

CONSUMPTION INEQUALITY AND PARTIAL INSURANCE*

Richard Blundell
University College London and Institute for Fiscal Studies.

Luigi Pistaferri
Stanford University and CEPR.

Ian Preston
University College London and Institute for Fiscal Studies.

February 2008

Abstract

This paper examines the link between income inequality and consumption inequality through the degree of insurance to income shocks. A new panel data consumption series is created for the PSID using an imputation procedure that maps food data into consumption data using the estimates of a demand function for food obtained from repeated CEX cross-sections. We document a disjuncture between income and consumption inequality in the US over the 1980s and show that it can be explained by changes in the persistence of income shocks. In particular, an initial growth in the variance of permanent shocks is then replaced by a continued growth in the variance of transitory income shocks. Although we find important differences in the degree of insurance by wealth, education, and birth cohort, the overall interpretation of the relationship between consumption and income inequality is maintained. We find some partial insurance of permanent income shocks with more insurance possibilities for the college educated and those approaching retirement. We find little evidence against full insurance for transitory income shocks except among low wealth households. Taxes and transfers as well as family labor supply are found to play an important role in insuring permanent shocks.

Key words: Consumption, Insurance, Inequality.

JEL Classification: D52; D91; I30.

*Blundell: Department of Economics, University College London, and Institute for Fiscal Studies, Gower Street, London WC1E 6BT, UK (email: r.blundell@ucl.ac.uk); Pistaferri: Department of Economics, Stanford University, Stanford, CA 94305 (email: pista@stanford.edu); Preston: Department of Economics, University College London, and Institute for Fiscal Studies, Gower Street, London WC1E 6BT, UK (email: i.preston@ucl.ac.uk). We would like to thank three anonymous referees, Joe Altonji, Orazio Attanasio, Giacomo De Giorgi, David Johnson, Arie Kapteyn, John Kennan, Robert Lalonde, Hamish Low, Bruce Meyer, Samuel Pienknagura, Gianluca Violante, Ken West and seminar participants at various institutions for helpful comments. Thanks are also due to Cristobal Huneus and Sonam Sherpa 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 (Blundell and Preston), the Joint Center for Poverty Research/Department of Health and Human Services, and the National Science Foundation under grant SES-0214491 (Pistaferri) is gratefully acknowledged. All errors are ours.

1 Introduction

While there is extensive work documenting changes in the wage and household income distributions over the 1980s and 1990s, there is relatively less work on the corresponding changes in the consumption distribution. Cutler and Katz (1992) and Johnson and Smeeding (1998) are notable exceptions. However, both studies are primarily descriptive and do not attempt to uncover the link between changes in income inequality and changes in consumption inequality. The goal of this paper is instead to analyze precisely such link.¹ To do it, we create a new panel series of consumption that combines information from the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX). We focus on the period between the end of the 1970s and the early 1990s when some of the largest changes in income inequality occurred. We show that the empirical relationship between the evolution of the consumption distribution and the evolution of the income distribution over this period can be characterized by the degree of persistence of the underlying income shocks and the degree of consumption insurance with respect to shocks of different durability. We argue that this representation provides a compelling framework for understanding the shifts in the consumption and income distributions.

Our analysis shows that, during the sampling period we study, income and consumption inequality diverged. We find that this can be explained by the change in the durability of income shocks over this period. In particular, an initial growth in the variance of permanent shocks was then replaced by a continued growth in the variance of transitory income shocks in the late 1980s. We find little evidence that the degree of insurance with respect to shocks of different durability changes over this period. In other words, rather than greater insurance opportunities, it is the relative increase in the variability of more insurable shocks that explains the disjuncture between income and consumption inequality over this period. We find important differences in the degree of insurance by wealth, education and birth cohort, but our interpretation of the relationship between consumption and income inequality is preserved.

The connection between consumption insurance and income shocks has a long history in economics. Two

¹Blundell and Preston (1998), Krueger and Perri (2006), and Heathcote, Storesletten and Violante (2004) have a similar goal. Below we discuss the relationship between these papers and ours.

polar models have dominated the agenda. On the one hand, the complete markets hypothesis assumes that consumption is fully insured against idiosyncratic shocks to income, both transitory and permanent. This hypothesis is typically rejected in micro data (Attanasio and Davis, 1996). On the other hand, the textbook permanent income hypothesis assumes that personal saving is the only mechanism available to agents to smooth income shocks. 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 (Deaton, 1992). In both aggregate and micro data, however, consumption appears to be excessively smooth, i.e., it reacts too little to permanent income shocks to be consistent with the theory (Campbell and Deaton, 1989; Attanasio and Pavoni, 2006). In other studies, consumption also exhibits excess sensitivity with respect to transitory shocks (Hall and Mishkin, 1982).² Models that feature complete markets and those that allow for just personal savings as a smoothing mechanism are clearly extreme characterizations of individual behavior and of the economic environment faced by the consumers. Deaton and Paxson (1994) notice this and envision “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 keeping with these remarks and empirical evidence, 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. We use the term *partial insurance* to denote the degree of transmission of income shocks to consumption.³ The paper makes three contributions to the existing literature. First, we address the issue of whether partial consumption insurance is available to agents and estimate the degree of partial insurance from the data, rather than imposing an *a priori* insurance configuration. Second, we estimate our model

²Hall and Mishkin (1982) use panel data on food consumption and income from the PSID and consider the covariance restrictions imposed by the PIH with quadratic utility. They impose the null of the PIH and do not study changes in inequality. See also Altonji, Martins and Siow (2003).

³Besides household saving and borrowing, there is scattered evidence on the role played by various partial insurance mechanisms 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, 2003), 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), 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).

using panel data on income and (imputed) non-durable consumption. The use of panel data allows us to relax a number of constraints which instead limit identification in repeated cross-sectional data. The use of non-durable consumption data avoids the ambiguities derived from basing the analysis on food consumption, which, besides being a necessity, represents a declining part of the household's budget. Finally, while we do not take a precise stand on the mechanisms (other than savings) that are available to smooth idiosyncratic shocks to income, we analyze empirically the mechanism behind the degree of insurance we find in the data, and in particular study the role of taxes and transfers, wealth, family labor supply, as well as heterogeneity by education and cohort of birth. Our aim is to provide 'structured facts' rather than a specific structural interpretation.⁴

Other papers have studied the joint evolution of the income and consumption distributions. 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. They use data on both income and consumption but lack a panel dimension. Our use of panel data on income *and* consumption allows us to identify the variance of the income shocks as well as the degree of insurance of consumption with respect to the two types of shocks. Krueger and Perri (2004) do not distinguish between transitory and permanent income shocks. As noted above, this is an important distinction, as we might expect to uncover less insurance for more persistent shocks. Moreover, this distinction plays an important role in separating changes in consumption inequality due to changing nature of income processes from changing availability of insurance. Krueger and Perri (2004) also propose a specific mechanism underlying the differences between consumption and income inequality (limited commitment), while we take a more agnostic approach. Finally, while they provide ample evidence on trends in consumption and income inequality, their exercise is primarily one of calibration (ours is one of estimation). Heathcote, Storesletten and Violante (2004) use the PSID to distinguish between less and more persistent shocks to male earnings. With this distinction, they show that a calibrated overlapping generations

⁴Our empirical approach is related to other papers in the literature, particularly Hall and Mishkin (1982), Altonji, Martins and Siow (2002), Deaton and Paxson (1994), and Moffitt and Gottschalk (1994). Hall and Mishkin (1982) use panel data on food consumption and income from the PSID and consider the covariance restrictions imposed by the PIH with quadratic utility. Altonji, Martins, and Siow (2002) improve on this by estimating a dynamic factor model of consumption, hours, wages, unemployment, and income, again using PSID data. Deaton and Paxson (1994) use repeated cross-section data from the US, the UK, and Taiwan to test the implications that the PIH imposes on consumption inequality. Moffitt and Gottschalk (1994) use PSID panel on income to identify the variance of permanent and transitory income shocks.

model with self-insurance and male labour supply is able to capture the broad pattern of consumption and wage inequality. These patterns are further examined in the recent study by Heathcote, Storesletten and Violante (2007) who, allowing for insurance beyond that in a simple bond economy, estimate a similar level of ‘partial insurance’ for persistent male earnings shocks as that recovered in our analysis. We derive the degree of insurance drawing a distinction between different measures of family income and earnings, using a new panel data series on consumption. Moreover, we offer an empirical evaluation of the mechanisms underlying the degree of insurance we find in the data. Nevertheless, our paper shares similar conclusions regarding the importance of insurance vs. durability of shocks.

The paper continues with a discussion of the underlying trends in income and consumption inequality and the development of the new panel data consumption series for the PSID. In Section 3 the consumption model is formulated and the identification strategy for recovering the insurance parameters and the inequality decomposition is discussed. Section 4 presents the empirical results concerning the evolution of volatility in permanent and transitory income shocks and estimates of the insurance parameters. The overall trends in inequality are similar to those found by Moffitt and Gottschalk (1995), Cutler and Katz (1992), Slesnick (2001) and Johnson, Smeeding, and Boyle Torrey (2005), among others.⁵ We disaggregate the data by different population groups to examine whether there are different changes in consumption inequality, and what mechanisms (institutions, labor market, credit market, etc.) are behind the estimated changes. Section 5 concludes.

2 Characteristics of Consumption and Income Inequality

While there are large panel data sets that track the distribution of wages and incomes for households over time, the same is not true for broad measures of consumption. The Panel Study of Income Dynamics (PSID) contains longitudinal income data, but the information on consumption is scanty (limited to food and few more items). Indeed, one of the reasons why consumption inequality has not been studied as extensively as income and wage inequality is the nature of data availability. In this section we first document some

⁵See Attanasio, Battistin and Ichimura (2004) and Primiceri and van Rens (2006) for other studies on consumption inequality in the US.

basic features of the evolution of consumption and income inequality that motivate our study. Repeated cross-section data such as the Consumer Expenditure Survey (CEX) are not enough to uncover the degree of persistence in income shocks or to identify the partial insurance model. For that we need panel data and in the second part of this section we describe our new panel data series.

2.1 The evolution of income and consumption inequality

There are two important features of the evolution of consumption and income inequality between the late 1970s and early 1990s which underpin our analysis. These are clearly evident from Figure 1, which uses PSID data on log income and CEX data on log consumption (see Section 2.2 below for details on sample selection and variable definitions). In this graph, we plot the actual estimates of the variances, as well as smoothing curves passing through the scatters (to ease legibility). In this figure the range of variation of the variance of PSID consumption is on the left-hand side; that of the variance of CEX consumption, on the right hand side. The first distinct feature is that the slope of the income variance (the solid line) is greater than the slope of the consumption variance (the dashed line). The second feature of these inequality figures is that consumption inequality flattens out completely in the second part of the 1980s, whereas income inequality continues to rise albeit at a much slower rate. Below we provide a framework for interpreting these changes. In particular, we show that the degree of detachment between consumption and income inequality depends on the persistence of income shocks and the availability of insurance to these shocks.

These overall patterns reflect what has also been found in previous analyses of inequality in income and consumption for this period, the most prominent study being that of Cutler and Katz (1992). See also the retrospective analysis in Johnson, Smeeding and Torrey (2005), and Dynarski and Gruber (1997). In the absence of panel data or a clear decomposition between low and high frequency shocks, none of these studies are able to relate the deviations in the two series to the durability of shocks (or the degree of insurance to shocks of different persistence), but the patterns they find do line up very closely with those in Figure 1. In particular Johnson, Smeeding and Torrey (2005) show the Gini for real equivalised disposable income rising from 0.34 to 0.40 in the period 1981 to 1985 and then up to 0.41 by 1992. The Gini for equivalised real

non-durable consumption rises from 0.25 to 0.28 over the first period and then hardly at all in the second period.⁶ Finally, Krueger and Perri (2006) document a rise in consumption inequality of a similar magnitude over this period with the variance of log consumption rising around 0.05 units over the 1980s. Their study uses data from the CEX exclusively and does not directly model the panel data dynamics of consumption and income jointly. In particular, they do not allow the degree of persistence in income shocks to vary over time.

In their ground-breaking study Deaton and Paxson (1994) present some detailed evidence on consumption inequality and interpret this within a life-cycle model. They note that consumption inequality should be monotonically increasing with age. Figure 2 shows this is broadly true for the cohorts in our sample. It also shows the large differences in initial conditions across birth cohorts with more recent cohorts experiencing a higher level of inequality at any given age. Initial conditions for different date of birth cohorts are extremely important to control for in understanding inequality.

Although Figure 1, and the discussion surrounding it, identifies two distinct episodes in the growth of income and consumption inequality, these overall trends do not help inform why these different episodes took place. Specifically they do not tell us anything about the nature of the changes in the income process or the nature of insurance that may have driven a wedge between consumption and income inequality. Studies that have investigated the impact of insurance either assume some external process for income or assume a specific form of insurance, typically the pure self-insurance model. Studies that have focussed on the durability of income shocks have exclusively focused on earnings among male workers and have not investigated the implications for consumption. For example, Moffitt and Gottshalk (1994, 2002) document a similar rise in male labor earnings inequality over the 1980s and attribute approximately half of this rise to changes in transitory earnings inequality. As we will see this is attributing rather more of the income inequality growth to transitory shocks than we find when combining family disposable income and consumption data. We explain the differences through labour supply reactions within the household.

⁶It is worth noting that the Gini and the variance of the log measures of inequality do not necessarily move in the same direction. Log normality is an exception. It is also useful to note in making these comparisons that a small transfer between two individuals a fixed income distance apart lower in the distribution will have a higher effect on the variance of logs than on the Gini coefficient.

2.2 A new panel consumption series

To further investigate the link between the evolution of income and consumption inequality and to estimate our partial insurance model we require panel data. The new panel data consumption series for PSID households that we develop here is derived by combining existing PSID data with data from the repeated cross-sections of the CEX. Previous studies have followed a similar approach. Skinner (1987), for example, imputes total consumption in the PSID using the estimated coefficients of a regression of total consumption on a series of consumption items (food, utilities, vehicles, etc.) that are present in both the PSID and the CEX. The regression is estimated with CEX data. Ziliak (1998) imputes consumption on the basis of income and the first difference of wealth (i.e., as the difference between income and savings). We depart from these studies by starting from a standard demand function for food (a consumption item available in both surveys). One novelty of our approach is to allow demands to change with relative prices, as well as non-durable expenditure and a host of demographic and socio-economic characteristics of the household. This demand function is estimated using CEX data. Food expenditure and total expenditure are modeled as jointly endogenous and, importantly, this relationship is allowed to change over time. Under monotonicity (normality) of food demand, this function can be inverted to obtain a measure of non durable consumption in the PSID. We find it attractive to work directly with the demand equation. However, as we allow for endogeneity and measurement error in both the total expenditure and the food expenditure variables, working directly with the inverse equation would also produce consistent estimates. 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).⁷

Before describing this procedure we first briefly describe the data and the sample selection. More details are provided in the Data Appendix A.1. For the main part of our analysis, we choose to select a PSID sample of continuously married couples headed by a male (with or without children) aged 30 to 65. We also

⁷After 1992 (or the 1993 survey year), PSID data are available in “early release” form and the interviews change from a pencil and paper telephone format to a computer-assisted telephone format, so we do not use them in the main part of our analysis. However, we do estimate the model using data up to 1996 as a sensitivity analysis, after which the panel became biennial.

eliminate households if the head or head’s spouse changes. Our sample selection therefore 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.⁸ To the extent that it is possible, we replicate this sample selection in the CEX. Finally, we should note that 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). For the most part we exclude SEO households and their split-offs. However, we do consider the robustness of our results in the low income SEO subsample.

We make use of two consumption measures: food and non-durables. In both data sets, food is the sum of annual expenditure on food at home and food away from home (in the PSID food data were not collected in 1987 and 1988).⁹ The definition of non durable consumption in the CEX is the same as in Attanasio and Weber (1995). It is the sum of food (defined above), alcohol, tobacco, and expenditure on other nondurable goods, such as services, heating fuel, public and private transports (including gasoline), personal care, and semidurables, defined as clothing and footwear. This definition 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.¹⁰

Table I compares the two data sets in terms of average demographic and socio-economic characteristics for selected years: 1980, 1983, 1986, 1989, and 1992. The 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

⁸Whether stable families have access to more or less insurance than non-stable families is an open question. 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 means-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.

⁹We are summing up expenditure on a luxury (food away from home) and on a necessity (food at home). Ideally, one could estimate a demand system and then work out a way to combine separate imputed values into one. We leave this to future work.

¹⁰We also experimented with a definition of nondurable consumption that includes services from some durables (housing and vehicles). We thank David Johnson at BLS for providing data on the latter.

similar, while the PSID tends to under-represent the proportion of people living in the West. Both male and female participation rates in the PSID are comparable to those in the CEX. 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. Total food expenditure (the sum of food at home and food away from home) is fairly similar in the two data sets.

To implement the imputation procedure, we pool all the CEX data from 1980 to 1992, and for any individual i in period t we write the following demand equation for food

$$f_{it} = W_{it}'\mu + p_t'\theta + \beta(D_{it})c_{it} + e_{it} \quad (1)$$

where f is the log of real food expenditure (which is available in both surveys), W and p contain a set of, respectively, demographic variables and relative prices (also available in both data sets), c is the log of 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 the elasticity $\beta(\cdot)$ (from now on, the budget elasticity) to vary with time and with observable household characteristics (D). The estimation results for our specification of (1) are reported in Table II. To account for measurement error of total expenditure we instrument the latter with the average (by cohort, year, and education) of the hourly wage of the husband and the average (also by cohort, year, and education) of the hourly wage of the wife. The budget elasticity is 0.85. The price elasticity is -0.98 . We test the overidentifying restrictions and fail to reject the null hypothesis (p-value of 28 percent). We also report statistics for judging the power of excluded instruments. They are all acceptable. Finally, we test whether the budget elasticity has remained constant over this period, and reject the hypothesis (p-value 1%). Generally the demographics have the expected sign. Armed with these estimates, we invert the demand function and derive a series of imputed non-durable consumption for all households in the PSID.

But how good is the imputation? In an annex to this paper, we review the conditions that make the

imputation procedure reliable.¹¹ Given that our preferred measure of inequality is the variance of the logs, we require that the evolution of the variance of the imputed log consumption series in the PSID mirrors that of the variance of the log consumption series in the CEX. A reliable imputation procedure requires that the variance of log consumption in the PSID differs from the CEX analog only by an additive factor (the variance of the error term of the demand equation scaled by the square of the budget elasticity); if this factor is constant over time, the *trends* in the two variances should be similar. Figure 3 shows that the variances line up extremely well. As in Figure 1, we eliminate the level effect, by rescaling the PSID consumption axis (on the left) to match that for CEX consumption (on the right). Trends in the variance of consumption are remarkably similar in the two data sets. In fact the reader can check that the variance of imputed PSID consumption is just an upward-translated version (by about 0.06 units) of the variance of CEX consumption. Both series suggest that between 1980 and 1986 consumption inequality grows quite substantially. Afterwards, both graphs are flat. In the annex, we show that this result is robust to variation in equivalence scales; we also show that our imputation procedure is capable of replicating quite well the trends in mean spending as long as account is made for differences in the mean of the input variable (food spending) in the two data sets.

3 Consumption Inequality, Insurance and the Durability of Income Shocks

To motivate the procedure for identifying the degree of transmission of income shocks to consumption, we propose a framework that focuses on the persistence of income shocks. We assume that the sole relevant source of idiosyncratic uncertainty faced by the consumer is net family income (defined as the sum of labor income and transfers, such as welfare payments, minus taxes paid). We also make the assumption of separability in preferences between consumption and leisure. This implies that all insurance provided through, say, an added worker effect, will pass through income. Similarly, insurance provided by taxes and transfers is accounted for in the net family income variable. In the discussion of the partial insurance results

¹¹The annex is available on the AER website.

we will, however, examine the importance of taxes and transfers, as well as married women’s labor market participation, as an insurance mechanism. Finally, 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.

3.1 The income process

Our aim here is to characterize changes in the persistence of shocks to income in a reasonably flexible but parsimonious way. For this we adopt a permanent-transitory model and allow the variances of the permanent and transitory factors to vary over time. In line with many previous empirical studies (MaCurdy, 1982; Abowd and Card, 1989; Moffitt and Gottschalk, 1994; Meghir and Pistaferri, 2004), we assume that the permanent component follows a random walk.¹²

Suppose real (log) income, $\log Y$, can be decomposed into a permanent component P and a mean-reverting transitory component v . The income process for each household i is:

$$\log Y_{it} = Z'_{it}\varphi_t + P_{it} + v_{it} \quad (2)$$

where t indexes time and Z is a set of income characteristics observable and known by consumers at time t . As we note below these will include demographic, education, ethnic and other variables. We allow the effect of such characteristics to shift with calendar time and we also allow for cohort effects.

We assume that the permanent component $P_{i,t}$ follows a martingale process of the form:

$$P_{i,t} = P_{i,t-1} + \zeta_{i,t} \quad (3)$$

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

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

¹²For example, Moffitt (1997) writes: “In the micro-level literature on earnings dynamics, Thomas MaCurdy, Abowd and Card, and Gottschalk and I all find evidence –also from the PSID– for a random walk in individual earnings in the United States” (p. 289, in *Comment to “Can families smooth variable earnings?”*, by S. Dynarski and J. Gruber, *Brookings Papers on Economic Activity*, 1, 229-303). Recent work on income dynamics, of which Guvenen (2006) is a leading example, has focused on models that allow less overall persistence and more general heterogeneous life-time income profiles. It would be a very useful exercise to extend the model of partial insurance we develop here to such alternative income processes. The key result of the *changing* persistence of income shocks and their impact on consumption inequality though seems unlikely to change.

with $\theta_0 \equiv 1$. It follows that (unexplained) income growth is

$$\Delta y_{i,t} = \zeta_{i,t} + \Delta v_{i,t} \quad (4)$$

where $y_{i,t} = \log Y_{i,t} - Z'_{i,t} \varphi_t$ denotes the log of real income net of predictable individual components.

3.2 The Transmission of Income Shocks to Consumption

We present a framework that allows us to study the degree of transmission of income shocks to consumption. We write (unexplained) change in log consumption as:

$$\Delta c_{i,t} = \phi_{i,t} \zeta_{i,t} + \psi_{i,t} \varepsilon_{i,t} + \xi_{i,t} \quad (5)$$

where $c_{i,t}$ is the log of real consumption net of its predictable components. We allow permanent income shocks $\zeta_{i,t}$ to impact consumption with a loading factor of $\phi_{i,t}$, which may potentially vary across individuals and time; the impact of transitory income shocks $\varepsilon_{i,t}$ is measured by the loading factor $\psi_{i,t}$. The random term $\xi_{i,t}$ represents innovations in consumption that are independent of those in income. This may capture measurement error in consumption, preference shocks, innovation to higher moments of the income process, etc. We call $\phi_{i,t}$ and $\psi_{i,t}$ *partial insurance* parameters.

Equation (5) nests the two extreme cases of full insurance of income shocks ($\phi_{i,t} = \psi_{i,t} = 0$) as contemplated by the complete markets hypothesis, and no insurance ($\phi_{i,t} = \psi_{i,t} = 1$) as in autarky, as well as intermediate cases in which $0 < \phi_{i,t} < 1$ and $0 < \psi_{i,t} < 1$. The closer the coefficient to zero, the higher the degree of insurance.

3.2.1 Self-Insurance

The most prominent intermediate case is the PIH with self-insurance through precautionary savings. The Appendix considers a version of the PIH with CRRA preferences, and shows that in this case approximation of the Euler equation for consumption gives $\phi_{i,t} \simeq \pi_{i,t}$ and $\psi_{i,t} \simeq \gamma_{t,L} \pi_{i,t}$, where $\pi_{i,t}$ is the share of future labor income in current human and financial wealth and $\gamma_{t,L}$ is an age-increasing annuitization factor.¹³ The

¹³See Appendix A.2. As far as we know, this is the first derivation of an analytical expression for the marginal propensity to consume with respect to permanent shocks in a model with CRRA preferences and transitory and permanent shocks. See Carroll (2001) for numerical simulations. Results from a simulation of a stochastic economy presented in Blundell, Low and Preston (2004) show that the approximation (17) can be used to accurately detect changes in the time series pattern of permanent and transitory variances to income shocks. These results are available on request (by email to: i.preston@ucl.ac.uk).

random term $\xi_{i,t}$ can be interpreted as the innovation to higher moments of the income process.¹⁴ Meghir and Pistaferri (2004) find evidence of this using PSID data.

The interpretation of the impact of income shocks on consumption growth in the PIH model with CRRA preferences is straightforward. 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, $\pi_{i,t} \simeq 1$, 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 self-insurance against permanent shocks only if the stock of assets built up is large relative to future labor income, which is to say $\pi_{i,t}$ is appreciably smaller than unity, in which case there will also be some smoothing of permanent shocks through self insurance. Carroll (2001) presents simulations that show for a buffer stock model the steady state value of $\pi_{i,t}$ is between 0.85 and 0.95. Blundell, Low and Preston (2007) simulate the model described in the Appendix using our estimates of the income process and find a value of $\pi_{i,t}$ of 0.8 or a little lower for individuals aged twenty years before retirement which corresponds to the average age in our sample. Finding that $\phi_{i,t} < \pi_{i,t}$ and/or $\psi_{i,t} < \gamma_{t,L}\pi_{i,t}$ represents evidence of partial insurance over and above self-insurance through savings.

3.2.2 Excess Smoothness and ‘Excess’ Insurance

A recent macroeconomic literature has explored a number of theoretical alternatives to the insurance configurations described above. These alternative models fall under two broad rubrics: those that assume public information but limited enforcement of contracts, and those that assume full commitment but private information. These models prove that the self-insurance case is Pareto-inefficient even conditioning on limited enforcement and private information issues. In both types of models, agents typically achieve more insurance than under a model with a single non-contingent bond, but less than under a complete markets environment. More importantly for our purposes, these models show that the relationship between income shocks and consumption depends on the degree of persistence of income shocks. Alvarez and Jermann (2000),

¹⁴This characterization follows Caballero (1990), who presents a model with stochastic higher moments of the income distribution. He shows that there are two types of innovation affecting consumption growth: innovation to the mean (the term $\pi_{i,t}\zeta_{i,t} + \pi_{i,t}\gamma_{t,L}\varepsilon_{i,t}$), and “a term that takes into account revisions in variance forecast” ($\xi_{i,t}$). Note that this term is not capturing precautionary savings *per se*, but the innovation to the consumption component that generates it (i.e., consumption growth due to precautionary savings will change to accommodate changes in the forecast of the amount of uncertainty one expects in the future).

for example, explores the nature of income insurance schemes in economies where agents cannot be prevented from withdrawing participation if the loss from the accumulated future 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. If income shocks are persistent enough and agents are infinitely lived, then participation constraints become so severe that no insurance scheme is feasible. With finite lived agents, the future benefits from a positive permanent shock exceed those from a comparable transitory shock. 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.

Another reason for partial insurance is moral hazard. This is the direction taken in Attanasio and Pavoni (2006). Here the economic environment is characterized by moral hazard and hidden asset accumulation, e.g., individuals have hidden access to a simple credit market. The authors show that, depending on the cost of shirking and the persistence of the income shock, some partial insurance is possible and a linear insurance rule can be obtained as an exact (closed form) solution in a dynamic Mirrlees model with CRRA utility. This provides a structural interpretation of the parameters in our estimated model. In particular, the response of consumption to permanent income shocks (what we call the partial insurance coefficient in our framework) could be interpreted as a measure of the severity of informational problems. Their empirical analysis finds evidence for “excess smoothness” of consumption with respect to permanent shocks.

3.2.3 Advance Information

In the analysis presented so far we have assumed that in the innovation process for income (4) the random variables $\zeta_{i,t}$ and $\varepsilon_{i,t}$ represent the arrival of new information to agent i in period t . If parts of these random terms were known in advance to the agent, then the intertemporal consumption model would argue that they should already be incorporated into current plans and would not directly effect consumption growth (5), see Cunha, Heckman and Navarro (2005). Suppose, for example, that only a proportion κ of the permanent shock was unknown to the consumer. Then the consumption growth relationship (5) would

become

$$\Delta c_{i,t} \simeq \tilde{\phi}_{i,t} \kappa \zeta_{i,t} + \psi_{i,t} \varepsilon_{i,t} + \xi_{i,t}. \quad (6)$$

where $\tilde{\phi}_{i,t}$ is the “true” insurance parameter. In this case $\tilde{\phi}_{i,t}$ would be underestimated by the information factor κ (i.e., we would call insurance what is, in fact and in part, advance information).¹⁵

The econometrician will treat $\zeta_{i,t}$ 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 $\phi_{i,t}$ has to be interpreted as reflecting a combination of insurance and information. In the absence of outside information (such as, say, subjective expectations), these two components cannot be separately identified. However, in our empirical analysis of the autocovariance structure of income and consumption we provide some evidence that advance information is not a serious problem during our sample period. In particular, we show that current consumption growth is not significantly correlated with future ‘shocks’ to income.

3.3 Evolution of Income and Consumption Variances

We assume that $\zeta_{i,t}$, $v_{i,t}$ and $\xi_{i,t}$ are mutually uncorrelated processes. As in Hall and Mishkin (1982) and others, one can impose covariance restrictions on the bivariate process (4) and (5) to identify the parameters of interest. In particular, equation (4) can be used to derive the following covariance restrictions in panel data

$$\text{cov}(\Delta y_t, \Delta y_{t+s}) = \begin{cases} \text{var}(\zeta_t) + \text{var}(\Delta v_t) & \text{for } s = 0 \\ \text{cov}(\Delta v_t, \Delta v_{t+s}) & \text{for } s \neq 0 \end{cases} \quad (7)$$

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 $\text{cov}(\Delta v_t, \Delta v_{t+s})$ depends on the serial correlation properties of v . If v is an $MA(q)$ serially correlated process, then $\text{cov}(\Delta v_t, \Delta v_{t+s})$ is zero whenever $|s| > q + 1$. Note also that if v is serially uncorrelated

¹⁵Another source of downward bias would result if the permanent component was less persistent than a martingale. As the π parameter reflects the annuity value of the shock, if the ζ shock was less persistent than implied by a unit root this would also lead to a value of ϕ less than unity.

($v_{i,t} = \varepsilon_{i,t}$), then $\text{var}(\Delta v_t) = \text{var}(\varepsilon_t) + \text{var}(\varepsilon_{t-1})$. Identification of the serial correlation coefficients does not hinge on the order of the process q . Allowing for an $MA(q)$ process, for example, adds $q - 1$ extra parameter (the $q - 1$ MA coefficients) but also $q - 1$ extra moments, so that identification is unaffected. Equation (7) shows that income inequality (obtained setting $s = 0$) may increase either because of increases in the variance of permanent shocks, or because of an increase in the variance of income growth due to transitory shocks.

The panel data restrictions on consumption growth from (5) are as follows¹⁶:

$$\text{cov}(\Delta c_t, \Delta c_{t+s}) = \phi_t^2 \text{var}(\zeta_t) + \psi_t^2 \text{var}(\varepsilon_t) + \text{var}(\xi_t) \quad (8)$$

for $s = 0$ and zero otherwise (due to the consumption martingale assumption). This equation shows that consumption growth inequality ($s = 0$) can rise for two reasons: a decline in the degree of insurance with respect to income shocks (for given variances), or an increase in the variances of income shocks (for given insurance). In other words (assuming $\xi_{i,t}$ is stationary), one can write the following decomposition for the time change in the variance of consumption growth:

$$\Delta \text{var}(\Delta c_t) = \text{var}(\zeta_t) \Delta \phi_t^2 + \phi_{t-1}^2 \Delta \text{var}(\zeta_t) + \text{var}(\varepsilon_t) \Delta \psi_t^2 + \psi_{t-1}^2 \Delta \text{var}(\varepsilon_t)$$

Our analysis below allows to separate the different forces at play visible in this equation. Finally, the covariance between income growth and consumption growth at various lags is:

$$\text{cov}(\Delta c_t, \Delta y_{t+s}) = \begin{cases} \phi_t \text{var}(\zeta_t) + \psi_t \text{var}(\varepsilon_t) \\ \psi_t \text{cov}(\varepsilon_t, \Delta v_{t+s}) \end{cases} \quad (9)$$

for $s = 0$, and $s > 0$ respectively. If v is an $MA(q)$ serially correlated process, then $\text{cov}(\Delta c_t, \Delta y_{t+s})$ is zero whenever $|s| > q + 1$. Thus, if v is serially uncorrelated ($v_{i,t} = \varepsilon_{i,t}$), then $\text{cov}(\Delta c_t, \Delta y_{t+s}) = -\psi_t \text{var}(\varepsilon_t)$ for $s = 1$ and 0 otherwise.

Note finally that it is likely that measurement error will contaminate the observed income and consumption data. Assume that both consumption and income are measured with multiplicative independent errors,

¹⁶The errors of approximation on these expressions are of the order of the expected values of the cubes of $|\zeta_t|$ and $|\varepsilon_t|$.

e.g.,

$$y_{i,t}^* = y_{i,t} + u_{i,t}^y \tag{10}$$

$$c_{i,t}^* = c_{i,t} + u_{i,t}^c \tag{11}$$

where x^* denote a measured variable, x its true, unobservable value, and u^x the measurement error.

In Appendix A.3 we discuss identification details of the model more in detail, and also show that the partial insurance parameter ϕ_t remains identified under measurement error, while only a lower bound for ψ_t is identifiable. A corollary of this is that the variance of measurement error in consumption can be identified (the theory suggests that consumption should be a martingale with drift, so any serial correlation in consumption growth can only be attributed to noise), but the variance of the measurement error in income can still not be identified separately from the variance of the transitory shock.¹⁷ The goal of the empirical analysis is to estimate features of the distribution of income shocks (variances of permanent and transitory shocks and the extent of serial correlation in the latter) and consumption growth (particularly the partial insurance parameters) using joint panel data on income and consumption growth on which the theoretical restrictions (7)-(9) have been imposed.

In the context of identifying sources of variation in household income and consumption, the availability of panel data presents several advantages over a repeated cross-sections analysis. With repeated cross-sections the variances and covariances of differences in income and consumption cannot be observed, though it is possible to make assumptions under which variances of shocks can be identified from differences in variances and covariances of their levels (assuming one knows the degree of insurance with respect to income shocks). For example, under the assumption that shocks are cross-sectionally orthogonal to past consumption and income, that transitory shocks are serially uncorrelated, and that $\phi_t = 1$ and $\psi_t = 0$, Blundell and Preston (1998) use repeated cross-section moments to separate the growth in the variance of transitory shocks to log income from the variance of permanent shocks (see also Deaton and Paxson, 1994). The assumed orthogonality assumption will be violated if aggregate consumption (or income) is not part of the

¹⁷Thus the variance of measurement error in consumption is identified by $-\text{cov}(\Delta c_t, \Delta c_{t+1})$.

consumer’s information set (see Deaton and Paxson, 1994). In panel data, identification 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 (see below). More crucially, with panel data one can estimate a richer model with the insurance parameter ϕ_t and ψ_t left free and thus test the validity of alternative explanations regarding the evolution of consumption inequality over time. In turn, knowledge of the extent of insurance is informative about the welfare effects of shifts in the income distribution. In our application we allow partial insurance parameters to differ by cohorts and interpret differences over time as year rather than age effects, although we appreciate that the choice is an arbitrary one made only for descriptive clarity.

Note finally that with panel data the 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:¹⁸

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

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

Using panel data on both consumption and income improves efficiency of these estimates because it provides extra moments for identification.

4 The Evidence

The parameters of interest in this study are the insurance parameters, ϕ and ψ , and the evolution of inequality in the permanent and transitory components to income. They are derived from the variance-covariance structure of changes in consumption and income. We consequently begin with the empirical characterisation of these autocovariances. We then evaluate the relative size and trends in the variance of permanent and transitory shocks to income and estimate the degree of insurance to these shocks for the whole sample and for different sub-groups of the population.

¹⁸See Meghir and Pistaferri (2004) for a generalization to serially correlated transitory shocks and measurement error in income. In their paper, they show that with an MA(q) process for the transitory shock, one needs $T = 4 + q$ years of data to identify the variances of interest. Given that we have access to a panel of 15 years, this condition is amply satisfied.

4.1 The Autocovariance of Consumption and Income

The impact of the deterministic effects Z_{it} on log income and (imputed) log consumption is removed by separate regressions 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 most of these characteristics to vary with calendar time. We then work with the residuals of these regressions, labelled $c_{i,t}$ and $y_{i,t}$.¹⁹

To pave the way to the formal analysis of partial insurance, Table III reports unrestricted minimum distance estimates of several moments of the income process for the whole sample: the variance of unexplained income growth, $\text{var}(\Delta y_t)$, the first-order autocovariances ($\text{cov}(\Delta y_{t+1}, \Delta y_t)$), and the second-order autocovariances ($\text{cov}(\Delta y_{t+2}, \Delta y_t)$). Estimates are reported for each year. Table IV repeats the exercise for our new panel data measure of consumption. Finally, Table V reports minimum distance estimates of contemporaneous and lagged consumption-income covariances. As noted above, some of the moments are missing because consumption data were not collected in the PSID in 1987-88.

Looking through Table III, one can notice the strong increase in the variance of income growth, rising by more than 30% by 1985. Also notice the 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 until the mid-1980s and then is stable or even declines. Second- and higher order autocovariances (which, from equation (7), are informative about the presence of serial correlation in the transitory income component) 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 in growth, as implied by an income process given by the sum of a martingale permanent component and a serially uncorrelated transitory component. Since evidence on second-order autocovariances is mixed, however, in estimation we

¹⁹To the extent that these regressions remove changes that are unexpected by the individuals, we might expect this to change the relative degree of persistence in the remaining shocks but not the insurance parameters. For example, by removing the effect of education-time on income and consumption, we are also removing the increase in inequality due to, say, changing education premiums (Attanasio and Davis, 1996). If we omit the education variables from our first stage, we find that it makes a little difference to the estimated insurance parameters (for example, the estimate of ϕ in Table VI below is 0.71 instead of 0.64). The same qualitative comment applies to the other variables whose effect is removed in the first stage.

allow for MA(1) serial correlation in the transitory component ($v_{i,t} = \varepsilon_{i,t} + \theta\varepsilon_{i,t-1}$).²⁰

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 IV and V. Table IV shows that the variance of imputed consumption growth also increases quite strongly in the early 1980s, peaks in 1985 and then it is essentially flat afterwards. Note the high value of the level 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.²¹ 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. Second-order and higher consumption growth autocovariances are mostly statistically insignificant and economically small.

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

From (9), the covariance between current consumption growth and one period ahead income growth $\text{cov}(\Delta c_t, \Delta y_{t+1})$ should reflect the extent of insurance with respect to transitory shocks (i.e., $\text{cov}(\Delta c_t, \Delta y_{t+1}) = 0$ if there is full insurance of transitory shocks). We note that in the pure self-insurance case with infinite horizon and MA(1) transitory component, the impact of transitory shocks on consumption growth is given by the annuity value $\frac{r(1+r-\theta)}{(1+r)^2}$. With a small interest rate, this will be indistinguishable from zero, at least statistically. In fact, this covariance is hardly statistically significant and economically close to zero. At the foot of the Table we present the p-values for the joint significance tests of the autocovariances $E(\Delta c_t, \Delta y_{t+j})$ ($j \geq 1$). These p-value also detect advance information. If future income shocks were known to the consumer

²⁰We also estimated the autocovariances of income growth at lags greater than 2 and find that none of them is statistically significant. These results are available from the authors on request.

²¹To 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 $\text{var}(\Delta c_t)$.

in earlier periods then consumption should adjust before the observed shock occurs. This should show up in significant autocovariances between changes in consumption and future incomes. However, we find no statistical evidence that this may be the case.

The covariance between current consumption growth and past income growth $\text{cov}(\Delta c_{t+1}, \Delta y_t)$ plays no role in the PIH model with perfect capital markets, but may be important in alternative models where liquidity constraints are present (a standard excess sensitivity argument, see Flavin, 1981). The estimates of this covariance in Table V are also close to zero.

To sum up, the evidence suggests that a simple permanent-transitory framework for income shocks with time varying second order moments in these shocks provides a good representation of the income process for families in the PSID over this period. Overall we find only weak evidence that transitory shocks impact consumption growth. However, in the sensitivity results reported below we find that there is evidence of significant responsiveness to transitory shocks for low wealth families and for the low income poverty sample of the PSID.

4.2 Insurance

Our focus here will be on the variances of the permanent and the transitory shock, σ_ζ^2 and σ_ε^2 , on the partial insurance coefficients for the permanent shock (ϕ) and for the transitory shock (ψ), and the way these parameters vary over time as well as among different groups in the population. Our estimates are based on a generalization of moments (7)-(9). In particular, to account for our imputation procedure, we allow consumption to be measured with error. We allow the variance of the measurement error in consumption to vary with time. This is to capture the fact that the imputation error is scaled by a time-varying budget elasticity which induces non-stationarity. We also consider an MA(1) process for the transitory error component of income ($v_{i,a,t} = \varepsilon_{i,t} + \theta\varepsilon_{i,t-1}$), and estimate the MA(1) parameter θ . Finally, we allow for i.i.d. unobserved heterogeneity in the individual consumption gradient, and estimate its variance (σ_ζ^2).

We present the results of three specifications: one for the whole sample (the “baseline” specification), one where the parameters are estimated separately by education (college vs. no college), and one where

parameters are estimated separately by cohort (born 1930s vs. born 1940s).²² We also allow for some time non-stationarity. In particular, in all specifications we let the variances of the permanent and the transitory shock, σ_ζ^2 and σ_ε^2 , respectively, vary with calendar time. As for the partial insurance coefficients for the permanent shock (ϕ) and for the transitory shock (ψ), we assume that they take on two different values, before and after 1985. This is consistent with the the evidence in Figure 1, which divides the sample period in a period of rapid growth in the variance (up until 1985), and one of relative stability afterwards. We test the null that the extent of insurance does not change over time, and with almost no exceptions we fail to reject the null. In the discussion of the results that follows we comment on the time variability of the insurance parameters where appropriate and present the results of the test in the tables.

The parameters are estimated by diagonally weighted minimum distance (DWMD). This estimation method is a simple generalization of equally minimum distance (EWMD). Unlike EWMD, it allows for heteroskedasticity. Moreover, it avoids the pitfalls of optimal minimum distance (OMD) remarked by Altonji and Segal (1996), which are primarily related to the terms outside the main diagonal of the optimal weighting matrix. Technical details are in Appendix A.4.²³

The first column of Table VI shows the results for the whole sample. We defer the discussion of the estimated variances of the permanent shock and the estimated variances of the transitory shock to the next paragraph. The MA parameter for the transitory shock is small. The estimates of the variance of the imputation error (not reported) are always precisely measured and suggest that the imputation error absorbs a large amount of the cross-sectional variability in consumption (the estimates vary between 0.05 and 0.10). The variance of unobserved heterogeneity in the consumption gradient is small but significant. In the whole sample the estimate of ϕ , the partial insurance coefficient for the permanent shock, provides evidence in favor of some partial insurance.²⁴ In particular, a 10 percent permanent income shock induces a

²²Results for the younger cohort (born in the 1950s) and the older cohort (born in the 1920s) are less reliable because these cohorts are not observed for the whole sample period. We thus omit them.

²³If we use EWMD we obtain extremely downward biased estimates of $\text{var}(\zeta_t)$ and extremely upward biased estimates of $\text{var}(\varepsilon_t)$ (compared to those we obtain using just income data, as in (12) and (13)). With DWMD the two sets of estimates are similar because we are effectively putting more “identification weight” for the income shock variances on the income moments and less on the consumption moments (which display more sampling variability due to the imputation procedure).

²⁴As shown in the Appendix, if income is measured with error, the estimate of $\sigma_\xi^2(\psi)$ is upward (downward) biased. However, the bias is likely negligible (see the Appendix for an example).

6.4 percent permanent change in consumption.²⁵ The evidence on ψ accords with a simple PIH model with a long horizon.²⁶ If we allow the partial insurance parameters to vary across time then we find a slightly lower estimate of ϕ - indicating more insurance - in the later part of the 1980s. This would be in line with the idea developed in Krueger and Perri (2006) that a higher variance provides additional incentives to insure. However, the differences in the partial insurance parameters over this time period are small and are not statistically significant. Hence we decided to restrict the coefficient to be constant over the whole period. The p-values for the test of constant insurance parameters over the two sub-periods are given in the last two rows of the table.²⁷

There is much discussion in the literature on the reasons for the increase in income inequality over the 1980s. In particular, there is much debate on whether the rise can be labeled permanent or transitory. In Figure 4 we plot the minimum distance estimate of the variance of the permanent shock, $\text{var}(\zeta_t)$, against time. There are two sets of estimates. One uses the full set of consumption and income moments for the baseline specification in Table VI, and another just utilizes the income data. There is a close accordance between the two series which provides a check on the validity of our specification. The figure points to a strong growth in permanent income shocks during the early 1980s. The variance of permanent shocks levels off thereafter. It is also worth noting that from trough to peak the variance of the permanent shock more than doubles.²⁸ This evidence on permanent shocks is similar to that reported by Moffitt and Gottschalk (1994) using PSID data on male earnings. However, as we will document below, the precise evolution of inequality in transitory shocks depends on the source of income under study. Male labor earnings data will be shown to display a higher transitory variance in the earlier part of this time period.

Table VI also reports the results of the model for two education groups (with and without college education), and for two representative birth cohorts (born in the 1940s and born in the 1930s).²⁹ The

²⁵This “excess smoothness” result has been replicated in recent papers by Attanasio and Pavoni (2006), Primiceri and Van Rens (2006), and Heathcote, Storesletten and Violante (2007).

²⁶If 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.

²⁷We note that the overall results are maintained by extending the data forward until 1996. These results are available from the authors on request.

²⁸An even more strikingly accordance between the two alternative estimates is found for the estimated variances of the transitory shock, which we omit here.

²⁹Since we stratify the sample by exogenous characteristics and estimate different parameters for different groups, we are

partial insurance parameter estimates point to interesting differences in insurance by type of household. In particular there appears to be less insurance in response to permanent shocks among the group with no college education (indeed, we would not statistically reject the null hypothesis that there is *no* insurance in this group). In contrast, the evidence on ψ accords with a simple PIH model and we cannot reject the null that there is full smoothing with respect to transitory shocks ($\psi = 0$) for both education groups, though for the less well educated the point estimate is higher.

When the sample is stratified by year of birth, we find qualitatively similar results: there is evidence for full insurance with respect to transitory shocks and differences in the extent of insurance with respect to the permanent shocks.³⁰ It is worth considering whether the presence of precautionary asset accumulation is an explanation for the pattern of results. Recall that the insurance coefficients may reflect differences in $\pi_{i,t}$ (the share of future labor income in the present value of lifetime wealth), which in our framework reflects how close an individual is to retirement age. Thus, $\pi_{i,t}$ is likely to be lower for older cohort because older cohorts have both more accumulated financial wealth and lower prospective human capital wealth. Indeed, we find some evidence that permanent shocks for the older cohort are smoothed to a greater extent than for younger cohorts, although these sub-group estimates are less precise. Whether this is due to the effect played by precautionary wealth accumulation remarked above or by greater availability of insurance (such as social security, disability insurance or even insurance provided by adult children) in the group of people born in the 1930s is something that we cannot address in the absence of additional information, such as panel data on assets and age-specific estimates of human capital wealth. Later we provide some suggestive evidence that wealth accumulation is a potentially important explanation for the degree of insurance with respect to permanent income shocks.³¹

How good is the fit of our model? In Figure 5 we plot the actual variance of income growth and its

effectively considering the insurability of shocks within groups.

³⁰We find qualitatively similar results if we relax the age requirement (including those aged between 25 and 30). The estimate of ϕ is 0.70 (s.e. 0.10), indicating slightly less insurance to permanent shocks. This can be interpreted as reflecting a longer horizon among younger aged individuals. The estimate of ψ is 0.06 (s.e. 0.04).

³¹In a separate experiment (not reported for brevity), we exploited variability across cohorts and allowed the insurance parameters ϕ and ψ to depend on age. We fit a linear age trend by minimum distance: $\phi_a = \phi_0 + \phi_1 age + e$, where e is an error. We find evidence of a decline in the value of ϕ by age (consistent with precautionary saving), but the estimates are not very precise. We also tried a quadratic age trend, but the fit worsened. A difference statistic would favor the linear trend specification.

predicted value from our baseline model. We repeat the exercise for the variance of consumption growth and the covariance between income and consumption growth. Our model appears to fit the model quite well in all three dimensions.

Before delving in to more detail concerning the underlying mechanisms at work in our results we ask the question: Could these baseline results have been obtained using food data alone? With almost no exception, all the papers in the literature (including Hall and Mishkin, 1982, Hayashi, Altonji, and Kotlikoff, 1996) use the PSID data on food, so it is worth asking what is the value added of using our imputed measure of consumption. A possible argument in favor of this simpler approach is that food is a constant fraction of nondurable expenditure, so that the degree of insurance of food with respect to income shocks (transitory and permanent) reflects partly the true degree of insurance of nondurable consumption (i.e., ϕ and ψ) and partly the relationship between food and nondurable consumption (the budget elasticity). If the latter is known (for example, from demand studies), the former can be backed out easily. The pitfall here is that the assumption of a constant budget elasticity (β in (1)) is rejected (see Table II). We re-estimated the model using food rather than our imputed measure of consumption. The results, not reported for brevity, show that using food would provide evidence that: (a) there is more insurance than with imputed consumption data, and (b) there is evidence that insurance is increasing over time, while with imputed consumption there is none (the value of ϕ falls from 0.57 to 0.29 and the p-value of the test of constant insurance is 1.6 percent). It is straightforward to prove that the insurance parameters we are identifying here is $\phi_t = \phi\beta_t$. Since β_t declines over time, there is evidence of increasing insurance. Thus, what is really a changing budget elasticity is interpreted as changing insurance (for which we do not find statistically significant evidence when using a measure of non-durable consumption). Of course, things would be even worse if insurance *was* also changing. A study using food data would be unable to separate changing insurance of income shocks from changing elasticity of food consumption. The conclusion is that using food may give misleading evidence on the size and the stability of the insurance parameters.

4.3 Taxes and Transfers and Labor Supply

To examine the role of alternative insurance mechanisms, Table VII presents an analysis that replaces family net income with two alternative income measures: total family earnings and male earnings. Here we focus exclusively on the two insurance parameters ϕ and ψ . The reduction in the permanent insurance coefficient ϕ in the second column (a 50% reduction) indicates the important role of taxes and transfers in providing insurance to permanent shocks. This happens because consumption still incorporates any insurance value of taxes and transfers but the new measure of income no longer does. This insurance is also reflected through changes in the estimated variance of permanent and transitory shocks.³² With taxes and transfers excluded, the variances of income shocks are indeed much higher. There is also a further decline in the estimated ϕ coefficient when we only consider male earnings.³³ This is indication that family labour supply may also have played an important insurance role during this period.

It is interesting to note at this point the different pattern of transitory income inequality recovered from the baseline model versus the male earnings only specification. This is presented in Figure 6 which plots the path of the two variances over this period. Once total net income is considered, rather than male labor earnings alone, there is a much shallower rise in transitory income uncertainty. This reconciles the differences with the results from the male earnings literature, in particular Moffitt and Gottschalk (1994) who, using male earnings in the PSID, document a much steeper rise in transitory inequality earlier in the 1980s. As noted above their pattern of permanent inequality is closely in accord with Figure 4. The most interesting aspect of Figure 6 is that in the early 1980s there is little or no growth in the variance of the transitory shock to net income. Most of the growth occurs in the second half of the sample. This is in sharp contrast with the trends in the variance of the permanent shock to net income, which rises in the early 1980s and flattens out afterwards. Thus we may conclude that the increase in income inequality of the early 1980s is of permanent nature, while the growth in the second half of the sample is more temporary in nature.

³²The results for the variance estimates are not reported, but available on request.

³³Heathcote, Storesletten and Violante (2007) estimate a similar response of consumption to permanent shocks to male earnings. As they note, endogenous male labor supply drives a further wedge between the transmission from earnings to that from wages. Permanent shocks to earnings passing through much less than wages due to the insurance value of labour supply.

4.4 A variance of consumption growth decomposition

At this point, we can go back to the decomposition of the variance of consumption growth proposed in Section 3.3:

$$\Delta \text{var}(\Delta c_t) \simeq \text{var}(\zeta_t) \Delta \phi_t^2 + \phi_{t-1}^2 \Delta \text{var}(\zeta_t) + \text{var}(\varepsilon_t) \Delta \psi_t^2 + \psi_{t-1}^2 \Delta \text{var}(\varepsilon_t)$$

and propose an explanation for our findings. We have argued that there is no evidence that insurance has changed over the sample period we examine. Thus $\Delta \phi_t^2 = \Delta \psi_t^2 = 0$. In the first half of our sample period there is a strong growth in the variance of permanent income shocks and little growth in the variance of transitory shocks, implying $\Delta \text{var}(\Delta c_t) \simeq \phi^2 \Delta \text{var}(\zeta_t)$. If there were no insurance with respect to permanent income shocks $\Delta \text{var}(\Delta c_t) = \Delta \text{var}(\zeta_t)$, but in fact we find empirically that $\phi < 1$, and so there is some attenuation, although as we saw earlier consumption inequality rises substantially. In the second half of the sample, the variance of permanent income shocks is stable while the variance of transitory shocks grows. This implies $\Delta \text{var}(\Delta c_t) \simeq \psi^2 \Delta \text{var}(\varepsilon_t)$. Since we find that $\psi \simeq 0$, there is little overall growth of consumption inequality in this period. This provides a simple explanation for the trends reported in Figures 1-3, as well as those in Table III and IV and Figure 5.

These results show that the change in the degree of persistence in income shocks is a key characteristic of the income distribution in the US over this period and an important link in the relationship with consumption inequality. Suppose that one ignores this change in persistence and simply specifies a single transmission parameter linking income shocks to consumption growth, as in Krueger and Perri (2006), for example. It is straightforward to show that with the weight of income variance shifting progressively towards more transitory shocks, one would have the impression that the degree of insurance is increasing over time, even though ϕ and ψ are both constant. The reason is that the single insurance coefficient ends up being a weighted average of ϕ and ψ , with weights given by the relative importance of permanent and transitory shocks in the overall income growth variance. If the weight on ψ rises, the fact that transitory shocks are easier to insure will provide misleading evidence regarding insurance. The disjunction between consumption inequality that we have documented occurs not because it has become easier to insure consumption against

income shocks, but because the rise of income inequality over part of this period is of temporary nature, and temporary shocks are generally easier to insure than permanent shocks.

One important question is what may have caused the shift in the persistence of the income process, i.e., a rise in what has been termed “income instability”. Gottschalk and Moffitt (1994) conclude that part of the rise in instability they observe in longitudinal PSID data is due to compositional effects, i.e., employment shifts from sector with less variable earnings (manufacturing) to sectors with more (services), or from unionized to non-unionized jobs. Another part is due to greater mobility between jobs and the increase in self-employment and part-time or temporary work. However, the bulk of the increase in transitory variance appears to have been idiosyncratic.

4.5 Private transfers, Low wealth, and Total Expenditure

Next we focus attention on help from relatives (private transfers) and on the degree of insurance among low income families. The impact of measured help from friends and relatives is negligible, as the first two columns in Tables VIII show. This result is reminiscent of Hayashi, Altonji and Kotlikoff (1996), who find little evidence of insurance within the family.

Examining groups stratified by wealth provides more interesting deviations from the baseline specification. In the third column of Table VIII we consider low wealth households. We define as “low wealth” households those whose wealth, in the first year they are observed, is in the bottom 20% of the distribution of initial wealth. Wealth is given by $\left(\frac{\text{asset income}_{i,t}}{r_t} + \text{housing}_{i,t}\right)$ where t corresponds to the first year when household i is observed in the sample. We assume r_t is equal to the T-bill return for that year. Given that the level of wealth in the initial period is pre-determined (with respect to consumption growth decisions taken thereafter), the corresponding sample stratification we adopt does not suffer from endogeneity problems.³⁴ We now find that there is a significant impact of transitory shocks on consumption. Not surprisingly this group has less ability to self insure even transitory income fluctuations. This estimate is not far from the 0.2 benchmark found by other researchers, such as Hall and Mishkin (1982), who impute this excess sensitivity

³⁴Using the actual wealth data available in the PSID in 1984 and 1989 is a possible alternative. However, given that we want to stratify the sample on the basis of initial wealth, we would end up with much reduced sample sizes.

of consumption to transitory income shocks to binding liquidity constraints. We also find that there is no statistical evidence of insurance with respect to permanent shocks. In contrast, insurance to permanent shocks is much more important for the higher wealth group, again in accord with the modelling framework outlined above. Accumulated wealth can in fact be run down to smooth consumption against persistent income shocks.

For low wealth households with limited access to credit markets, is it possible that durable purchase and the timing of durable replacement might act as some form of insurance to transitory shocks. This argument is developed in Browning and Crossley (2003), who show that with small costs of accessing the credit market (or small transaction costs in the second-hand market for durables), the replacement of not fully collateralized durables could be used to smooth non-durable consumption in the face of short-run income shocks. This would imply that with a measure of consumption that includes durables we should find less evidence for insurance, i.e., the estimated ψ would rise. The penultimate column of Table VIII, which uses a consumption measure including durable purchases and focuses on a low wealth sample likely to face credit restrictions, provides some confirmation of that. It suggests that durables are particularly useful as a smoothing mechanism in response to transitory shocks for low wealth individuals.³⁵

Finally, in the last column of Table VIII, we extend our sample to the families of the SEO (the low-income subsample in the PSID). In comparison with the baseline we would again reject full insurance with respect to transitory shocks. This confirms the finding that in low-income or low-wealth samples, the evidence for insurance against transitory shock is basically absent. Interestingly, the overall pattern of permanent income inequality is similar across various specifications and samples (with the exception of education - the growth in the variance of permanent shocks does appear to have continued in to the late 1980s for those with college education), as displayed in Figure 7. One possible interpretation of this is that the differences in the estimates of ϕ that we find reflect genuine economic differences in access to insurance rather than differences

³⁵See Meyer and Sullivan (2001) for a detailed discussion of the measurement of durables in the CEX. Our measure of total consumption includes food, alcohol, tobacco, services, heating fuel, public and private transports (including gasoline), personal care, semidurables (clothing and footwear), and expenditure on durables, namely housing (mortgage interests, property tax, rents, other lodging, textiles, furniture, floor coverings, appliances), new and used cars, vehicle finance charges and insurance, car rentals and leases, health (insurance, prescription drugs, medical services), education, cash contributions, and personal insurance (life insurance and retirement).

in the variance of permanent shocks.

5 Conclusions

The aim of this paper has been to evaluate the link between consumption and income inequality through the degree of consumption insurance with respect to income shocks, both temporary and permanent. This was achieved by investigating the extent to which the distribution of income shocks is transmitted to the distribution of consumption. For this we created a new panel consumption series for the PSID using an imputation procedure that maps food data into consumption data using the estimates of a demand equation for food, estimated from repeated CEX cross-sections. We document a disjuncture between income and consumption inequality that occurred in the middle of the 1980s in the US. We argue that this disjuncture can be explained by the change in the persistence of income shocks over this period. In particular, an initial growth in permanent shocks which was then replaced by growth in transitory income shocks.

The analysis uncovers a strong growth in permanent income shocks in the US during the early 1980s (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 only mildly increasing in the period where the variance of permanent shock is increasing, and it increases only when the variance of permanent shock slows down. Although we find important differences in the degree of insurance to these shocks by wealth, education and birth cohort, the interpretation of the relationship between consumption and income inequality is maintained.

The economic framework in this paper allowed for self-insurance, in which consumers smooth idiosyncratic shocks through saving, and complete markets in which all idiosyncratic shocks are insured. Neither of these models were found to accord with the evidence. Instead we find some partial insurance for permanent shocks and almost complete insurance of transitory shocks. Only for low wealth households do we find significant sensitivity, and therefore only partial insurance, with respect to transitory income shocks. Interestingly, there appears to be a much greater degree of insurance of permanent shocks among the college educated.

Not surprisingly we find also more insurance of such shocks for older cohorts. Our model suggests that we should see more insurance, even for permanent shocks, among those nearing retirement, especially where they have built up sufficient precautionary savings. The tax and welfare system are also found to play an important insurance role for permanent shocks. When we include durables in our measure of consumption we find much less evidence of insurance of transitory shocks suggesting that durables may be acting as alternative smoothing mechanism for low wealth families.

Recent work on income and consumption dynamics, building on the earlier studies of earnings dynamics by Lillard and Weiss (1979) and by Baker (1997), focuses on models that allow general heterogeneous life-time income profiles (Guisarri, 2006). These studies also find lower overall persistence. As we have noted the unit root assumption follows findings from many papers in the literature on labor earnings, but alternative processes with less persistence and individual trends are increasingly common. The introduction of heterogeneous income trends is an important development and it would be a very useful exercise to extend the model of partial insurance we develop here to such alternative income processes.³⁶ The main point in this study, however, is that it is the change in the degree of persistence of income shocks in the 1980s, rather than the level itself, that explains the observed disjuncture between the evolution of income and consumption inequality.

These 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

³⁶Baker and Solon (2003) use a large Canadian tax administrative data set to show that the random walk component remains of key importance even in the heterogeneous trends specification.

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] Altonji, J., A.P. Martins, and A. Siow (2002), “Dynamic factor models of consumption, hours, and income”, *Research in Economics*, 56, 3-59.
- [3] Altonji, J., and L. Segal (1996), “Small-sample bias in GMM estimation of covariance structures”, *Journal of Business and Economic Statistics*, 14, 353-66.
- [4] 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.
- [5] Alvarez, F. and U. Jermann (2000), “Efficiency, equilibrium and asset pricing with risk of default”, *Econometrica*, 68, 775-97.
- [6] Attanasio, O., E. Battistin and H. Ichimura (2004), “What really happened to consumption inequality in the US?”, University College London, mimeo.
- [7] Attanasio, O., and S. Davis (1996), “Relative wage movements and the distribution of consumption”, *Journal of Political Economy*, 104, 1227-62.
- [8] Attanasio, O. and N. Pavoni (2006), “Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption”, mimeo University College London, May.
- [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] Baker, M. (1997), "Growth-rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings", *Journal of Labor Economics*, 15, 338-375.
- [13] Baker, M. and G. Solon (2003), "Earnings Dynamics and Inequality among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records," *Journal of Labor Economics*, vol. 21(2), pages 267-288.
- [14] Blundell, R., H. Low and I. Preston (2007), "Income Risk and Consumption Inequality: A Simulation Study", Institute for Fiscal Studies WP04/26, revised November 2007.
- [15] Blundell, R., and L. Pistaferri (2003), "Income volatility and household consumption: The impact of food assistance programs", *Journal of Human Resources*, Vol 38, 1032-1050.
- [16] Blundell, R., L. Pistaferri and I. Preston (2004), "Imputing consumption in the PSID using food demand estimates from the CEX", Institute for Fiscal Studies WP04/27 (<http://www.ifs.org.uk/workingpapers/wp0427.pdf>).
- [17] Blundell, R., and I. Preston (1998), "Consumption inequality and income uncertainty", *Quarterly Journal of Economics* 113, 603-640.
- [18] Browning, M., and T. Crossley (2003), "Shocks, stocks and socks: Consumption smoothing and the replacement of durables", McMaster Economics Working Paper 2003-7.
- [19] Caballero, R. (1990), "Consumption Puzzles and Precautionary Savings", *Journal of Monetary Economics*, 25(1), 113-136.
- [20] Campbell, J.Y. and A.S. Deaton (1989), "Why is consumption so smooth?", *Review of Economic Studies* 56, 357-373.
- [21] Carroll, C. (2001), "Precautionary saving and the marginal propensity to consume out of permanent income", NBER Working Paper 8233.

- [22] Chamberlain, G. (1984), "Panel data", in Handbook of Econometrics, vol. 2, edited by Zvi Griliches and Michael D. Intriligator. Amsterdam: North-Holland.
- [23] Cunha, F., J. Heckman and S. Navarro (2005), "Separating uncertainty from heterogeneity in life cycle earnings", Oxford Economic Papers 57, 191-261.
- [24] 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.
- [25] Deaton, A. (1992), Understanding Consumption. Baltimore: John Hopkins University Press.
- [26] Deaton, A. (1997), The Analysis of Household Surveys. Baltimore: John Hopkins University Press.
- [27] Deaton, A., and C. Paxson (1994), "Intertemporal choice and inequality", Journal of Political Economy, 102, 384-94.
- [28] Dynarski, S., and J. Gruber (1997), "Can families smooth variable earnings?", Brooking Papers on Economic Activity, 1, 229-305.
- [29] Engen, E., and J. Gruber (2001), "Unemployment insurance and precautionary savings", Journal of Monetary Economics, 47, 545-79.
- [30] Fay, S., E. Hurst, and M. White (2002), "The consumer bankruptcy decision", American Economic Review, 92, 706-18.
- [31] Flavin, M. A. (1981), "The Adjustment of Consumption to Changing Expectations about Future Income," Journal of Political Economy, vol. 89, no.5.
- [32] Gruber, J. (2000), "Cash welfare as a consumption smoothing mechanism for single mothers", Journal of Public Economics, 75, 157-82.
- [33] Gruber, J., and A. Yelowitz (1999), "Public Health Insurance and Private Savings," Journal of Political Economy, 107(6), 1249-74.

- [34] Guiso, L., L. Pistaferri and F. Schivardi (2005), “Insurance within the firm”, *Journal of Political Economy* 113(5), 1054-87.
- [35] Guvenen, F. (2007), “Learning your Earning: Are Labor Income Shocks Really Very Persistent?”, *American Economic Review*, vol. 97(3), pages 687-712, June.
- [36] Hall, R., and F. Mishkin (1982), “The sensitivity of consumption to transitory income: Estimates from panel data of households”, *Econometrica*, 50, 261-81.
- [37] Hayashi, F., J. Altonji, and L. Kotlikoff (1996), “Risk sharing between and within families”, *Econometrica*, 64, 261-94.
- [38] Heathcote, J., K. Storesletten and G. Violante (2004), “The Macroeconomic Implications of Rising Wage Inequality in the United States”, Working Paper, New York University, August.
- [39] Heathcote, J., K. Storesletten and G. Violante (2007), “Consumption and Labour Supply with Partial Insurance: An Analytical Framework”, CEPR Discussion Paper, May.
- [40] Hill, M. (1992), *The Panel Study of Income Dynamics: A user’s guide*, Newbury Park, California: Sage Publications.
- [41] Johnson, D. and T.M. Smeeding (1998), “Measuring the Trends in Inequality of Individuals and Families: Income and Consumption”, BLS manuscript.
- [42] Johnson, D.S., T.M. Smeeding and B.R.Torrey (2005), “United States Inequality through the Prisms of Income and Consumption”, *Monthly Labor Review*, 11-24, April.
- [43] Kniesner, T., and J. Ziliak (2002), “Tax reform and automatic stabilization”, *American Economic Review*, 92, 590-612.
- [44] Kotlikoff, L., and A. Spivak (1981), “The family as an incomplete annuities market”, *Journal of Political Economy*, 89, 372-91.

- [45] Krueger, D., and F. Perri (2006), “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory”, *Review of Economic Studies*, 73, 163-193.
- [46] Lillard, L.A., and Weiss, Y.A. (1979), “Components of Variation in Panel Earnings Data: American Scientists, 1960-70”, *Econometrica*, 47, 437-454.
- [47] 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.
- [48] Mankiw, M., and M. Kimball (1992), “Precautionary saving and the timing of taxes”, *Journal of Political Economy*, 97, 863-79.
- [49] Meghir, C., and L. Pistaferri (2004), “Income variance dynamics and heterogeneity”, *Econometrica*, 72(1), 1-32.
- [50] Meyer, B. and J. Sullivan (2001), “The living conditions of single mothers in the 19080s and 1990s”, NBER Working Paper 8298, May.
- [51] Moffitt, R., and P. Gottschalk (1994), “Trends in the autocovariance structure of earnings in the US: 1969-1987”, Brown University, mimeo.
- [52] Moffitt, R., and P. Gottschalk (2002), “Trends in the Transitory Variance of Earnings in the United States”, *Economic Journal*, vol. 112(127), pages C68-C73, March.
- [53] Primiceri, G. and T. van Rens (2006), “Predictable Life-Cycle Shocks, Income Risk and Consumption Inequality”, mimeo, CREI UPF, February.
- [54] Skinner, J. (1987), “A superior measure of consumption from the Panel Study of Income Dynamics”, *Economic Letters*, 23, 213-16.
- [55] Slesnick, D.T. (2001), “Consumption and Social Welfare”, Cambridge University Press, Cambridge, April.

- [56] Stephens, M. Jr. (2001), "The Long-Run Consumption Effects of Earnings Shocks", *The Review of Economics and Statistics*, vol.83, n.1, p.28-36.
- [57] Stephens, M. Jr. (2002), "Worker displacement and the added worker effect", *Journal of Labor Economics*, 20, 504-37.
- [58] Ziliak, James P. (1998), "Does the Choice of Consumption Measure Matter? An Application to the Permanent-Income Hypothesis", *Journal of Monetary Economics*, 41(1), 201-16.

Table I
Comparison of means, PSID and CEX

	1980		1983		1986		1989		1992	
	PSID	CEX	PSID	CEX	PSID	CEX	PSID	CEX	PSID	CEX
Age	42.94	43.71	43.43	45.01	43.86	46.03	44.03	45.26	45.95	46.88
Family size	3.61	3.95	3.52	3.74	3.48	3.64	3.44	3.61	3.42	3.56
# of children	1.32	1.47	1.25	1.26	1.21	1.19	1.18	1.17	1.14	1.15
White	0.91	0.89	0.92	0.88	0.93	0.88	0.94	0.89	0.94	0.88
HS dropout	0.21	0.20	0.18	0.20	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	0.32	0.30
College dropout	0.49	0.48	0.51	0.48	0.53	0.52	0.54	0.55	0.55	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	0.31	0.29
South	0.31	0.28	0.31	0.28	0.30	0.27	0.30	0.27	0.30	0.25
West	0.15	0.24	0.17	0.21	0.18	0.25	0.18	0.23	0.18	0.23
Husband working	0.96	0.97	0.94	0.92	0.93	0.91	0.94	0.93	0.93	0.89
Wife working	0.69	0.68	0.71	0.67	0.74	0.71	0.78	0.73	0.77	0.74
Disposable income	29,333	25,083	35,427	31,628	42,374	39,204	50,684	45,382	58,841	49,609
Food expenditure	4,447	4,554	4,868	4,543	5,294	5,079	5,872	6,021	6,604	6,289

Table II
The demand for food in the CEX

This table reports IV estimates of the demand equation for (the logarithm of) food spending in the CEX. We instrument the log of total nondurable expenditure (and its interaction with time, education, and kids dummies) with the cohort-education-year specific average of the log of the husband's hourly wage and the cohort-education-year specific average of the log of the wife's hourly wage (and their interactions with time, education and kids dummies). Standard errors are in round parenthesis; the Shea's partial R^2 for the relevance of instruments in square brackets. In all cases, the p-value of the F-test on the excluded instrument is <0.01 percent.

Variable	Estimate	Variable	Estimate	Variable	Estimate
$\ln c$	0.8503 (0.1511) [0.012]	$\ln c * 1992$	0.0037 (0.0056) [0.083]	Family size	0.0272 (0.0090)
$\ln c * \text{High School dropout}$	0.0730 (0.0718) [0.050]	$\ln c * \text{One child}$	0.0202 (0.0336) [0.150]	$\ln p_{food}$	-0.9784 (0.2160)
$\ln c * \text{High School graduate}$	0.0827 (0.0890) [0.027]	$\ln c * \text{Two children}$	-0.0250 (0.0383) [0.120]	$\ln p_{transports}$	5.5376 (8.0500)
$\ln c * 1981$	0.1151 (0.1123) [0.053]	$\ln c * \text{Three children+}$	0.0087 (0.0340) [0.197]	$\ln p_{fuel+utils}$	-0.6670 (4.7351)
$\ln c * 1982$	0.0630 (0.0837) [0.052]	One child	-0.1568 (0.3215)	$\ln p_{alcohol+tobacco}$	-1.8684 (4.1425)
$\ln c * 1983$	0.0508 (0.0704) [0.048]	Two children	0.3214 (0.3650)	Born 1955-59	-0.0385 (0.0554)
$\ln c * 1984$	0.0478 (0.0662) [0.051]	Three children+	0.0132 (0.3259)	Born 1950-54	-0.0085 (0.0477)
$\ln c * 1985$	0.0304 (0.0638) [0.064]	High school dropout	-0.7030 (0.6741)	Born 1945-49	-0.0060 (0.0406)
$\ln c * 1986$	0.0223 (0.0587) [0.068]	High school graduate	-0.8458 (0.8298)	Born 1940-44	-0.0051 (0.0348)
$\ln c * 1987$	0.0528 (0.0599) [0.065]	Age	0.0122 (0.0085)	Born 1935-39	-0.0044 (0.0273)
$\ln c * 1988$	0.0416 (0.0458) [0.049]	Age ²	-0.0001 (0.0001)	Born 1930-34	0.0032 (0.0193)
$\ln c * 1989$	0.0370 (0.0373) [0.046]	Northeast	0.0087 (0.0065)	Born 1925-29	-0.0051 (0.0140)
$\ln c * 1990$	0.0187 (0.0295) [0.060]	Midwest	-0.0213 (0.0105)	White	0.0769 (0.0129)
$\ln c * 1991$	-0.0004 (0.0318) [0.111]	South	-0.0269 (0.0096)	Constant	-0.6404 (0.9266)
OID test				20.92 (d.f. 18; χ^2 p-value 28%)	
Test that income elasticity does not vary over time				27.69 (d.f. 12; χ^2 p-value 0.6%)	

Table III
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	0.0832 (0.0089)	-0.0196 (0.0035)	-0.0018 (0.0032)
1981	0.0717 (0.0075)	-0.0220 (0.0034)	-0.0074 (0.0037)
1982	0.0718 (0.0051)	-0.0226 (0.0035)	-0.0081 (0.0026)
1983	0.0783 (0.0066)	-0.0209 (0.0034)	-0.0094 (0.0042)
1984	0.0805 (0.0055)	-0.0288 (0.0036)	-0.0034 (0.0032)
1985	0.1090 (0.0180)	-0.0379 (0.0074)	-0.0019 (0.0038)
1986	0.1023 (0.0077)	-0.0354 (0.0054)	-0.0115 (0.0038)
1987	0.1116 (0.0097)	-0.0375 (0.0051)	0.0016 (0.0046)
1988	0.0925 (0.0080)	-0.0313 (0.0042)	-0.0021 (0.0032)
1989	0.0883 (0.0067)	-0.0280 (0.0059)	-0.0035 (0.0034)
1990	0.0924 (0.0095)	-0.0296 (0.0049)	-0.0067 (0.0050)
1991	0.0818 (0.0059)	-0.0299 (0.0040)	NA
1992	0.1177 (0.0079)	NA	NA

Table IV
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	0.1275 (0.0097)	-0.0526 (0.0076)	0.0022 (0.0056)
1981	0.1197 (0.0116)	-0.0573 (0.0084)	0.0025 (0.0043)
1982	0.1322 (0.0110)	-0.0641 (0.0087)	0.0006 (0.0060)
1983	0.1532 (0.0159)	-0.0691 (0.0100)	-0.0056 (0.0067)
1984	0.1869 (0.0173)	-0.1003 (0.0163)	-0.0131 (0.0089)
1985	0.2019 (0.0244)	-0.0872 (0.0194)	NA
1986	0.1628 (0.0184)	NA	NA
1987	NA	NA	NA
1988	NA	NA	NA
1989	NA	NA	NA
1990	0.1751 (0.0221)	-0.0602 (0.0062)	-0.0057 (0.0067)
1991	0.1646 (0.0142)	-0.0696 (0.0100)	NA
1992	0.1467 (0.0130)	NA	NA

Table V
The consumption-income growth covariance matrix

Year	$\text{cov}(\Delta y_t, \Delta c_t)$	$\text{cov}(\Delta y_{t+1}, \Delta c_t)$	$\text{cov}(\Delta y_t, \Delta c_{t+1})$
1980	0.0040 (0.0041)	0.0013 (0.0039)	0.0053 (0.0037)
1981	0.0116 (0.0036)	-0.0056 (0.0032)	-0.0043 (0.0036)
1982	0.0165 (0.0036)	-0.0064 (0.0031)	-0.0006 (0.0039)
1983	0.0215 (0.0045)	-0.0085 (0.0049)	-0.0075 (0.0043)
1984	0.0230 (0.0052)	-0.0030 (0.0043)	-0.0119 (0.0050)
1985	0.0197 (0.0068)	-0.0035 (0.0047)	-0.0035 (0.0065)
1986	0.0179 (0.0048)	-0.0015 (0.0052)	NA
1987	NA	NA	NA
1988	NA	NA	NA
1989	NA	NA	0.0030 (0.0040)
1990	0.0077 (0.0045)	0.0045 (0.0065)	-0.0016 (0.0042)
1991	0.0112 (0.0044)	0.0011 (0.0049)	-0.0071 (0.0042)
1992	0.0082 (0.0048)	NA	NA
Test $\text{cov}(\Delta y_{t+1}, \Delta c_t) = 0$ for all t		p-value 25%	
Test $\text{cov}(\Delta y_{t+2}, \Delta c_t) = 0$ for all t		p-value 27%	
Test $\text{cov}(\Delta y_{t+3}, \Delta c_t) = 0$ for all t		p-value 74%	
Test $\text{cov}(\Delta y_{t+4}, \Delta c_t) = 0$ for all t		p-value 68%	

Table VI
Minimum distance partial insurance and variance estimates

This table reports DWMD results of the parameters of interest. We also estimate time-varying variances of measurement error in consumption (results not reported for brevity). See the main text for details. Standard errors in parenthesis.

		Whole sample	No College	College	Born 1940s	Born 1930s
σ_{ζ}^2 (Variance perm. shock)	1979-81	0.0102 (0.0035)	0.0067 (0.0037)	0.0099 (0.0053)	0.0074 (0.0035)	0.0057 (0.0072)
	1982	0.0207 (0.0041)	0.0154 (0.0053)	0.0252 (0.0060)	0.0210 (0.0061)	0.0166 (0.0075)
	1983	0.0301 (0.0057)	0.0317 (0.0075)	0.0233 (0.0089)	0.0184 (0.0058)	0.0246 (0.0086)
	1984	0.0274 (0.0049)	0.0333 (0.0074)	0.0176 (0.0060)	0.0219 (0.0077)	0.0224 (0.0102)
	1985	0.0293 (0.0096)	0.0287 (0.0073)	0.0204 (0.0151)	0.0187 (0.0066)	0.0333 (0.0225)
	1986	0.0222 (0.0060)	0.0173 (0.0068)	0.0312 (0.0101)	0.0222 (0.0077)	0.0111 (0.0114)
	1987	0.0289 (0.0063)	0.0202 (0.0073)	0.0354 (0.0098)	0.0307 (0.0080)	0.0079 (0.0111)
	1988	0.0157 (0.0069)	0.0117 (0.0079)	0.0183 (0.0110)	0.0155 (0.0076)	0.0007 (0.0099)
	1989	0.0185 (0.0059)	0.0107 (0.0101)	0.0274 (0.0061)	0.0176 (0.0082)	0.0217 (0.0182)
	1990-92	0.0134 (0.0042)	0.0092 (0.0045)	0.0216 (0.0065)	0.0081 (0.0059)	0.0063 (0.0091)
σ_{ε}^2 (Variance trans. shock)	1979	0.0415 (0.0059)	0.0465 (0.0096)	0.0302 (0.0056)	0.0314 (0.0054)	0.0342 (0.0070)
	1980	0.0318 (0.0039)	0.0330 (0.0053)	0.0284 (0.0059)	0.0269 (0.0056)	0.0306 (0.0072)
	1981	0.0372 (0.0035)	0.0364 (0.0053)	0.0253 (0.0046)	0.0319 (0.0058)	0.0267 (0.0064)
	1982	0.0286 (0.0039)	0.0376 (0.0063)	0.0214 (0.0042)	0.0264 (0.0049)	0.0342 (0.0078)
	1983	0.0286 (0.0037)	0.0372 (0.0063)	0.0186 (0.0037)	0.0190 (0.0045)	0.0284 (0.0077)
	1984	0.0351 (0.0039)	0.0405 (0.0059)	0.0305 (0.0051)	0.0223 (0.0047)	0.0453 (0.0100)
	1985	0.0380 (0.0075)	0.0356 (0.0056)	0.0496 (0.0130)	0.0280 (0.0062)	0.0504 (0.0115)
	1986	0.0544 (0.0058)	0.0474 (0.0076)	0.0452 (0.0085)	0.0261 (0.0060)	0.0672 (0.0153)
	1987	0.0480 (0.0054)	0.0520 (0.0082)	0.0421 (0.0071)	0.0440 (0.0093)	0.0499 (0.0095)
	1988	0.0383 (0.0047)	0.0472 (0.0074)	0.0343 (0.0060)	0.0386 (0.0068)	0.0543 (0.0148)
1989	0.0369 (0.0068)	0.0539 (0.0126)	0.0219 (0.0051)	0.0360 (0.0070)	0.0493 (0.0132)	
1990-92	0.0506 (0.0040)	0.0536 (0.0062)	0.0345 (0.0049)	0.0429 (0.0060)	0.0753 (0.0127)	
θ (Serial correl. trans. shock)		0.1132 (0.0247)	0.1268 (0.0318)	0.1086 (0.0341)	0.1324 (0.0442)	0.1706 (0.0470)
σ_{ξ}^2 (Variance unobs. slope heterog.)		0.0105 (0.0041)	0.0074 (0.0079)	0.0141 (0.0040)	0.0122 (0.0064)	0.0001 (0.0090)
ϕ (Partial insurance perm. shock)		0.6423 (0.0945)	0.9439 (0.1783)	0.4194 (0.0924)	0.7928 (0.1848)	0.6889 (0.2393)
ψ (Partial insurance trans. shock)		0.0533 (0.0435)	0.0768 (0.0602)	0.0273 (0.0550)	0.0675 (0.0705)	-0.0381 (0.0737)
P-value test of equal ϕ		23%	99%	8%	81%	18%
P-value test of equal ψ		75%	33%	29%	76%	4%

Table VII
Minimum distance partial insurance and variance estimates

This table reports DWMD results of the parameters of interest. We also estimate time-varying variances of measurement error in consumption (results not reported for brevity). See the main text for details. Standard errors in parenthesis.

Consumption: Income: Sample:	Non dur. Net Income Baseline	Non dur. Earnings only Baseline	Non dur. Male earnings Baseline
ϕ (Partial insurance perm. shock)	0.6423 (0.0945)	0.3100 (0.0574)	0.2245 (0.0493)
ψ (Partial insurance trans. shock)	0.0533 (0.0435)	0.0633 (0.0309)	0.0502 (0.0294)

Table VIII
Minimum distance partial insurance and variance estimates, various sensitivity analyses

This table reports DWMD results of the parameters of interest. We also estimate time-varying variances of measurement error in consumption (results not reported for brevity). See the main text for details. Standard errors in parenthesis.

Consumption: Income: Sample:	Non dur. Net Income Baseline	Non dur. Excluding help Baseline	Non dur. Net income Low wealth	Non dur. Net income High wealth	Total Net income Low wealth	Non dur. Net income Baseline+SEO
ϕ (Partial insurance perm. shock)	0.6423 (0.0945)	0.6215 (0.0895)	0.8489 (0.2848)	0.6248 (0.0999)	1.0342 (0.3517)	0.7652 (0.1031)
ψ (Partial insurance trans. shock)	0.0533 (0.0435)	0.0500 (0.0434)	0.2877 (0.1143)	0.0106 (0.0414)	0.3683 (0.1465)	0.1211 (0.0354)

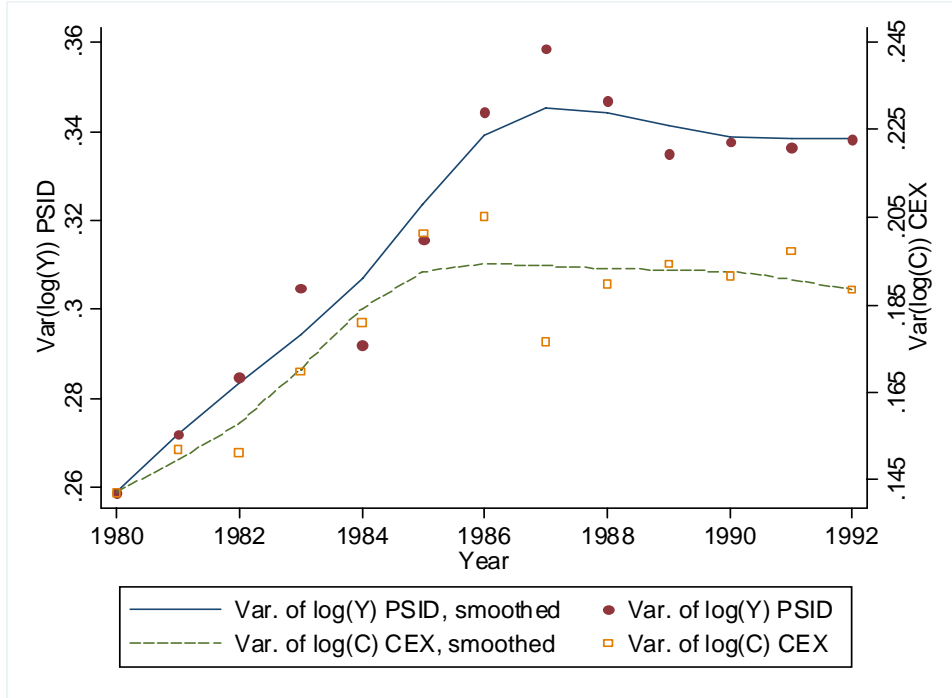


Figure 1: The Overall Pattern of Inequality.

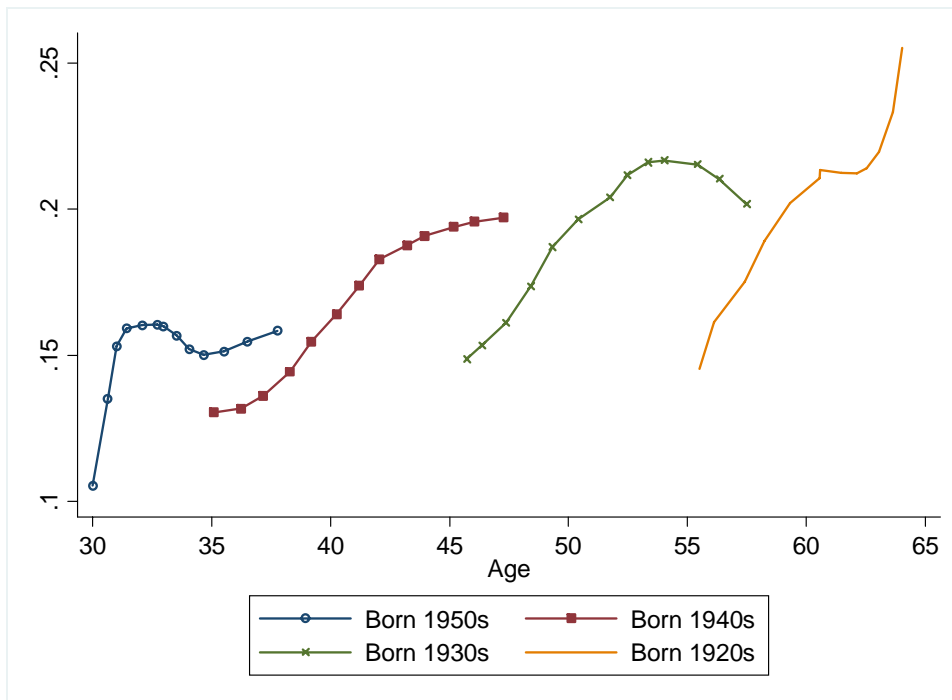


Figure 2: The Variance of Log Consumption Over the Life Cycle

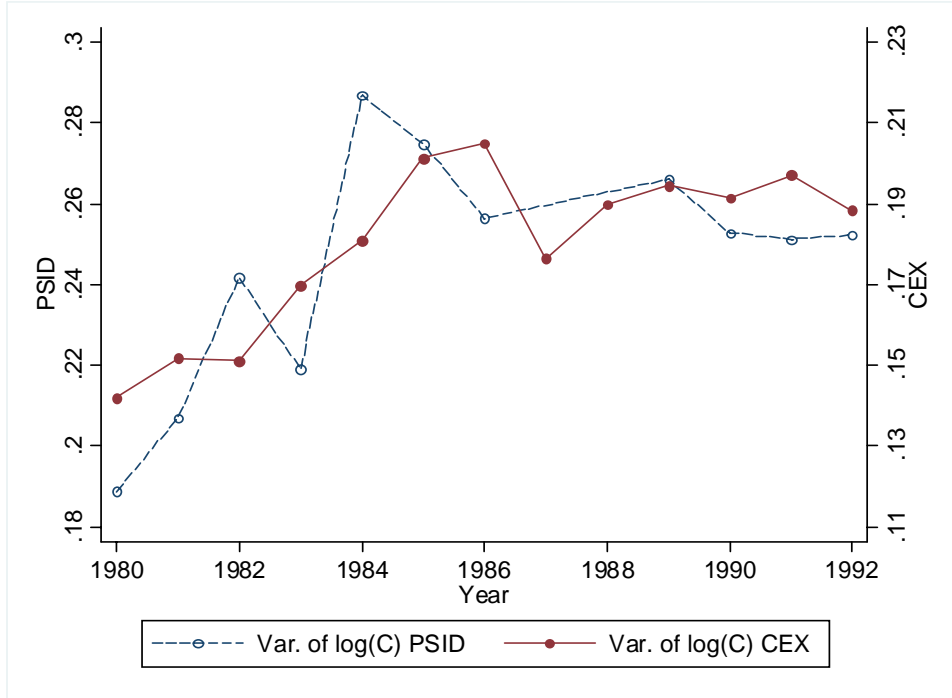


Figure 3: The CEX and New PSID Compared.

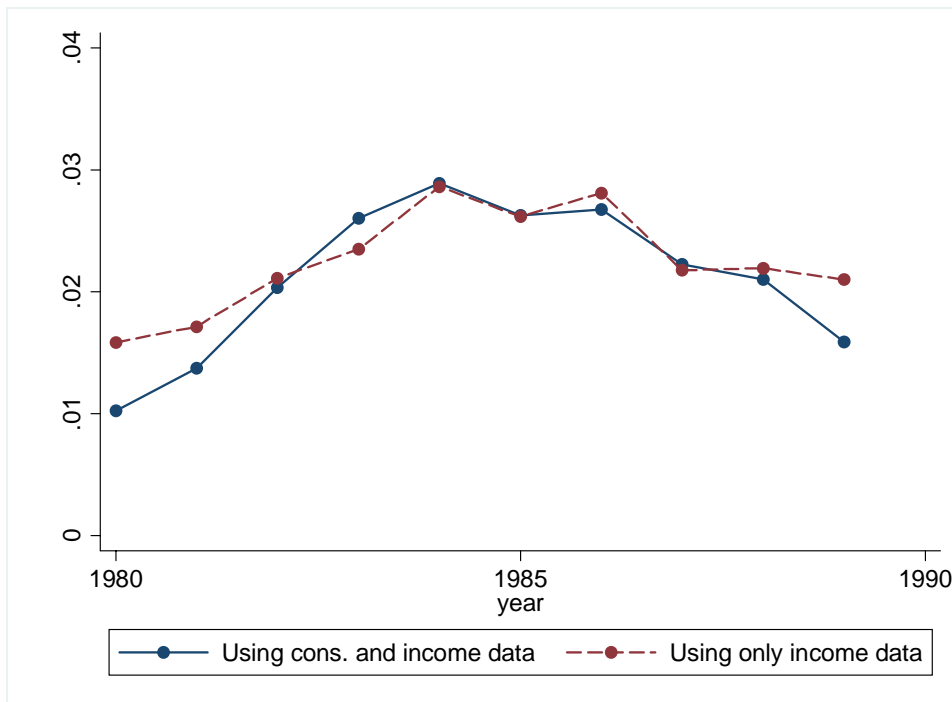


Figure 4: The Variance of Pemanent Shocks in the 1980s

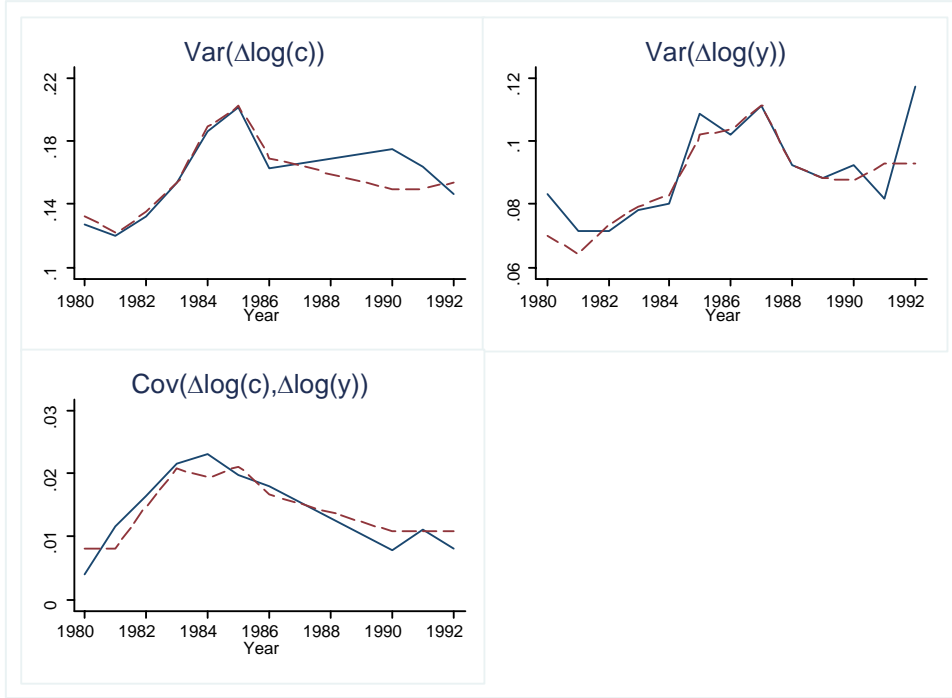


Figure 5: The Goodness of Fit of the Model.

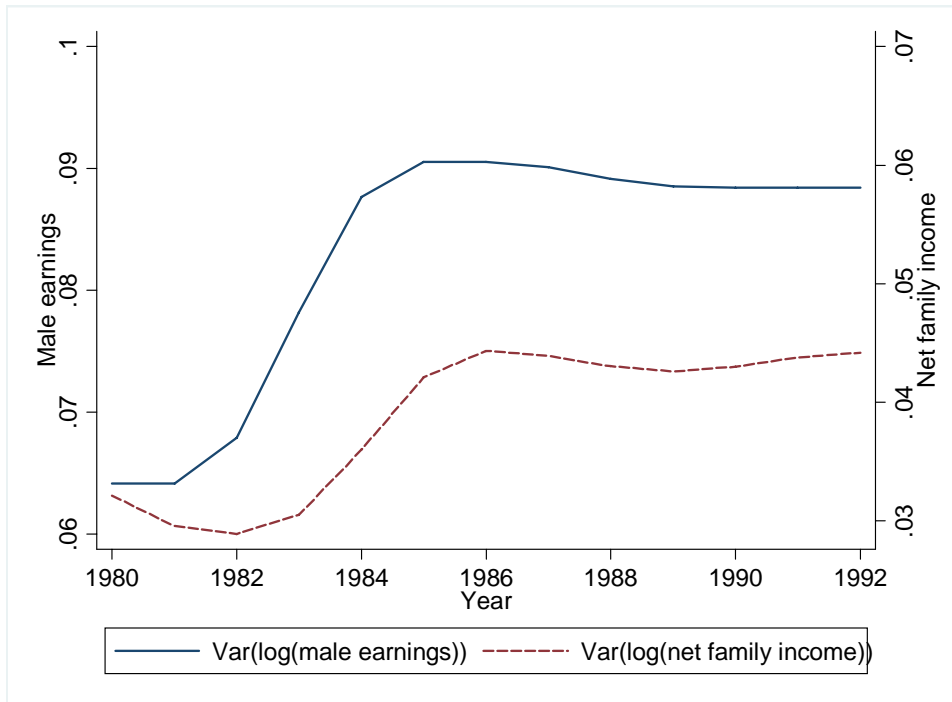


Figure 6: The Variance of Transitory Shocks.

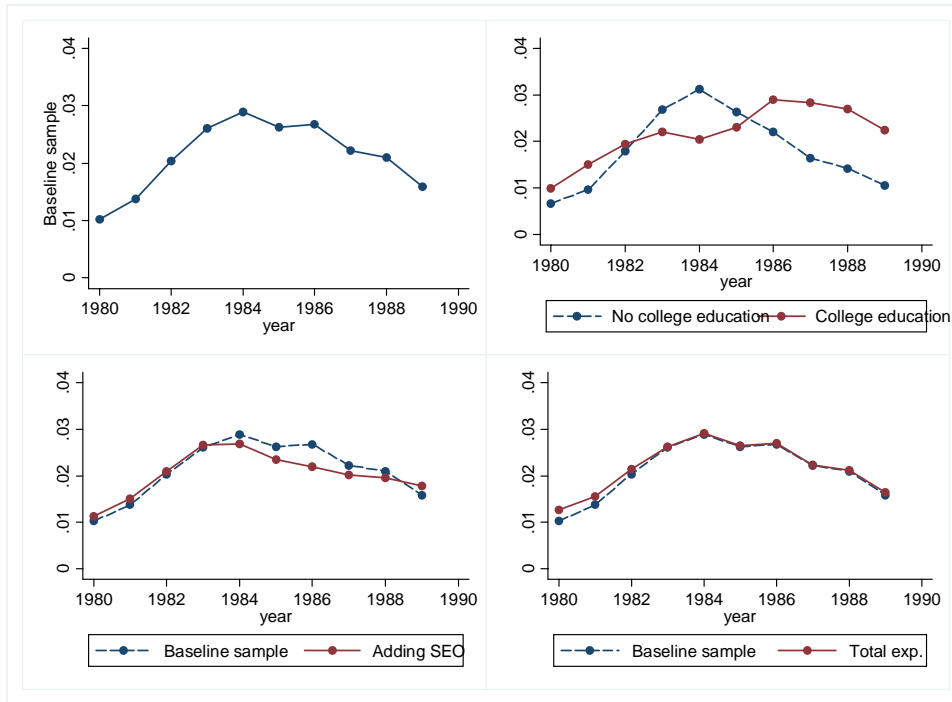


Figure 7: The Variance of Permanent Shocks in Various Specifications and Samples.

A.1 Appendix: Data

A.1.1 The PSID

Since the PSID has been widely used for microeconomic research, we shall only sketch the description in this appendix.³⁷ 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 education, food spending, 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, the timing of the survey questions on food expenditure is much less clear (see Hall and Mishkin, 1982, and Altonji and Siow, 1987, for two alternative views). 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 federal taxes on non-financial income and deflates the corresponding value by the CPI. We assume that federal taxes on non-financial income are a proportion of total federal taxes given by the ratio between non-financial income and total income. Before 1991, federal taxes 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. For these two years, we impute taxes using the TAXSIM program at the NBER. 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 more than high school and high school 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 A.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 and those with missing report on race, education, and region³⁸. We keep continuously married couples and drop some income outliers.³⁹ We then drop those born before 1920 or after 1959.

For most of the analysis we exclude 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. We also check sensitivity of results to including people aged 25-30. The final sample used in the minimum distance exercise below is composed of 17,604 observations and 1,765 households. Our income regressions do not use 36 observations with topcoded income, financial income, or federal taxes. 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, 1940s, and 1950s. Years where cell size is less than 50 are discarded.

³⁷See Hill (1992) for more details about the PSID.

³⁸When possible, we impute values for education and region of residence using adjacent records on these variables.

³⁹An 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 in a given year.

A.1.2 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.⁴⁰ 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 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.⁴¹

As the PSID, the CEX collects information on a variety of socio-demographic variables, including income and consumer expenditure. 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 are used.

Our initial 1980-2004 CEX sample includes 1,848,348 monthly observations, corresponding to 192,564 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, and those with outlier income (defined as a level of income below the amount spent on food) or incomplete income responses. The final sample used to estimate the food demand equation in Table I contains 14,430 households. Table AII details the sample selection process in the CEX.

Table A.I
Sample selection in the PSID

Reason for exclusion	# dropped	# remain
(Initial sample from family files, 1967-1992)	NA	172,274
Latino subsample	8,403	163,871
Intermittent headship	17,855	146,016
Interviewed prior to 1978	52,448	93,568
Change in family composition	18,561	75,007
Female head	23,806	51,201
Missing values	1,071	50,130
Change in marital status	5,737	44,393
Income outliers	2,386	42,007
Born before 1920 or after 1959	8,362	33,645
Poverty subsample	12,455	21,190
Aged less than 30 or more than 65	3,586	17,604

⁴⁰A description of the survey, including more details on sample design, interview procedures, etc., may be found in “Chapter 16: Consumer Expenditures and Income”, from the BLS Handbook of Methods.

⁴¹There is some evidence that trends in consumption inequality measured in the two CEX surveys have diverged in the 1990s (Attanasio, Battistin and Ichimura, 2004). While research on the reasons for this divergence is clearly warranted, our analysis, which uses data up to 1992, will only be marginally affected.

Table A.II
Sample selection in the CEX

Reason for exclusion	# dropped	# remain
(Initial sample, 1980-2004)	NA	192,564
Missing expenditure data	1,488	191,076
Present for less than 12 months	98,926	92,150
Observed after 1992	47,901	44,249
Zero before-tax income	1,321	42,928
Missing values	4,015	38,913
Marital status	16,141	22,772
Born before 1920 or after 1959	4,696	18,076
Aged less than 30 or more than 65	1,860	16,216
Income outliers and incomplete income response	1,786	14,430

A.2 Appendix: The Euler Equation Approximation

If preferences are quadratic (and interest rates are not subject to uncertainty), it is possible to obtain a closed form solution for consumption. It is also straightforward to derive an exact mapping between the expectation error of the Euler equation for consumption and income shocks. See Hall and Mishkin (1982), for example. Quadratic preferences have well known undesirable features, such as increasing risk aversion and lack of a precautionary motive for saving. More realistic preferences, such as the CRRA functional form used here, solve these problems but deliver no closed form solution for consumption. The Euler equation can be linearized to describe the behavior of consumption growth. In this appendix we derive an approximation of the mapping between the expectation error of the Euler equation and the income shock.

Consider the consumption problem faced by household i of age t . Assuming that preferences are of the CRRA form, the objective is to choose a path for consumption C so as to:

$$\max_C E_t \sum_{j=0}^{T-t} \frac{1}{(1+\delta)^j} \frac{C_{i,t+j}^\beta - 1}{\beta} e^{Z'_{i,t+j} \vartheta_{t+j}}. \quad (14)$$

where $Z_{i,t+j}$ incorporates taste shifters. Maximization of (14) is subject to the budget constraint which in the self-insurance model assumes individuals have access to a risk free bond with real return r_{t+j}

$$A_{i,t+j+1} = (1+r_{t+j})(A_{i,t+j} + Y_{i,t+j} - C_{i,t+j}) \quad (15)$$

$$A_{i,T} = 0 \quad (16)$$

with $A_{i,t}$ given. We set the retirement age after which labor income falls to zero at L , assumed known and certain, and the end of the life-cycle at age T . We assume that there is no uncertainty about the date of death. With budget constraint (15) optimal consumption choices can be described by the Euler equation

$$C_{i,t-1}^{\beta-1} = \frac{1+r_{t-1}}{1+\delta} e^{\Delta Z'_{i,t} \vartheta_t} E_{t-1} C_{i,t}^{\beta-1}. \quad (17)$$

As it is, equation (17) is not useful for empirical purposes. We then consider approximating it in the following way. In general, the logarithm of the sum of an arbitrary series X_t, X_{t+1}, \dots, X_S can be written as:

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

Taking a Taylor expansion around $\ln X_{t+k} = \ln X_t + \sum_{i=0}^k \delta_{t+i}$, $k = 1, \dots, S-t$ for some path of increments

$\delta_t, \delta_{t+1}, \dots, \delta_S$ with $\delta_t = 0$,

$$\begin{aligned} \ln \sum_{k=0}^{S-t} X_{t+k} &\simeq \ln X_t + \ln \left[1 + \sum_{k=1}^{S-t} \exp\left(\sum_{i=0}^k \delta_{t+i}\right) \right] \\ &+ \sum_{k=1}^{S-t} \frac{\exp\left(\sum_{i=0}^k \delta_{t+i}\right)}{\left[1 + \sum_{k=1}^{S-t} \exp\left(\sum_{i=0}^k \delta_{t+i}\right) \right]} \left(\ln X_{t+k} - \ln X_t - \sum_{i=0}^k \delta_{t+i} \right) \\ &\simeq \sum_{k=0}^{S-t} \alpha_{t+k,S}^{\delta} [\ln X_{t+k} - \ln \alpha_{t+k,S}^{\delta}] \end{aligned} \quad (18)$$

where $\alpha_{t+k,S}^{\delta} = \exp\left(\sum_{i=0}^k \delta_{t+i}\right) / \left[1 + \sum_{k=1}^{S-t} \exp\left(\sum_{i=0}^k \delta_{t+i}\right) \right]$ and the error in the approximation is $\mathcal{O}\left(\sum_{k=0}^{S-t} (\ln X_{t+k} - \ln X_t - \sum_{i=0}^k \delta_{t+i})^2\right)$.

Applying this approximation to the Euler equation (17) above gives:

$$\Delta \log C_{i,t} \simeq \Delta Z'_{i,t} \vartheta'_t + \eta_{i,t} + \Omega_{i,t}$$

where $\vartheta'_t = (1 - \beta)^{-1} \vartheta_t$, $\eta_{i,t}$ is a consumption shock with $E_{t-1} \eta_{i,t} = 0$, $\Omega_{i,t}$ captures any slope in the consumption path due to interest rates, impatience or precautionary savings and the error in the approximation is $\mathcal{O}(E_t \eta_{i,t}^2)$. Suppose that any idiosyncratic component to this gradient to the consumption path can be adequately picked up by a vector of deterministic characteristics $\Gamma_{i,t}$ and a stochastic individual element $\xi_{i,t}$

$$\Delta \log C_{i,t} - \Gamma_{i,t} - \Delta Z'_{i,t} \vartheta'_t = \Delta c_{i,t} \simeq \eta_{i,t} + \xi_{i,t}.$$

From (4) we also have

$$\Delta y_{i,t+k} = \zeta_{i,t+k} + \sum_{j=0}^q \theta_j \varepsilon_{i,t+k-j}.$$

The intertemporal budget constraint is

$$\sum_{k=0}^{T-t} Q_{t+k} C_{i,t+k} = \sum_{k=0}^{L-t} Q_{t+k} Y_{i,t+k} + A_{i,t}$$

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 $Q_t = 1$).

Applying the approximation (18) appropriately to each side

$$\begin{aligned} &\sum_{k=0}^{T-t} \alpha_{t+k,T}^{\Omega + \Delta Z \vartheta' - r} [\ln C_{i,t+k} - \ln Q_{t+k} - \ln \alpha_{t+k,T}^{\omega - r}] \\ &\simeq \pi_{i,t} \sum_{k=0}^{L-t} \alpha_{t+k,L}^{\Delta Z \vartheta' - r} [\ln Y_{i,t+k} - \ln Q_{t+k} - \ln \alpha_{t+k,L}^{\Delta Z \varphi - r}] \\ &+ (1 - \pi_{i,t}) \ln A_{i,t} - [(1 - \pi_{i,t}) \ln(1 - \pi_{i,t}) + \pi_{i,t} \ln \pi_{i,t}] \end{aligned}$$

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

Taking differences in expectations gives

$$\eta_{i,t} \simeq \pi_{i,t} [\zeta_{i,t} + \gamma_{t,L} \varepsilon_{i,t}]$$

where $\gamma_{t,L} = \left(\sum_{j=0}^q \alpha_{t+j,L}^{\Delta Z \vartheta' - r} \theta_j \right)$ and the error on the approximation is $\mathcal{O}([\zeta_{i,t} + \gamma_{t,L} \varepsilon_{i,t}]^2 + E_{t-1} [\zeta_{i,t} + \gamma_{t,L} \varepsilon_{i,t}]^2)$.

Then⁴²

$$\Delta c_{i,t} \simeq \xi_{i,t} + \pi_{i,t} \zeta_{i,t} + \gamma_{t,L} \pi_{i,t} \varepsilon_{i,t}$$

with a similar order of approximation error.

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

⁴²Blundell, Low and Preston (2004) contains a lengthier derivation of such an expression, including discussion of the order of magnitude of the approximation error involved.

A.3 Appendix: Identification

The simplest model

Here we show how the model can be identified with four years of data $(t+1, t, t-1, t-2)$, and discuss various extensions. We start with the simplest model with no measurement error, serially uncorrelated transitory component, and stationarity. This is relaxed later. (Unexplained) consumption and income growth in period s ($s = t-1, t, t+1$) are, respectively:

$$\begin{aligned}\Delta c_s &= \xi_s + \phi \zeta_s + \psi \varepsilon_s \\ \Delta y_s &= \zeta_s + \Delta \varepsilon_s\end{aligned}$$

(where for simplicity we have assumed that the transitory shock to income is i.i.d.).⁴³ The parameters to identify are: $\phi, \psi, \sigma_\xi^2, \sigma_\zeta^2$, and σ_ε^2 .

As in Meghir and Pistaferri (2004), we can prove that:

$$E(\Delta y_t (\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})) = \sigma_\zeta^2 \quad (19)$$

and that:

$$E(\Delta y_t \Delta y_{t-1}) = E(\Delta y_{t+1} \Delta y_t) = -\sigma_\varepsilon^2 \quad (20)$$

Identification of σ_ε^2 through (20) rests on the idea that income growth rates are autocorrelated due to mean reversion caused by the transitory component (the permanent component is subject to i.i.d. shocks). Identification of σ_ζ^2 through (19) rests on the idea that the variance of income growth ($E(\Delta y_t \Delta y_t)$), subtracted the contribution of the mean reverting component ($E(\Delta y_t \Delta y_{t-1}) + E(\Delta y_t \Delta y_{t+1})$), coincides with the variance of innovations to the permanent component.

In general, if one has T years of data, only $T-3$ variances of the permanent shock can be identified, and only $T-2$ variances of the i.i.d. transitory shock can be identified. As said in the text, with panel data on income, the variances of permanent and transitory shock can be identified without recourse to consumption data.

One can also prove that:

$$\frac{E(\Delta c_t (\Delta y_{t-1} + \Delta y_t + \Delta y_t))}{E(\Delta y_t (\Delta y_{t-1} + \Delta y_t + \Delta y_t))} = \phi \quad (21)$$

$$\frac{E(\Delta c_t \Delta y_{t+1})}{E(\Delta y_t \Delta y_{t+1})} = \psi \quad (22)$$

$$E(\Delta c_t (\Delta c_{t-1} + \Delta c_t + \Delta c_{t+1})) - \frac{[E(\Delta c_t (\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1}))]^2}{E(\Delta y_t (\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1}))} + \frac{[E(\Delta c_t \Delta y_{t+1})]^2}{E(\Delta y_t \Delta y_{t+1})} = \sigma_\xi^2 \quad (23)$$

These moment conditions provide complete identification of the parameters of interest. Identification of ψ using (22) uses the fact that income and lagged consumption may be correlated through the transitory component ($E(\Delta c_t \Delta y_{t+1}) = \psi \sigma_\varepsilon^2$). Scaling this by $E(\Delta y_t \Delta y_{t+1}) = \sigma_\varepsilon^2$ identifies the loading factor ψ . Note that there is a simple IV interpretation here: ψ is identified by a regression of Δc_t on Δy_t using Δy_{t+1} as an instrument. A similar reasoning applies to (21): the current covariance between consumption and income growth ($E(\Delta c_t \Delta y_t)$), stripped of the contribution of the transitory component, reflects the arrival of permanent income shocks ($E(\Delta c_t (\Delta y_{t-1} + \Delta y_t + \Delta y_t)) = \phi \sigma_\zeta^2$). Scaling this by the variance of permanent income shock, identified by using income moments alone, identifies the loading factor ϕ . Note that here too there is a simple IV interpretation: ϕ is identified by a regression of Δc_t on Δy_t using $(\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})$ as an instrument. Finally, (23) identifies the variance of the component σ_ξ^2 using a residual variability idea: the variance of consumption growth, stripped of the contribution of permanent and transitory income shocks, reflects heterogeneity in the consumption gradient.

Measurement error in consumption

Consider now the realistic case in which consumption is measured with error, i.e.,

$$c_{i,t}^* = c_{i,t} + u_{i,t}^c$$

⁴³The proof mechanism can be easily extended to deal with MA(q) transitory shock processes as in Meghir and Pistaferri (2004).

where c^* denote measured consumption, c is true consumption, and u^c the measurement error. Measurement error in consumption induces serial correlation in consumption growth. Because consumption is a martingale with drift in the absence of measurement error, the variance of measurement error can be readily recovered using

$$E(\Delta c_t^* \Delta c_{t-1}^*) = E(\Delta c_t^* \Delta c_{t+1}^*) = -\sigma_{u^c}^2 \quad (24)$$

The other parameters of interest are still identified by (21)-(23), replacing the unobserved c with the measured c^* . One obvious reason for the presence of measurement error in consumption is our imputation procedure. From (1), we can write $c_t^* = c_t + \beta(D_{it})^{-1} e_{it}$. Note that the measurement error $\beta(D_{it})^{-1} e_{it}$ will be non-stationary (which we account for in estimation).

Measurement error in income

Assume now that income is also measured with error, i.e.,

$$y_{i,t}^* = y_{i,t} + u_{i,t}^y$$

Now it can be proved that ϕ and $\sigma_{u^c}^2$ are still identified by (21) and (24) (replacing the unobserved y and c with the measured y^* and c^*). However, σ_ε^2 and $\sigma_{u^y}^2$ cannot be told apart, and ψ (as well as σ_ξ^2) thus remains unidentified.⁴⁴ It is possible however to put a lower bound on ψ using the fact that:

$$\psi \geq \frac{E(\Delta c_t^* \Delta y_{t+1}^*)}{E(\Delta y_t^* \Delta y_{t+1}^*)}$$

Thus it is possible to argue that the estimate of ψ in Tables VI-VIII is downward biased due to measurement error in income. Using estimates contained in Meghir and Pistaferri (2004), a back-of-the-envelope calculation shows that the variance of measurement error in earnings accounts for approximately 30% of the variance of the overall transitory component of earnings. Given that our estimate of ψ is close to zero in most cases, an adjustment using this inflation factor would make little difference empirically. To give an example, the estimate of ψ in Table VI, column 1, would increase from 0.055 to 0.079.

Using a similar reasoning, one can argue that we have an upper bound for σ_ξ^2 , in that

$$\sigma_\xi^2 \leq E(\Delta c_t^* (\Delta c_{t-1}^* + \Delta c_t^* + \Delta c_{t+1}^*)) - \frac{[E(\Delta c_t^* (\Delta y_{t-1}^* + \Delta y_t^* + \Delta y_{t+1}^*))]^2}{E(\Delta y_t^* (\Delta y_{t-1}^* + \Delta y_t^* + \Delta y_{t+1}^*))} + \frac{[E(\Delta c_t^* \Delta y_{t+1}^*)]^2}{E(\Delta y_t^* \Delta y_{t+1}^*)}$$

The bias, however, is likely negligible. Using the same out-of-the-envelope calculation above, we calculate that the estimate of σ_ξ^2 in Table VI, column 1, would decrease from 0.0122 to 0.0121. For this reason, in what follows we assume for simplicity that income is measured without error in deriving the various identification restrictions.

Non-stationarity

Allowing for non-stationarity and with T years of data, the following moments can be used to identify the variance of the permanent shock:

$$E(\Delta y_s^* (\Delta y_{s-1}^* + \Delta y_s^* + \Delta y_{s+1}^*)) = \sigma_{\zeta,s}^2 \quad (25)$$

for $s = 3, 4, \dots, T-1$. The variance of the transitory shock can be identified using:

$$-E(\Delta y_s^* \Delta y_{s+1}^*) = \sigma_{\varepsilon,s}^2 \quad (26)$$

for $s = 2, 3, \dots, T-1$. With an MA(1) process for the transitory component, the analog of (25) and (26) become:

$$E(\Delta y_s^* (\Delta y_{s-2}^* + \Delta y_{s-1}^* + \Delta y_s^* + \Delta y_{s+1}^* + \Delta y_{s+2}^*)) = \sigma_{\zeta,s}^2$$

for $s = 4, 5, \dots, T-2$, and (assuming θ is being already identified)⁴⁵

$$-E(\Delta y_s^* \Delta y_{s+2}^*) = \theta \sigma_{\varepsilon,s}^2$$

⁴⁴If one is willing to assume that ε is non-stationary while u^y is stationary, then ψ is effectively identified using

$$\psi = \frac{E(\Delta c_t^* \Delta y_{t+1}^* - \Delta c_{t-1}^* \Delta y_t^*)}{E(\Delta y_t^* \Delta y_{t+1}^* - \Delta y_{t-1}^* \Delta y_t^*)}$$

⁴⁵The parameter θ is identified by non-linear moment conditions, which we omit here.

for $s = 2, 3, \dots, T-2$. In our case, $s = 1$ corresponds to 1978 and $s = T$ corresponds to 1992. These are the restrictions that we impose in the empirical analysis.

The other parameters of interest ($\sigma_{u^c}^2, \phi, \psi, \sigma_\xi^2$) can be identified using:

$$\begin{aligned} -E(\Delta c_s^* \Delta c_{s+1}^*) &= \sigma_{u^c}^2 \\ \frac{E(\Delta c_s^* \Delta y_{s+1}^*)}{E(\Delta y_s^* \Delta y_{s+1}^*)} &= \psi \end{aligned}$$

pooling data for $s = 2, 3, \dots, T-1$, and

$$\begin{aligned} \frac{E(\Delta c_s^* (\Delta y_{s-1}^* + \Delta y_s^* + \Delta y_{s+1}^*))}{E(\Delta y_s^* (\Delta y_{s-1}^* + \Delta y_s^* + \Delta y_{s+1}^*))} &= \phi \\ E(\Delta c_s^* (\Delta c_{s-1}^* + \Delta c_s^* + \Delta c_{s+1}^*)) - \frac{[E(\Delta c_s^* (\Delta y_{s-1}^* + \Delta y_s^* + \Delta y_{s+1}^*))]^2}{E(\Delta y_s^* (\Delta y_{s-1}^* + \Delta y_s^* + \Delta y_{s+1}^*))} + \frac{[E(\Delta c_s^* \Delta y_{s+1}^*)]^2}{E(\Delta y_s^* \Delta y_{s+1}^*)} &= \sigma_\xi^2 \end{aligned}$$

pooling data for $s = 3, 4, \dots, T-2$.

Note that we can allow for time-varying insurance parameters:

$$\Delta c_s = \xi_s + \phi_s \zeta_s + \psi_s \varepsilon_s + \Delta u_s^c$$

which would be identified by the moment conditions:

$$\begin{aligned} \frac{E(\Delta c_s^* \Delta y_{s+1}^*)}{E(\Delta y_s^* \Delta y_{s+1}^*)} &= \psi_s \\ \frac{E(\Delta c_s^* (\Delta y_{s-1}^* + \Delta y_s^* + \Delta y_{s+1}^*))}{E(\Delta y_s^* (\Delta y_{s-1}^* + \Delta y_s^* + \Delta y_{s+1}^*))} &= \phi_s \end{aligned}$$

for all $s = 2, 3, \dots, T-1$ and $s = 3, 4, \dots, T-2$ respectively. These are the moment conditions that we use when we allow the insurance parameters to vary over time.⁴⁶ Equation (24) still identifies the variance of measurement error in consumption pooling all available data. Again, this is because consumption growth is autocorrelated only because of measurement error (in the absence of it, it would be a martingale).

More general model

Suppose consumption growth is now given by

$$\Delta c_s^* = \xi_s + \phi_0 \zeta_s + \phi_1 \zeta_{s-1} + \psi_0 \varepsilon_s + \psi_1 \varepsilon_{s-1} + \Delta u_s^c$$

while income growth is still:

$$\Delta y_s^* = \zeta_s + \varepsilon_s - \varepsilon_{s-1}$$

In this case, we assume consumption growth depends on current and lagged income shocks. The parameters to identify (in the stationary case for simplicity) are $\phi_0, \phi_1, \psi_0, \psi_1, \sigma_\xi^2, \sigma_{u^c}^2, \sigma_\zeta^2$, and σ_ε^2 . The variances of the income shocks are still identified by:

$$E(\Delta y_t^* (\Delta y_{t-1}^* + \Delta y_t^* + \Delta y_{t+1}^*)) = \sigma_\zeta^2$$

and:

$$E(\Delta y_t^* \Delta y_{t-1}^*) = E(\Delta y_{t+1}^* \Delta y_t^*) = -\sigma_\varepsilon^2$$

However, one can prove that, of all the ‘‘insurance’’ coefficients, only ψ_0 can be identified, using

$$\frac{E(\Delta c_t^* \Delta y_{t+1}^*)}{E(\Delta y_t^* \Delta y_{t+1}^*)} = \psi_0$$

while all the others remain not identified in the absence of further restrictions. For example, the expression we used to identify ϕ in the baseline scenario,

⁴⁶We experienced convergence problems in the most flexible specification when we allow for yearly variation in ϕ and ψ . We thus imposed the restriction that they are the same across the sub-periods 1979-84 and 1985-92.

$$\frac{E(\Delta c_t^* (\Delta y_{t-1}^* + \Delta y_t^* + \Delta y_{t+1}^*))}{E(\Delta y_t^* (\Delta y_{t-1}^* + \Delta y_t^* + \Delta y_{t+1}^*))}$$

would now identify the sum $(\phi_0 + \phi_1)$. Increasing the number of lags of income shocks in the consumption income growth equation has no effects: ψ_0 is still identified, while the other insurance parameters are not.

A.4 Appendix: Estimation details

The two basic vectors of interest are:

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

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. Since PSID consumption data were not collected in 1987 and 1988, the vector $\Delta \mathbf{c}_i$ is understood to have $\dim(\Delta y_i) - 3$, i.e., the rows with missing consumption data have already been swept out from $\Delta \mathbf{c}_i$. Moreover, if the individual was not interviewed in year t , we replace the unobservable $\Delta c_{i,t}$ and $\Delta y_{i,t}$ with zeros. Conformably with the vectors above, we define:

$$\mathbf{d}_i^c = \begin{pmatrix} d_{i,1}^c \\ d_{i,2}^c \\ \dots \\ d_{i,T}^c \end{pmatrix} \text{ and } \mathbf{d}_i^y = \begin{pmatrix} d_{i,1}^y \\ d_{i,2}^y \\ \dots \\ d_{i,T}^y \end{pmatrix}$$

where $d_{i,t}^c = 1 \{ \Delta c_{i,t} \text{ is not missing} \}$ and $d_{i,t}^y = 1 \{ \Delta y_{i,t} \text{ is not missing} \}$. Overall, this notation allows us to handle in a simple manner the problems of unbalanced panel data and of missing consumption data in 1987 and 1988.

Stacking observations on Δy and Δc (and on d^c and d^y) for each individual we obtain the vectors:

$$\mathbf{x}_i = \begin{pmatrix} \Delta \mathbf{c}_i \\ \Delta \mathbf{y}_i \end{pmatrix} \text{ and } \mathbf{d}_i = \begin{pmatrix} \mathbf{d}_i^c \\ \mathbf{d}_i^y \end{pmatrix}$$

Now we can derive:

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

where \oslash denotes an elementwise division. The vector \mathbf{m} contains the estimates of $\text{cov}(\Delta y_t, \Delta y_{t+s})$, $\text{cov}(\Delta y_t, \Delta c_{t+s})$, and $\text{cov}(\Delta c_t, \Delta c_{t+s})$, a total of $T(2T+1)$ unique moments.⁴⁷ To obtain the variance-covariance matrix of \mathbf{m} , define conformably with \mathbf{m} the individual vector:

$$\mathbf{m}_i = \text{vech} \{ \mathbf{x}_i \mathbf{x}_i' \}$$

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

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

where $\mathbf{D}_i = \text{vech} \{ \mathbf{d}_i \mathbf{d}_i' \}$ 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} :

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

where $\mathbf{\Upsilon}$ 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:

⁴⁷In practice there are less than $T(2T+1)$ moments because data on consumption are not available all years.

$$\begin{pmatrix} \text{var}(\Delta c_1) \\ \text{cov}(\Delta c_1, \Delta c_2) \\ \dots \\ \text{cov}(\Delta c_1, \Delta c_T) \\ \dots \end{pmatrix} = \begin{pmatrix} \phi^2 \text{var}(\zeta_1) + \psi^2 \text{var}(\varepsilon_1) + \text{var}(\xi_1) + \text{var}(u_1^c) + \text{var}(u_0^c) \\ -\text{var}(u_1^c) \\ \dots \\ 0 \\ \dots \end{pmatrix} + \Upsilon$$

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 diagonally-weighted minimum distance (DWMD) 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 $\widehat{\mathbf{\Lambda}}$.