,t} + \Delta u_{i,a,t}^y \\
\Delta c_{i,a,t}^* &\simeq \phi_{b,t} \zeta_{i,a,t} + \psi_{b,t} \alpha_a \varepsilon_{i,a,t} + \Delta u_{i,a,t}^c
\end{aligned}$$

Consider first the case $L-a \rightarrow \infty$, $r \rightarrow 0$, and $\gamma_{a,L} \rightarrow 0$. Assume that measurement error in income is orthogonal to measurement error in consumption. Note that the ‘‘measured’’ ratio (16) $\frac{\text{var}(\Delta c_{a,t}^*)}{\text{cov}(\Delta y_{a,t}^*, \Delta c_{a,t}^*)}$ no longer identifies $\phi_{b,t}$ because of measurement error in consumption. However, the expression:

$$\phi_{b,t} = \frac{\text{cov}(\Delta c_{a,t}^*, \Delta y_{a,t}^*)}{\text{cov}(\Delta y_{a,t}^*, \Delta y_{a-1,t-1}^* + \Delta y_{a,t}^* + \Delta y_{a+1,t+1}^*)}$$

³²Quasi-classical in the sense that the variance of the measurement error is not (necessarily) constant over time.

still identifies the partial insurance parameter $\phi_{b,t}$. Similarly:

$$\phi_{b,t} = \frac{\text{cov}(\Delta c_{a,t}^*, \Delta c_{a-1,t-1}^* + \Delta c_{a,t}^* + \Delta c_{a+1,t+1}^*)}{\text{cov}(\Delta c_{a,t}^*, \Delta y_{a,t}^*)}$$

showing that $\phi_{b,t}$ is again overidentified. Under the martingale assumption for consumption, $\text{var}(u_{a,t}^c)$ can be identified using the covariance of current and lagged consumption growth:

$$\text{var}(u_{a,t}^c) = -\text{cov}(\Delta c_{a+1,t+1}^*, \Delta c_{a,t}^*) \quad (\text{A2.1})$$

However, $\text{var}(\varepsilon_{a,t})$ and $\text{var}(u_{a,t}^y)$ cannot be told apart, and $\psi_{b,t}$ thus remains unidentified. If $L - a$ is finite and $r \neq 0$, then:

$$\phi_{b,t} = \frac{\text{cov}(\Delta c_{a,t}^*, \Delta y_{a,t}^* + \Delta y_{a+1,t+1}^*)}{\text{cov}(\Delta y_{a,t}^*, \Delta y_{a-1,t-1}^* + \Delta y_{a,t}^* + \Delta y_{a+1,t+1}^*)}$$

This identifies the partial insurance parameter $\phi_{b,t}$. Once more, $\text{var}(\varepsilon_{a,t})$ and $\text{var}(u_{a,t}^y)$ cannot be separately identified and $\psi_{b,t}$ remains unidentified. It is possible however to put an upper bound on $\psi_{b,t}$ using the fact that:

$$\psi_{b,t} \leq \alpha_a^{-1} \frac{\text{cov}(\Delta c_{a,t}^*, \Delta y_{a+1,t+1}^*)}{\text{cov}(\Delta y_{a+1,t+1}^*, \Delta y_{a,t}^*)}.$$

Thus it is possible to argue that the estimate of $\psi_{b,t}$ is upward biased due to measurement error in income. Given that in the empirical analysis $\psi_{b,t}$ is close to zero in most cases, this problem does not seem to be empirically important.

A.3 Appendix: Estimation details

The two basic vectors of interest are:

$$\mathbf{c}_i = \begin{pmatrix} \Delta c_{i,1} \\ \Delta c_{i,2} \\ \dots \\ \Delta c_{i,T} \end{pmatrix} \text{ and } \mathbf{y}_i = \begin{pmatrix} \Delta y_{i,1} \\ \Delta y_{i,2} \\ \dots \\ \Delta y_{i,T} \end{pmatrix}$$

where, for simplicity, we indicate with 0 the first year in the panel (1978) and with T the last (1992), and the reference to age has been omitted. Conformably with the vectors above, define:

$$\mathbf{d}_i = \begin{pmatrix} d_{i,1} \\ d_{i,2} \\ \dots \\ d_{i,T} \end{pmatrix}$$

where $d_{i,t} = 1$ $\{y_{i,t}, c_{i,t}$ are not missing $\}$. This means that only complete observations on these two variables are used. This notation allows us to handle the problem of unbalanced panel data and the fact that the PSID did not collect consumption data in 1987 and 1988 in a simple manner.

Stacking observations on Δy and Δc for each individual we obtain the vector:

$$\mathbf{x}_i = \begin{pmatrix} \mathbf{c}_i \\ \mathbf{y}_i \end{pmatrix}$$

Now we can derive:

$$\mathbf{m} = \text{vech} \left\{ \sum_{i=1}^N (\mathbf{x}_i \mathbf{x}_i') \oslash \mathbf{d}_i \mathbf{d}_i' \right\}$$

where \oslash denotes an elementwise division. The vector \mathbf{m} contains the estimates of $cov(\Delta y_t, \Delta y_{t+s})$, $cov(\Delta y_t, \Delta c_{t+s})$, and $cov(\Delta c_t, \Delta c_{t+s})$, a total of $T(2T+1)$ unique moments). To obtain the variance-covariance matrix of \mathbf{m} , define conformably with \mathbf{m} the individual vector:

$$\mathbf{m}_i = (\mathbf{x}_i \mathbf{x}_i') \oslash \mathbf{d}_i \mathbf{d}_i'$$

The variance-covariance matrix of \mathbf{m} that can be used for inference is:

$$\mathbf{V} = \left[\sum_{i=1}^N (\mathbf{m}_i - \mathbf{m})(\mathbf{m}_i - \mathbf{m})' \otimes (\mathbf{d}_i \mathbf{d}_i') \right] \oslash \mathbf{D} \mathbf{D}'$$

where $\mathbf{D} = \text{vech} \left\{ \sum_{i=1}^N \mathbf{d}_i \mathbf{d}_i' \right\}$ and \otimes denotes an elementwise product. The square roots of the elements in the main diagonal of \mathbf{V} provide the standard errors of the corresponding elements in \mathbf{m} .

What we do in the empirical analysis is to estimate models for \mathbf{m} :

$$\mathbf{m} = f(\mathbf{\Lambda}) + \Upsilon$$

where Υ captures sampling variability and $\mathbf{\Lambda}$ is the vector of parameters we are interested in (the variances of the permanent shock and the transitory shock, the partial insurance parameters, etc.). For instance the mapping from \mathbf{m} to $f(\mathbf{\Lambda})$ is:

$$\begin{pmatrix} \text{var}(\Delta c_1) \\ \text{cov}(\Delta c_1, \Delta c_2) \\ \dots \\ \text{cov}(\Delta c_1, \Delta c_T) \\ \dots \end{pmatrix} = \begin{pmatrix} \phi^2 \text{var}(\zeta_1) + \psi^2 \alpha^2 \text{var}(\varepsilon_1) + \text{var}(\xi_1) + \text{var}(\eta_1^c) + \text{var}(\eta_0^c) \\ -\text{var}(\eta_1^c) \\ \dots \\ 0 \\ \dots \end{pmatrix} + \Upsilon$$

We solve the problem of estimating $\mathbf{\Lambda}$ by minimizing:

$$\min_{\mathbf{\Lambda}} (\mathbf{m} - f(\mathbf{\Lambda}))' \mathbf{A} (\mathbf{m} - f(\mathbf{\Lambda}))$$

where \mathbf{A} is a weighting matrix. Optimal minimum distance (OMD) imposes $\mathbf{A} = \mathbf{V}^{-1}$, equally weighted minimum distance (EWMD) imposes $\mathbf{A} = \mathbf{I}$, and variance-weighted minimum distance (VWMD) requires that \mathbf{A} is a diagonal matrix with the elements in the main diagonal given by $\text{diag}(\mathbf{V}^{-1})$.

For inference purposes we require the computation of standard errors. Chamberlain [1984] shows that these can be obtained as:

$$\widehat{\text{var}}(\widehat{\mathbf{\Lambda}}) = (\mathbf{G}' \mathbf{A} \mathbf{G})^{-1} \mathbf{G}' \mathbf{A} \mathbf{V} \mathbf{A} \mathbf{G} (\mathbf{G}' \mathbf{A} \mathbf{G})^{-1}$$

where $\mathbf{G} = \frac{\partial f(\mathbf{\Lambda})}{\partial \mathbf{\Lambda}} \Big|_{\mathbf{\Lambda}=\widehat{\mathbf{\Lambda}}}$ is the Jacobian matrix evaluated at the estimated parameters $\widehat{\mathbf{\Lambda}}$. The validity of the overidentifying restrictions can be tested by constructing the test statistic:

$$J = (\mathbf{m} - f(\widehat{\mathbf{\Lambda}}))' \mathbf{R}^- (\mathbf{m} - f(\widehat{\mathbf{\Lambda}}))$$

in the EWMD case. Here \mathbf{R}^- is a generalized inverse of $\mathbf{R} = \mathbf{W} \mathbf{V} \mathbf{W}$ with $\mathbf{W} = \mathbf{I} - \mathbf{G} (\mathbf{G}' \mathbf{G})^{-1} \mathbf{G}'$. Under the null hypothesis that the model is valid, the test statistic J is asymptotically distributed χ^2 with degrees of freedom equal the difference between the number of elements in \mathbf{m} and the rank of the gradient matrix \mathbf{G} . This test statistics is reported at the bottom of Tables 9 and 10.

A.4 Appendix: The imputation procedure

Consider the demand equation for food at home (21) in the CEX:

$$f_{i,x} = W'_{i,x} \mu + \beta (Z_{i,x}) c_{i,x} + e_{i,x}$$

where we use the subscript x to indicate an observation from the CEX and the subscript p to indicate an observation from the PSID. Define *imputed* consumption in the CEX by inverting assuming $\hat{\beta}(Z_{i,x}) \neq 0$:

$$\hat{c}_{i,x} = \hat{\beta}(Z_{i,x})^{-1} (f_{i,x} - W'_{i,x} \hat{\mu})$$

where a caret indicates a consistent estimate. The corresponding *imputed* measure of consumption in the PSID is

$$\hat{c}_{i,p} = \hat{\beta}(Z_{i,p})^{-1} (f_{i,p} - M'_{i,p} \hat{\mu})$$

To understand under which conditions moments of imputed PSID consumption mirror those of “true” consumption, note that we are confronted with a (non-standard) measurement error problem of the form:

$$\hat{c}_{i,x} = \hat{\beta}(Z_{i,x})^{-1} \beta(Z_{i,x}) c_{i,x} + W'_{i,x} \hat{\beta}(Z_{i,x})^{-1} (\mu - \hat{\mu}) + u_{i,x}$$

and $u_{i,x} = \hat{\beta}(Z_{i,x})^{-1} e_{i,x}$. Consider for simplicity the single-variable regression case:

$$\hat{c}_{i,x} = \frac{(\mu - \hat{\mu})}{\hat{\beta}} + \frac{\beta}{\hat{\beta}} c_{i,x} + u_{i,x}$$

and define with $M(x) = \frac{\sum_{i=1}^N x_i}{N}$ and $V(x) = \frac{\sum_{i=1}^N (x_i - M(x))^2}{N}$ the sample cross-sectional mean and variance of the variable x . Let us consider two cases of interest.

The first case is when $c_{i,x}$ is measured without error. In this case, $\text{plim } \hat{\beta} = \beta$ and $\text{plim } \hat{\mu} = \mu$. It follows that:

$$\text{plim } M(\hat{c}_x) = \text{plim } M(c_x)$$

and:

$$\text{plim } V(\hat{c}_x) = \text{plim } V(c_x) + \frac{1}{\beta^2} \text{plim } V(e_x)$$

Thus the sample mean of predicted CEX consumption converges to the same limit of the sample mean of true consumption, while the sample variance of predicted CEX consumption converges to the limit of the variance of true consumption up to an additive term. The latter decreases with the value of the expenditure elasticity. If the demand for food at home is relatively inelastic ($\beta \rightarrow 0$) the additive term may be potentially quite large.

As for PSID imputed consumption, it is easy to prove that:

$$\text{plim } M(c_p) = \text{plim } M(c_x) + \frac{1}{\beta} [\text{plim } M(f_p) - \text{plim } M(f_x)]$$

and:

$$\text{plim } V(c_p) = \text{plim } V(c_x) + \frac{1}{\beta^2} \text{plim } V(e_x) + \frac{1}{\beta^2} [\text{plim } V(f_p) - \text{plim } V(f_x)] \quad (\text{A4.1})$$

Thus the sample mean of imputed PSID consumption converges to the limit of the sample mean of true consumption up to an additive term (the mean difference in the input variable available in both surveys, e.g., food consumption, scaled by the expenditure elasticity). If food consumption is on average the same in the two data sets, the sample mean of imputed PSID consumption converges to the same limit of the sample mean of true consumption. Otherwise, the sample mean of imputed PSID consumption may overestimate or underestimate the sample mean of true consumption. It is possible to correct for this discrepancy by using, e.g., $\frac{M(f_p) - M(f_x)}{\beta}$ as a correction factor. This is what we do in Figure 2.

Note also that the sample variance of imputed PSID consumption differs from the variance of true consumption because of two factors: the difference between the variance of food consumption in the two data sets, scaled by the square of the expenditure elasticity, and the variance of food heterogeneity, again

scaled by the square of the elasticity.³³ Our minimum distance procedure is designed to estimate this factor using, e.g., (A2.1).

In the second case, $c_{i,x}$ is measured with classical error: $c^* = c + u$. It follows that $\text{plim } \hat{\beta} = \frac{\beta}{\lambda}$ and $\text{plim } \hat{\mu} = \mu + \beta \frac{(\lambda-1)}{\lambda} \text{plim } M(c)$, where $\lambda = \frac{\text{plim } V(c^*)}{\text{plim } V(c)} \geq 1$. Repeating the same steps above:

$$\text{plim } M(\hat{c}_x) = \text{plim } M(c_x) = \text{plim } M(c_x^*)$$

and:

$$\begin{aligned} \text{plim } V(\hat{c}_x) &= \lambda \text{plim } V(c_x) + \left(\frac{\lambda}{\beta}\right)^2 \text{plim } V(e_x) \\ &= \lambda^2 \text{plim } V(c_x^*) + \left(\frac{\lambda}{\beta}\right)^2 \text{plim } V(e_x) \end{aligned}$$

As for PSID imputed consumption,

$$\begin{aligned} \text{plim } M(\hat{c}_p) &= \text{plim } M(c_x) + \frac{\lambda}{\beta} [\text{plim } M(f_p) - \text{plim } M(f_x)] \\ &= \text{plim } M(c_x^*) + \frac{\lambda}{\beta} [\text{plim } M(f_p) - \text{plim } M(f_x)] \end{aligned}$$

and:

$$\begin{aligned} \text{plim } V(c_p) &= \lambda \text{plim } V(c_x) + \left(\frac{\lambda}{\beta}\right)^2 \text{plim } V(e_x) + \left(\frac{\lambda}{\beta}\right)^2 [\text{plim } V(f_p) - \text{plim } V(f_x)] \\ &= \lambda^2 \text{plim } V(c_x^*) + \left(\frac{\lambda}{\beta}\right)^2 \text{plim } V(e_x) + \left(\frac{\lambda}{\beta}\right)^2 [\text{plim } V(f_p) - \text{plim } V(f_x)] \quad (\text{A4.2}) \end{aligned}$$

In general, the presence of classical measurement error makes the discrepancy between moments of imputed consumption and moments of true consumption worse than in the case where c is measured without error.

This discussion suggests the use of an instrumental variable procedure in the attempt of minimizing the impact of biases in the estimated coefficients β and μ (see Table 4 for the results).

Panel A of Figure 2 shows that our imputed PSID consumption overestimates CEX consumption. This may have various explanations. First, despite our focus on homogeneous samples, the demographic characteristics in the PSID may still be quite different than those in the CEX (see Table 3). The second possibility is that our IV procedure may not eliminate the bias in the parameters of the demand equation due to measurement error in non-durable expenditure. In this case, overestimation is expected. Finally, and more importantly, the input variable (food expenditure) is measured differently in the two data sets, and measurement of this variable has changed over time, as we noticed in the main text.³⁴ This warrants the correction shown in Figure 2.

³³It also follows that the variances of first differences $V(\Delta\hat{c}_x)$ will exceed $V(\Delta c_x)$ by the term $\frac{V(\Delta e_x)}{\beta^2}$, and that $V(\Delta\hat{c}_p)$ will exceed $V(\Delta c_x)$ by the term

$$\frac{V(\Delta e_x)}{\beta^2} + \frac{1}{\beta^2} [\text{plim } V(\Delta f_p) - \text{plim } V(\Delta f_x)].$$

³⁴It is worth noting that while the matching of mean consumption is desirable, it is not necessary, in that the scope of the empirical analysis is to estimate models for the variance of consumption and its covariance with income.

To see whether measurement error is an issue in our imputations procedure, we compute the slope of the relationship between $V(\hat{c}_p)$ and $V(c_x^*)$ in two cases: OLS and IV. The OLS slope is 1.15, the IV is 0.99.³⁵ The OLS intercept is also larger than the IV intercept. These findings are indeed predictable by the analysis above. To see this point, note that equation (A4.1) and (A4.2) may be interpreted as referring to the OLS and IV case, respectively, if consumption in the CEX is measured with classical error and a valid instrument is available. From the comparison of these two equations two things can be noticed. First, the slope of the relationship between $V(\hat{c}_p)$ and $V(c_x^*)$ is unity in the IV case and greater than one in the OLS (biased) case. Second, the intercept of the relationship is greater in the OLS case than in the IV case. Both predictions appear to be supported.

While this evidence can only be taken as suggestive, it provides some useful information. First, an IV adjustment seems warranted. Second, once the adjustment is made, the moments of imputed PSID consumption mirror closely those of “true” consumption (see Figures 1 and 2).

³⁵These are obtained by simple OLS regressions of $V(\hat{c}_p^{\text{OLS}})$ and $V(\hat{c}_p^{\text{IV}})$ on $V(c_x^*)$.