

CONSUMPTION AND INCOME PERSISTENCE ACROSS GENERATIONS

Giovanni Gallipoli* Hamish Low† Aruni Mitra‡

February 15, 2017

Abstract

We examine the persistence of economic outcomes within families. First, we link parents and children in U.S. longitudinal data and document the cross-generational association in both incomes and expenditures. Next, we develop a richer model of intra-family persistence and derive a set of theoretical moment restrictions. This allows us to identify and estimate key parameters describing how different shocks influence long-term economic outcomes. We use these estimates to quantify the contribution of parental factors to the observed dispersion of income and consumption among adult children.

1 Introduction

Households make consumption, savings and labour supply decisions in an uncertain environment characterized by stochastic income fluctuations. It is therefore reasonable to expect that shocks to parental income have significant and possibly delayed implications for their children. This might happen through their impact on human capital formation or through other, often unobserved, transfers. Such impacts are likely to vary by family characteristics like race, family size and education.

Much of the existing research linking family outcomes across generations has focused on income persistence, and very rarely on wealth (Charles and Hurst 2003). However,

*Vancouver School of Economics, UBC, Canada. E-mail: gallipol@mail.ubc.ca. Financial support from SSHRC is gratefully acknowledged.

†Faculty of Economics, University of Cambridge, UK. E-mail: hamish.low@econ.cam.ac.uk.

‡Vancouver School of Economics, UBC, Canada. E-mail: aruni.mitra@alumni.ubc.ca.

partly due to data constraints, relatively little is known about the persistence of consumption expenditures and, more generally, about the joint determination of income and consumption across generations.

Changes in a household's financial circumstances can induce intergenerational effects with empirically testable implications in both expenditure and income. In this paper we explore these empirical implications using a life-cycle model with intergenerational linkages. This model delivers a set of moment restrictions that we use to estimate key parameters dictating the responses of income, wealth and consumption expenditures to different shocks. We use the estimates of these parameters to quantify the importance of parental factors for the dispersion of different economic outcomes in the child's generation.

We begin the analysis by documenting the robust reduced-form relationships linking parental and child income and consumption expenditures. Measuring the reduced-form elasticities of intergenerational persistence is done by regressing the child variable on the parental variable, and interpreting the regression coefficient as the measure of persistence. In fact, motivated by the large increase in income inequality, much empirical research has looked for evidence of a significant decline in intergenerational income mobility in the United States over the past few decades.¹ In our initial reduced-form analysis we adopt the procedure suggested by Lee and Solon (2009, henceforth LS). Their approach stands out for the use of multiple years of income data for each individual in the sample, which makes the intergenerational elasticity estimates more precise. We use the same regression methodology as LS, but extend their analysis in the sense that we not only examine the impact of parental permanent family income on the child's adult family income, but also consider the impact of other parental variables like permanent labour earnings and consumption on the (adult) child's consumption expenditures and labour earnings.

Empirical analyses involving consumption expenditures were made impossible before the 1980s by the absence of longitudinal micro-level data on individual expenditures. Even after 1980, when the Consumption Expenditure Survey (CE) started providing detailed consumption expenditure data at the micro-level, the scope of these data was

¹This is usually done by running regressions like the one described above on data for different periods. The estimates obtained in such empirical studies are rather diverse - not least because of the different estimation techniques employed to measure the elasticity (see Aaronson et al. (2008) for a survey of the empirical findings).

insufficient for intergenerational studies because individuals were followed for a maximum of four quarters. In contrast, the Panel Study of Income Dynamics (PSID), while following individuals for longer periods of time, suffered from the fact that data was consistently collected only for housing and food-related expenditures. Only after 1999 new and more detailed consumption categories were added to the PSID. Since 2005 the PSID expenditure survey collects all the categories captured by the CE. However, using the detailed consumption data available after 2005 (or 1999) it is not feasible to perform a long-term analysis linking parental consumption during the formative years of the child to the consumption of the child when he/she grows up to be a household head. One alternative way to proceed is to step outside of the standard LS framework and, instead, regress current parental consumption on current child's consumption as an adult. This is the approach taken by Charles et al. (2014), as they regress the log of 2005-2009 average consumption expenditure of the adult children on that of the parents in the *same years*.² However, this type of contemporaneous analysis cannot identify the impact on adult children of shocks to parental income and consumption when the child was growing up. To address this shortcoming, we pursue a different strategy. That is, we impute the consumption for years prior to 1999 using the fairly detailed consumption expenditure information available in the PSID from the 1999 interview wave onwards, and use this imputed consumption data series in the LS analysis.

Having obtained reliable measures of the elasticities linking income and consumption across generations of the same family, we examine the mechanism underlying this persistence. To study the mechanics of the interdependence between income and consumption we posit a model with infinitely-lived agents who make consumption-savings choices in an environment with exogenous, persistent shocks to permanent (earned) income. The assumption of infinitely-lived agents is made to account for the fact that individuals care about their offspring, above and beyond any human capital transfers made to the children during their formative years. This framework allows us to decompose the total variance in consumption and income of a child into deterministic parental factors, exogenous shocks, and other idiosyncratic child-specific factors. The intergenerational linkages in income and consumption are assumed to stem from intra-family persistence of earned

²Along with a host of controls like cubic polynomials in ages of the child and the parent, race, marital status, family size, household head dummy and homeowner status.

income as well as from saving and transfer decisions.³ To our knowledge, this is the first study in the literature of intra-family linkages, to model an optimal consumption decision along with a standard income process. Crucially, this modelling choice introduces an endogenous covariance between unobserved income components (assets) and permanent (earned) income while jointly estimating family persistence in earnings, consumption and unobserved income. We provide a discussion of model identification (in the Appendix), and estimate the structural parameters using the Generalized Method of Moments (GMM) by minimizing the distance between observed and analytical second order moments of the income and consumption processes of parents and their adult children.

We assess the impact of human capital, as measured through the educational level of the individuals, in the overall intergenerational persistence of permanent incomes.⁴ This is done by comparing the estimates of intergenerational elasticity of permanent income obtained from running the model separately with and without controlling for the impact of educational status of individuals on their income and consumption measures. Also, this model assumes that all extra-economic linkages within the family, e.g., genetic or emotional influences, are completely reflected through the persistence in earnings and unobserved income.

Section 2 of the paper overviews the PSID dataset, used throughout our analysis, and also provides a detailed account of the consumption imputation exercise. Section 3 describes the reduced-form analysis of intergenerational persistence within families, while Section 4 develops the model and sets out the details of the GMM estimation exercise. Section 5 reports the results of our counterfactual analysis and inequality accounting, shedding some light on the role of observable and unobservable parental factors in shaping inequality in long-term children's outcomes. Section 6 concludes with a brief discussion and summary of the main findings.

³Due to the absence of comprehensive asset data, the role of wealth is often ignored and not modelled. For example, Cox et al. (2004) effectively treat savings as an unobserved factor called '*tastes*', and do not consider any covariance of such tastes with the permanent incomes of individuals.

⁴We do not further decompose the intergenerational linkages into the transmission channel of parental human capital or that of financial resources, nor we examine the indirect role of human capital in the determination of parental permanent incomes (see Lefgren et al. (2012)).

2 Data: Availability, Quality and Scope

Data for the current analysis is sourced from the PSID, administered by the University of Michigan's Survey Research Center. The longitudinal survey began in 1968 with a national probability sample of almost 5,000 U.S. families. The sample families were reinterviewed annually from 1968 through 1997, and since then biennially. We focus our study only on the non-Latino, non-immigrant households within the Survey Research Center (SRC) component of the PSID. For reasons discussed in Brown (1996), we exclude the households in the Survey of Economic Opportunity (SEO) component of the PSID where poor households were over-sampled.

The dataset from PSID is particularly conducive for intergenerational analyses because it has followed children from the original sample as they have grown into adulthood and become household heads themselves. Our analysis pertains to children born between 1952 and 1985. To avoid over-representation of children who left their homes at a later stage of their lives, the sample excludes children born before 1952 (that is, those children who were older than 16 years at the time of the first 1968 PSID interview). Also, it is recognized that an individual's income (consumption) to be at least somewhat indicative of his/her long run income (consumption), one must observe that person's income (consumption) no earlier than the age of 25 years. Hence, the sample is cut off after the cohort born in 1985 as the last year of available data is 2010 (in the 2011 interview). The first year in which children's income is observed is 1977 (as reported in the 1978 interview) - the year in which the 1952 cohort reached age 25. Consequently, we can observe the 1952 cohort for ages between 25 and 58, but the 1985 cohort is only observed at their age 25. On the other hand, parents who have ages greater than 65 years are dropped from the analysis to avoid complications arising from retirement. In our structural analysis, we use an even stricter sample restriction where we keep observations for only male individuals who have at least 5 years of data in the PSID, have non-negative labour earnings, work for less than 5840 hours annually, are not self-employed, and do not have annual or biennial earnings growth rate of more than 400 percent or less than -200 percent.

The labour earnings data for the household head and the total family income data are readily available for most survey rounds of the PSID. Nevertheless the consumption data is quite sparse across the years and not even presented as a single variable in the PSID.

Different consumption expenditure categories have to be suitably summed up (using appropriate weights depending on the frequency of consumption in a particular category, e.g., yearly, monthly, weekly, etc.) to arrive at an aggregate measure of consumption expenditure. Now, there are 11 major categories of consumption variables, namely, (i) food, (ii) housing, (iii) child-care, (iv) education, (v) transportation, (vi) healthcare, (vii) recreation and entertainment, (viii) trips and vacation, (ix) clothing and apparel, (x) home repairs and maintenance, and (xi) household furnishings and equipment. Of these 11 major categories, food and housing are most consistently observed across the years - expenditure on food is observed from 1968 through 2013 interviews barring only 1973, 1988 and 1989, while housing expenditure is observed on all years except 1978, 1988 and 1989. Child-care expenditure data is available for 24 rounds of interview - 1970-1972 (3 interview years), 1976, 1977, 1979 and 1988-2013 (18 interview years). Education, transportation and health-care are only observed for the last 8 PSID interviews (biennially from 1999 through 2013). The rest of the categories from (vii) through (xi) are observed for the last 5 interviews (biennially from 2005 to 2013). From this list of availability of expenditure categories for different waves of the PSID survey we understand that a simple sum of the expenditure categories for different years will not be the correct way to create the total consumption variable because every year will have different consumption categories available. There are two ways to correct this systematic bias in the calculation of the total consumption variable - either take the measure of consumption to be equal to just the expenditure on food because it is the most consistently observed category, or impute the consumption of the missing categories using the information on food expenditure. Since the first technique will involve ignoring variation in the consumption of non-durable goods other than food, we choose the other alternative. While performing the imputation we have skipped the consumption expenditure categories that were added to the PSID from the 2005 wave. This is done to keep the measure of consumption consistent over years and also maximize the number of categories that can be used. Moreover, the categories added from the 2005 wave collectively constitute a very small fraction of the total consumption. Below we present the detailed imputation procedure.

2.1 Imputation of Consumption Data

Andreski et al. (2014) compares the consumption expenditure data from CE and PSID by mapping the CE expenditure items into the broader categories of the PSID and imputing exact values for categorical responses and missing/invalid responses. They find that expenditures in almost all categories of consumption vary widely across the two datasets, e.g., home repairs and maintenance in the PSID are approximately twice as large as the CE estimates, and the PSID home insurance expenditures are 40 to 50 percent higher than their CE counterparts. In spite of these inconsistencies (obviously due to differences in survey methodologies and sampling techniques), the mean of expenditure data since 1999 in PSID and CE were shown to be comparing favourably by Li et al. (2010), and the consumption expenditures in the two datasets vary broadly in a similar way with observable household characteristics like age of household head, household size, educational attainment, marital status, race and homeownership.

Emboldened by this consistency amongst the PSID, CE, and also the aggregate consumption estimates in the NIPA (National Income and Product Accounts) data, Attanasio and Pistaferri (2014) (henceforth AP) impute consumption data for the missing consumption categories in the PSID before 1999 by using the more detailed data post 1999. Their extrapolation is consistent with theories of consumer demand in the sense that the allocation of total resources spent in a given period over different commodities is dependent on relative prices, taste-shifters (namely, demographic and socioeconomic variables), and total expenditure.

In this paper, we will follow the consumption imputation procedure of AP using the following approximated demand system:

$$\ln(N_{it}) = Z'_{it}\omega + p'_i\pi + g(F_{it}; \lambda) + u_{it}, \quad (2.1)$$

where N is consumption net of food expenditure, Z are the socioeconomic controls (viz., dummies for age, education, marital status, race, state of residence, employment status, self-employment, head's hours worked, homeownership, disability, family size, and the number of children in the household), p are the relative prices (the overall CPI and the CPIs for food at home, food away from home, and rent), F is the total food expenditure (i.e., sum of food at home, food away from home, and food stamps) that is observed in the

PSID consistently through the years, $g(\cdot)$ is a polynomial function, and u is the error term. This equation is estimated using data from the 1999-2011 PSID waves (or 1998-2010 calendar years), where the net consumption measure N_{it} is the sum of annualized expenditures on home insurance, electricity, heating, water, other miscellaneous utilities, car insurance, car repairs, gasoline, parking, bus fares, taxi fares, other other transportation, school tuition, other school expenses, child care, health insurance, out-of-pocket health, and rent.⁵ The measure for rent equals the actual annual rent payments for renters and is imputed to 6 percent of the self-reported house value (Flavin et al. (2002)) for the homeowners.

After estimating the logarithm of the net consumption by running a pooled OLS regression on equation (2.1), we construct a measure of imputed total consumption as

$$\hat{C}_{it} = F_{it} + \exp \{ Z'_{it} \hat{\omega} + p'_i \hat{\pi} + g(F_{it}; \hat{\lambda}) \} \quad (2.2)$$

This measure is then corrected for inflation by dividing it by the overall CPI, and finally it is transformed into adult-equivalent terms using the OECD adult equivalence scale, $(1 + 0.7(A - 1) + 0.5K)$, where A is the number of adults and K the number of children in the household unit.

One natural question to ask is how well the imputed consumption values match with the observed values during the period when both data series are available. A natural choice for a measure of the goodness of fit is the R^2 of the regression (2.1), which is found to be 0.47. However, what we are really interested in is matching the standard deviations of the observed and imputed series because we would be using only the second order moments of income and consumption for estimating our semi-structural model in Section 4. We find that our imputed consumption series can match the observed series quite closely both in terms of standard deviation, and a more general non-linear measure like the Gini coefficient. Figure A1 in Appendix A presents the Gini coefficients (normalized to their initial values in 1998) of log imputed and actual consumption (in Panel C), and also compares the standard deviations of actual and imputed consumption, with those of real income and labour earnings (in Panels A, B, and D). The top-coded values for total

⁵One thing to note in the definition of net consumption is that we have excluded food expenditure from it to avoid endogeneity issues in the regression. Also, despite the availability of the 2013 PSID wave, we decided to stop at the 2011 wave because some of the sub-categories of expenditure like health insurance are still not updated in the PSID.

family income in the PSID are replaced with the estimates obtained from fitting a Pareto distribution to the upper tail of the family income distribution.

3 Measuring Intergenerational Mobility

With the imputed consumption series and the two measures of income - labour earnings and total family income, we are now in a position to study the intergenerational mobility with respect to all these three measures of economic well-being. We will divide our analysis in two parts - reduced form descriptive analysis using regressions and mobility matrices, and a more structural way where we would assume specific data-generating processes for income and consumption. The next two subsections in this section deal with the reduced form analyses, while Section 4 will focus on the semi-structural model and its estimation using the GMM technique.

3.1 Elasticities of Family Income, Labour Earnings, and Consumption

The most natural way to measure the impact of parental economic well-being on the child is to look at the intergenerational elasticities of income and consumption, which, by definition, answers by what percentage the child's income or consumption changes when there is one percentage change in the corresponding parental variable. The elasticities are nothing but the regression coefficients when the child's log income or log consumption is regressed on that of the parent. However, there are many identification issues that one must take care of while estimating such a regression. We will be using the most widely used regression specification in the literature, which we will refer to as the LS specification, but also point out its possible limitations following Hertz (2007).

3.1.1 Regression Specification

Intergenerational elasticities are ideally measured as the impact of permanent parental income and expenditure on the child's permanent income and consumption. However, we do not observe the permanent income, earnings or consumption of any individual. So, we must proxy these permanent variables by some function of the current variables. Since we wish to produce a time trend of the intergenerational elasticities, we use the

current income (consumption) variable of the child as a proxy for the child's permanent variable. A child's income (consumption) data are observed over a range of ages, which changes over time. This necessitates controlling for the child's age variation. One way of dealing with this issue is to take into account the income (consumption) data at a particular age (say 30 years) for all children. This approach is adopted by Lopoo and Mayer (2005). However, this is clearly not efficient as it throws out most of the data available. So, we adopt the LS technique of taking the child data for all the years available, along with a full set of age controls in the regression specification. One point of concern that still remains after controlling for child age effects is that of the measurement error in the child variable. Usually a measurement error in the regressand (in this case the child variable) does not bias the parameter estimates, but Haider et al. (2006) show why using current income (consumption) of the child as a proxy for his or her permanent lifetime income (consumption) does not fit into this classical measurement error case. As LS summarizes, the systematic heterogeneity across children in their growth rates of income (consumption) over the life-cycle lead to the measurement error in log of the current child variable to be mean-reverting only in the early years of the life-cycle, and mean-departing later on. Because children with higher permanent income (consumption) have steeper growth trajectories, this means that the inequality in current child variable is likely to under-state or over-state the inequality in the permanent child variable depending on which stage of the life-cycle one is observing the current variable. Since we observe the children from ages 25 to 58 years, which essentially covers most of their working lives, we must take care of this life-cycle heterogeneity in variance of the child's current variable. We do this by including as controls the interaction of the permanent parental variable (income or consumption) with the quartic of the child's age. However, there is an implicit identifying assumption in the claim that the inclusion of this interaction term will take care of the heterogeneity in economic growth profiles, viz., the age-income (or age-consumption) profile is the same for all 34 cohorts of children born between 1952 and 1985. Haider et al. (2006) suggest that the measurement error bias from the regressand is inconsequential if the current variable is measured at around age 40. So, we also centre the child's age around 40 years, that is, the age of a child in year t , who was born in year c , is calculated as $(t - c - 40)$. The actual age of the child is $(t - c)$ years, but it is normalized to equal 0 at age 40 years. This helps in interpreting the time-series estimates of the intergenerational elasticity as

how the intergenerational elasticity at age 40 of the child evolves as successive cohorts of children pass through that same age. However, one must be careful not to confuse these yearly estimates of intergenerational elasticities with the cohort-specific elasticities. In fact, these intergenerational elasticities at age 40 for a given year can be interpreted as an asymmetrical moving average of the cohort-specific elasticities for the cohorts of adult children who are observed for that particular year. It is asymmetrical because the older cohorts weigh more in a particular year's estimate owing to the fact that cohorts enter as they turn 25 years of age but never leave till the end of the PSID dataset. This asymmetry, which makes the interpretation of the elasticity estimates a bit difficult, can be easily solved by making cohorts exit after a certain age, but that would lead to missing out on valuable information for those omitted cohorts. The only meaningful alternative to our estimation of time trend of this cohort-weighted intergenerational elasticity is to estimate cohort-specific elasticities using lifetime average of income (or consumption) for the adult children. However, since we are mainly interested in the time series of intergenerational persistence to compare with the time trend of income and consumption inequality, we will be estimating the yearly elasticities instead of the cohort-specific ones.

Once we have solved the issues arising from proxying the child permanent variable by the corresponding current variable, we now need to come up with a suitable proxy for long-run income (consumption) data for the parents that will serve as the principal regressor. A short-run measure of parental income (consumption) used as a proxy introduces an attenuation bias in the estimation of the intergenerational elasticity in the long-run income (consumption). However, so long as this downward bias is roughly stable over time, it will not distort the estimation of changes in intergenerational mobility. We take the average log annual family income (consumption) over the years when the child was between 15 and 17 years old as the proxy for the long-run income (consumption) of the parent.⁶ Obviously, the age of the parents of different children born in a particular cohort will not be the same when the children reach the age of 15 to 17. So, one needs to control for the average age of the parental household head over those ten years.

⁶Note that the choice of 15 years as the starting child age for parental observation is dictated by data availability because that is the earliest for the oldest 1952 cohort when their parental data is available (starting from 1967, as documented in the 1968 interview). Also, notice that we could have taken the average of parental income (consumption) for the parents' entire lifetime (till 65 years of age), but that would rule out any variation in the permanent parental variable across siblings born at different life-stages of the parents.

Keeping in mind the above platitude about the different controls that we should have in our regression specification, we include the quartics for parental age and the child's age as controls, as also the interaction between permanent parental income (consumption) and the quartic of the child's age,

$$\zeta_{ict} = \mu' D_t + \beta_t x_{ic} + \gamma(PA) + \delta(CA) + \theta X + \epsilon_{ict} \quad (3.1)$$

where

$$\gamma = [\gamma_1, \gamma_2, \gamma_3, \gamma_4], \delta = [\delta_1, \delta_2, \delta_3, \delta_4], \theta = [\theta_1, \theta_2, \theta_3, \theta_4],$$

$$PA = [PA_{ic}, PA_{ic}^2, PA_{ic}^3, PA_{ic}^4]',$$

$$CA = [(t - c - 40), (t - c - 40)^2, (t - c - 40)^3, (t - c - 40)^4]', \text{ and}$$

$$X = [x_{ic}(t - c - 40), x_{ic}(t - c - 40)^2, x_{ic}(t - c - 40)^3, x_{ic}(t - c - 40)^4]'$$

In equation (3.1) above, which is the LS specification, D_t denotes the time dummies for each $t = 1977(1)1996$ and $t = 1998(2)2010$; x_{ic} denotes the average parental log income (consumption) variable described earlier for child i born in cohort c for all $c = 1952(1)1985$; PA_{ic} and its higher order polynomials control for the average parental age at the time x_{ic} is observed; and CA is the vector of the quartic in child's age centred around 40 years. Peters (1992) points out that the error term ϵ_{ict} reflects factors like luck in labour and marriage markets, intergenerational transmission of genetic traits and other environmental factors. We are mainly interested in the coefficients β_t because these values capture the time variation in intergenerational economic mobility. We will be calculating five different types of the β_t coefficient vector - $\beta_{ee}, \beta_{yy}, \beta_{cc}, \beta_{ce}$ and β_{cy} , where the first letter of the subscript denotes the choice of ζ_{ict} , and the second letter corresponds to x_{ic} in equation (3.1); the subscript e refers to log labour earnings, y denotes log total income, and c denotes log imputed consumption. For example, β_{ce} denotes the intergenerational elasticity between parental labour earnings and child consumption. Section 3.1.2 below discusses the estimates from the regressions.

Before analyzing the regression estimates, we should discuss some of the possible criticisms of the regression specification in equation (3.1). Hertz (2007) points out that there may be observable sources of heterogeneity, such as parental and child educational attainment, number of adult family members, race, etc., that affect the age-income (or age-consumption) profile of the child even after controlling for parental income through

the interaction terms. Also, since we are estimating an average cross-sectional effect of parental income for every year, our regression specification cannot allow the inclusion of individual fixed effects for the age-income (or age-consumption) profile in the same regression. The only way to take care of all these heterogeneities is to run a separate first stage regression of the child economic variable on a polynomial of child age, individual fixed effects, and the interaction of variables like education, race, etc. with the polynomial in child age, and finally regress the fitted values from this first-stage regression on the parental variables. While the omission of such heterogeneity can make the elasticity estimates inconsistent, correcting for them is not without its own issues. This is because if one includes as controls the interaction of such variables with the child's age in the first stage regression as noted above, then one should not include the interaction terms of parent's income (or consumption) because that would introduce a spurious covariance between parental income (or consumption) and the child's predicted outcome. Thus, we find that one essentially faces a choice between the parental economic status variable on one hand, and the family-specific observed characteristics on the other, to include as controls. We decided to include the interaction of parental income (or consumption) with the child's age to maintain comparability with the majority of the literature.


3.1.2 Results from Regressions

Figure 1 shows the time trends of the five different intergenerational elasticities. Even though the estimates are noisy, some long term patterns do emerge from them.⁷ Clearly, we find that both labour earnings mobility and total family income mobility have reduced over the years for the period under study. There appears to be two major time periods when this reduction in mobility occurred in the U.S. - once prior to the mid 1980s, and again from the early 2000s. Between 1985 and 2000, the intergenerational persistence was largely stable. Also, we find that the intergenerational elasticity of total income is always higher in magnitude than that of labour earnings, implying significant intergenerational economic linkages beyond earned income outcomes. This reduced-form finding, in par-

⁷Table B1 in Appendix B shows the estimates of the five different intergenerational elasticities across the years. While interpreting these estimates, one should not attach any meaning to the absolute values of the elasticities, rather the time trend of each series. This is because the normalization of the age of the adult child to 40 years affects the magnitude of the elasticities, that is, any other normalization would change the absolute estimates but keep the general time trend intact.

ticular, will serve as the motivation for the consideration of intra-family linkages in assets (arising from persistences in saving behaviour, inter-vivos transfers, etc.) in our structural model in Section 4.

Figure 1. Estimates of Intergenerational Elasticities (1977-2010)



Note: Values on either side of missing estimates were interpolated by a straight line to maintain continuity of the time lines. The missing estimates are indicated in Table B1 in Appendix B. The trend lines ignore statistical significance.

One way to interpret the regression coefficient β is to think of it as a product of the correlation coefficient between the parental and child variables (income or consumption), and the standard deviation of the child variable relative to the standard deviation of the parental variable, $\beta_t = \frac{cov(\zeta_{ict}, x_{ic})}{[s.d.(x_{ic})]^2} = corr(\zeta_{ict}, x_{ic}) \frac{s.d.(\zeta_{ict})}{s.d.(x_{ic})}$.⁸ This interpretation shows that a

⁸Of course, this simple linear regression coefficient formula is not exactly correct in the multiple regression setting of equation (3.1), but the decomposition between correlation and relative standard deviation discussed here, still holds conceptually. Specifically, in this context, one can think of the child variable as a measure net of the controls for parental and child ages, and also the interaction of the child's age with the parental variable.

rising intergenerational elasticity can be a direct consequence of the growing inequality in the U.S. with a constant or even falling intergenerational linkage as captured by the correlation component. While this decomposition of the regression coefficient β helps us in understanding the role of inequality in shaping intergenerational persistence, one must be careful not to link the phenomenon of higher intergenerational persistence of total income than that of labour earnings with the well-established empirical finding of higher inequality in total income than in labour income.⁹ This is because intergenerational elasticity is not exactly a function of the absolute level of inequality, but the inequality in the children's generation relative to that in the parents' generation. Even if inequality in total income is always higher than that of labour earnings, it does not tell us anything about the standard deviations of income and earnings for the children relative to their parents'. Also, one must be careful while interpreting the above comparison between the magnitudes of earnings elasticity and income elasticity because the data for these two series do not pertain to the same unit of observation — while earnings is only for the household head, total family income not only includes the non-labour income component of the household head, it also captures the labour and non-labour incomes of other members of the household. Thus, an income elasticity higher than earnings elasticity can be attributed not only to a higher persistence in non-labour or wealth income but also to a higher degree of linkages in the working habits of the parental and child households, e.g., the likelihood of the wife earning labour income from outside work, etc.

As far as the intergenerational elasticities involving child consumption are concerned, the long-run trends are less clear. While there have been obvious short term trends of unambiguous rise or fall in mobility, the overall intra-family persistence involving child consumption has remained stable for the last three decades. However, there is one interesting observation to be made from the comparison of these time trends with those of income and earnings — while all the five trends have more or less followed each other quite closely till the mid 1990s, there has been a stark divergence ever since. On one

⁹Owing to the fact that the bottom half of the U.S. population does not hold any significant amount of wealth (see Saez (2014)), there is much higher inequality in total income (which includes capital income in the form of return on wealth, along with labour earnings) than in labour income. However, this fact does not hold up in the PSID sample. A cursory look at the time series of the standard deviations of log total income and log labour earnings in Panel D of Figure A1 shows that in the PSID data-sample, the standard deviation of log total income is actually smaller than that of log labour earnings for most of the years under study.

hand the income and earnings elasticities kept increasing, while the consumption elasticity started to fall and the cross elasticities of child consumption with respect to parental income and earnings decreased, or at least did not increase. Combining these findings we have an interesting situation where parental income and earnings explain more and more of child's income and earnings, and yet a smaller fraction of the child's consumption in the new millennium. This divorce between income and earnings on one hand, and consumption of the child on the other, can be due to various factors — improved access to credit market being the most obvious one. If the children's generation can use the credit market to insure away the income risks transmitted from their parents then one can expect a divergence between intergenerational persistences in income and consumption expenditures.

4 A Model of Intergenerational Persistence

Once we have studied the intergenerational mobility in income and consumption separately, we now use a semi-structural set-up with joint evolution of income and consumption of parents and their adult children to study what percentage of the total variation in child's income and consumption is explained by parental factors. The basic building blocks for this analysis include assuming a time series process for individual income, and a link between the incomes of the parent and the child. Each individual is allowed to choose his per-period consumption so as to maximize the discounted expected utility subject to a budget constraint. Thus, no additional ad hoc assumption is made regarding the time series evolution of consumption, which is in sharp contrast to Cox et al. (2004). Another important distinction with Cox et al. (2004) is our use of labour earnings instead of total family income as the measure of income in the analysis. Our consumption is derived as a sum of non-human asset wealth and human capital wealth. Now, to make the distinction between non-human wealth on which interest is earned and the human capital income, we chose labour earnings as the measure of income.

We denote a time point (or a particular year) by t . We denote the parent and the child by superscripts P and C , and a parent-child pair is denoted by the subscript i . Note that since one parent can have multiple children, the i subscript uniquely identifies the child, but not the parent. Income and consumption are denoted by Y and C , and their loga-

rithms are y and c respectively. In the analysis that follows we use measures of income and consumption which are net of their predictable individual components. In particular, we regress log-income and log-consumption on a set of observable idiosyncratic and economy-wide characteristic variables and take the residuals of those regressions as our measures of y and c respectively. The set of regressors include the full set of dummies for year, year of birth, family size, marital status, number of children, state of residence, employment status and whether a person is white or not. While Cox et al. (2004) does not include the educational level of the individuals as a control variable, there is a separate branch of literature, as exemplified most recently by Heckman et al. (2016), that looks into the importance of the persistence of education in determining intra-family economic linkages, and finds that education roughly accounts for 30 percent (see Hertz (2006)) of the total intergenerational correlation in income. To see if our structural estimation also corroborates such findings, we estimate our model once with, and again without controlling for individual educational attainment.

At any year t , the parent i will have an income $Y_{i,t}^P$ such that log income $y_{i,t}^P$ can be expressed as a sum of his log permanent income, $\bar{y}_i^P$, and an independent and identically distributed transitory shock process, $v_{i,t}^P$ with mean zero and variance $\sigma_{v^P}^2$. The child will have the same basic income process as the parents, $y_{i,t}^C = \bar{y}_i^C + v_{i,t}^C$ (where $v_{i,t}^C$ is a mean zero i.i.d. process with variance $\sigma_{v^C}^2$), except that their permanent income is going to be partly determined by their parental permanent income. In particular, we assume $\bar{y}_i^C = \rho \bar{y}_i^P + \epsilon_i^C$, where ϵ_i^C follows an i.i.d. process with variance $\sigma_{\epsilon^C}^2$. The parameter ρ captures the intergenerational elasticity in permanent income. No matter whether ρ is less than or greater than unity, the permanent income of the adult child is allowed be greater or less than that of his parent by the presence of the additive term ϵ_i^C .

To summarize, the log-income and growth rate of income of the parents and their adult children are given by the following set of equations:

$$y_{i,t}^P = \bar{y}_i^P + v_{i,t}^P \quad (4.1)$$

$$y_{i,t}^C = \rho \bar{y}_i^P + \epsilon_i^C + v_{i,t}^C \quad (4.2)$$

With the income process in place, we are now in a position to solve the utility maximization problem of a representative individual. Since the utility maximization problem

is the same for both the parents and their children, we drop the superscripts P and C in the following consumption analysis and derive a general expression for consumption ($C_{i,t}$) as a function of one's permanent income. We assume that any adult individual, standing at time point t , chooses his future stream of consumption so as to maximize the discounted present value of his expected utility subject to his budget constraint. In other words, when an individual makes his consumption decision he has the knowledge of his permanent income, but not the future income shocks. The following is the utility maximization problem for any adult person:

$$\begin{aligned} & \max_{\{C_{i,k}\}_{k=t}^{\infty}} E_t \sum_{j=0}^{\infty} \beta^j u(C_{i,t+j}) \\ \text{s.t. } & A_{i,t+1} = (1+r)(A_{i,t} + Y_{i,t} - C_{i,t}), \end{aligned}$$

where β is the discount factor, r is the real interest rate (assumed to be constant), and $A_{i,t}$ is the savings in the form of some assets. Now, assuming a quadratic utility function¹⁰, and $\beta(1+r) = 1$, the permanent income hypothesis will hold thereby implying that the consumption at time t is given by the annuity value of his non-human wealth $A_{i,t}$ and his expected future income,

$$C_{i,t} = \frac{r}{1+r} \left[A_{i,t} + \sum_{j=0}^{\infty} \left(\frac{1}{1+r} \right)^j E_t (Y_{i,t+j}) \right]$$

Now, since we have assumed a time series process for log income, we would like to write the above expression for consumption level in terms of logs of the variables. This will involve a first order Taylor series approximation of the logarithm of each variable around unity - for any variable x , $\ln(x) \simeq \ln(1) + \frac{x-1}{1} = x - 1$.¹¹ Using this approximate relationship $x \simeq 1 + \ln(x)$ for any variable x in the above expression for consumption, and denoting

¹⁰This choice of quadratic utility function is innocuous so far as one is ready to ignore the precautionary saving motive of individuals. Appendix C.1 shows that a first order Taylor expansion of the Euler equation in the constant relative risk aversion (CRRA) utility case gives back the condition as the quadratic utility case.

¹¹Note that the choice of unity as the value around which the approximation is carried out gives a very specific interpretation to the variables in their levels - they can be interpreted as fractions of the corresponding economy-wide variables.

$\ln(A_{i,t})$ by $\tilde{a}_{i,t}$, we get,

$$\begin{aligned} 1 + c_{i,t} &\simeq \frac{r}{1+r} (1 + \tilde{a}_{i,t}) + \frac{r}{1+r} \sum_{j=0}^{\infty} \left(\frac{1}{1+r} \right)^j E_t (1 + y_{i,t+j}) \\ \implies c_{i,t} &\simeq \frac{r}{1+r} (1 + \tilde{a}_{i,t}) + \bar{y}_i + \frac{r}{1+r} v_{i,t} \end{aligned}$$

Denoting $a_{i,t} \equiv \frac{r}{1+r} (1 + \tilde{a}_{i,t})$ and the function $f(r) = \frac{r}{1+r}$ ¹², we can thus write the approximate log-consumption processes for the parent and the child as follows:

$$c_{i,t}^P = a_{i,t}^P + \bar{y}_i^P + f(r) v_{i,t}^P \quad (4.3)$$

and

$$c_{i,t}^C = a_{i,t}^C + \rho \bar{y}_i^P + \epsilon_i^C + f(r) v_{i,t}^C$$

Now, the asset of the child in every period might be related to his parent's assets due to reasons like inter vivos transfers (Altonji et al. (1997)). In particular, we assume $a_{i,t}^C = \phi a_{i,t}^P + \psi_{i,t}^C$, where $\psi_{i,t}^C$'s are drawn independently from an identical distribution with variance $\sigma_{\psi^C}^2$. Thus, the consumption equation for a child is given by

$$c_{i,t}^C = \phi a_{i,t}^P + \rho \bar{y}_i^P + \epsilon_i^C + \psi_{i,t}^C + f(r) v_{i,t}^C \quad (4.4)$$

Equations (4.1) through (4.6) specify the complete set of equations that characterize the intergenerational dependence in the economy in terms of income, consumption and assets. However, the PSID does not have consistent asset information prior to 1998, which makes it unsuitable for being used as an observed variable in any intergenerational study. Therefore, we will have to treat the terms $a_{i,t}^P$ and $a_{i,t}^C$ as unobserved variations that affect consumption every period. Now, the $a_{i,t}$'s will of course contain asset variations, but they might potentially contain other variations as well. For example, in case of more general non-linear utility functions (e.g., the CRRA utility function), any linear approximation of the Euler equation will lead to the omitted higher order preference terms (like the one capturing the precautionary saving motive) being loaded onto the unobserved

¹²For the estimation exercise discussed below, we take the real interest rate $r = 0.04$, which is the average real interest rate in the United States from 1967 till 2010.

$a_{i,t}$ term. This means that intergenerational correlation in $a_{i,t}$ stems not only from inter-vivos transfers but also preference linkages. Interestingly, such higher order preference terms, like assets, can also be correlated with permanent income. This necessitates an inclusion of a covariance term between $a_{i,t}$ and the permanent income as a parameter. Although we expect assets and permanent income to be positively correlated, the sign of the aforementioned covariance parameter cannot be predicted because the $a_{i,t}$ term contains other information as well. For example, if individuals with lower permanent income are credit-constrained¹³ then a strong precautionary saving motive might generate a negative correlation between $a_{i,t}$ and permanent income.

Having established the time series evolution of income and consumption, we are now in a position to have the list of second order moments that will be used in the GMM estimation of the parameters. The 9 parameters to be estimated are as follows: the intergenerational persistences - ρ, ϕ ; the variance of parental permanent income - $\sigma_{\bar{y}_i^P}^2$; the variance of the i.i.d. shocks - $\sigma_{v^P}^2, \sigma_{v^C}^2, \sigma_{\epsilon^C}^2, \sigma_{\psi^C}^2$; the variance of parental assets - $\sigma_{a^P}^2$; and the covariance between parental asset and parental permanent income - $\sigma_{a^P, \bar{y}_i^P}$. The moment conditions that we use for estimating these parameters are the six variances of income, consumption and growth rate of income for both the parent and the child, and the six covariances involving the consumption and income of the parents and children. Thus, we have an over-identified system of ten moment conditions in nine unknown parameters.¹⁴ We minimize the distance between the observed and model-implied moments by giving equal weight to each moment condition. This is in principle nothing but a GMM estimation exercise where the weighting matrix is chosen to be the identity matrix. The estimates that we get from using the identity matrix are not used to create the optimal weighting matrix because running the GMM exercise in the second stage with this optimal weighting matrix could lead to large finite sample bias in the parameter estimates, as shown in Altonji et al. (1996). The key is to use a weighting matrix that does not depend on the parameter estimates.

In our GMM specification we are ignoring the covariances amongst the various innovations - $\epsilon_i^C, \psi_{i,t}^C, v_{i,t}^C$ and $v_{i,t}^P$, and this orthogonality assumption is crucial for our identifica-

¹³Caballé (2016) provides a theoretical model of borrowing-constrained parents choosing child education leading to intergenerational persistence.

¹⁴Appendix D presents all the moment conditions, along with the over-identification argument.

tion. Also, we are ignoring any explicit measurement errors in income and consumption. Considering measurement error is especially crucial for consumption because the data series for consumption has been imputed by us. As a robustness exercise we estimate our model using only observed consumption data from the 1999 through 2011 PSID interview rounds. We do not need to impute any consumption data for these years, but the number of observations becomes significantly smaller, particularly because of much fewer observations for parents. Some of the parameter estimates lose statistical significance with this smaller dataset, but they are quantitatively similar to our baseline estimates. Besides this consumption imputation issue, Bound et al. (1994) find that the contribution of the variance of measurement error in the observed variance of labour earnings is quite substantial. However, in our current specification, we have no room to identify the variance of the measurement errors in income and consumption variables. Nevertheless, since we do not have data on assets, we are implicitly treating the measurement errors in parental and child consumption as being subsumed inside $a_{i,t}^P$ and $\psi_{i,t}^C$. One important feature of classical measurement error is that it becomes even more serious for models involving time differences of a panel data series.¹⁵ This is probably the reason why Blundell et al. (2008), who use first differences of the income and consumption series, is explicitly concerned about measurement error. We avoid this additional problem by choosing not to impose any autoregressive persistence on the transitory shock processes, because estimation of the intertemporal persistence parameter requires the second moments of the time-differenced earnings data-series.¹⁶

¹⁵The reason is as follows: If a time series x_t is observed with some error u_t (assumed to be independently and identically distributed with mean zero and variance σ_u^2) then we can write the observed series as $x_t^{obs.} = x_t^{true} + u_t$, and its first difference as $\Delta x_t^{obs.} = \Delta x_t^{true} + \Delta u_t$. Now, $Var(x_t^{obs.}) = Var(x_t^{true}) + \sigma_u^2$, whereas $Var(\Delta x_t^{obs.}) = Var(\Delta x_t^{true}) + 2\sigma_u^2$. So the contribution of the variance of the measurement error actually doubles in the first differences.

¹⁶Apart from this practical consideration, the choice of not imposing the autoregressive persistence on the transitory shocks is innocuous because the presence of such a persistence structure will mechanically reduce the estimated variance of the transitory shocks. We however estimate our model with additional AR(1) processes for the transitory shocks $v_{i,t}^P$ and $v_{i,t}^C$ as a robustness check, and confirm that the estimates of intergenerational persistence do not change significantly from our baseline model.

Table 1: GMM Estimates

Parameters	Unbalanced Panel Data	
	No education control	Education control
ρ	0.3652 (0.0776)	0.2116 (0.0946)
ϕ	0.3234 (0.1154)	0.2570 (0.1168)
$\sigma_{v^P}^2$	0.0940 (0.0441)	0.0746 (0.2597)
$\sigma_{v^C}^2$	0.2507 (0.0100)	0.2467 (0.0109)
$\sigma_{\epsilon^C}^2$	0.0560 (0.0048)	0.0350 (0.0022)
$\sigma_{\psi^C}^2$	0.0123 (0.0033)	0.0169 (0.0025)
$\sigma_{\bar{y}_i^P}^2$	0.1735 (0.0431)	0.1374 (0.2616)
$\sigma_{a^P}^2$	0.0978 (0.0415)	0.1086 (0.2422)
$\sigma_{a^P, \bar{y}_i^P}$	-0.1063 (0.0421)	-0.1057 (0.2515)

Note: Panel bootstrap standard errors using 200 bootstrap samples from 1089 unique child I.D.'s, are reported in parentheses.

One way to get around the problem of measurement error, without explicitly modelling it, is to estimate the model using time-averaged data (see Solon (1992), and Zimmerman (1992)). As long as there is no compelling reason to believe that the PSID consistently under-reports or over-reports the earnings and expenditure data across years and across the individuals, we can assume that this time-averaging actually gets rid of the problem of classical measurement error. We run the model with both panel and time-averaged data, and collect the panel data estimates in Table 1. The time-averaged data

gives qualitatively similar results with slightly higher estimates for ρ and ϕ . We compare the estimates with and without controlling for individual educational attainment. We find that controlling for educational attainment of the individuals reduces the intergenerational elasticity in permanent income by about 42 percent. The intergenerational persistence in assets also gets reduced by roughly 20 percent. Thus we find that years of education provide a strong basis for intra-family persistence in economic well-being, with a higher impact of education on earnings than on assets and personal tastes. However, even after purging out the effect of education, the intergenerational linkage in permanent income and assets are quite strong and significant.

Performing the variance decomposition between permanent and transitory components of residual earnings using the GMM estimates in Table 1, we see that while the variance of the permanent earnings component has gone down (from roughly 0.17 in the parents' generation to about 0.08 in the children's generation), there has been a more than compensatory increase in the transitory variance (from about 0.09 for the parents to about 0.25 for the children). To check if this increase in transitory variance over time holds in the data outside our intergenerational structural model, we perform a variance decomposition of earnings into permanent and transitory shocks following Heathcote, Perri and Violante (2010) where they assume the earnings process to be a sum of a random walk and a transitory shock. We find that the variance of the transitory shock to earnings has indeed risen over the period 1967 through 2008. The details of this analysis are presented in Appendix E. This trend implies that most earnings shocks have been insured away in the recent years for the children's generation.

One must note that the variance of the transitory earnings component in our model arises from two sources: cross-sectional variation across individuals and the yearly variation for each individual. In so far as we are ignoring any temporal persistence in this transitory component (viz., an AR(1) process for v_{it}^P and v_{it}^C), we are going to over-estimate the variance of these transitory components. For example, introducing the AR(1) persistence in the transitory shocks reduces the estimate of $\sigma_{v^P}^2$ to 0.0751 and that of $\sigma_{v^C}^2$ to 0.1277. To put these estimates in perspective, in one of their many specifications, Moffitt and Gottschalk (2011) estimates a log real earnings process similar to ours (using PSID data from 1969 till 1987) and finds the variance of the transitory shock to be 0.17 with the first order auto-regressive persistence parameter being 0.35. This shows that the esti-

mates we are getting are close to those found in the literature. Moffitt and Gottschalk also report an increasing variance of earnings of white men over the 1970s and 1980s that can be attributed almost equally to rising variances of permanent and transitory components of earnings. Note that our paper is silent about whether or not the transitory and permanent components have become more variable across a person's lifetime, but that the permanent component appears to have become less variable for the children's generation.

Once we have compared the parameter estimates and their implications from the two different data specifications (panel and time-averaged cross-sectional), we will now see how both of them perform in terms of matching the empirical moments. Crucially, what helps our variance estimates to be positive while fitting the moments so well, is the presence of the negative covariance term between assets and permanent income — the justification of which has already been provided while discussing the properties of the variable $a_{i,t}$.

Table 2: Actual and Fitted Moments

Moments	Actual	Fitted
$Var\left(y_{i,t}^P\right)$	0.2676	0.2675
$Var\left(y_{i,t}^C\right)$	0.3299	0.3298
$Var\left(c_{i,t}^P\right)$	0.0595	0.0588
$Var\left(c_{i,t}^C\right)$	0.0770	0.0769
$Cov\left(y_{i,t}^P, y_{i,t}^C\right)$	0.0633	0.0634
$Cov\left(y_{i,t}^P, c_{i,t}^C\right)$	0.0292	0.0290
$Cov\left(y_{i,t}^P, c_{i,t}^P\right)$	0.0633	0.0708
$Cov\left(y_{i,t}^C, c_{i,t}^C\right)$	0.0763	0.0762
$Cov\left(y_{i,t}^C, c_{i,t}^P\right)$	0.0262	0.0245
$Cov\left(c_{i,t}^P, c_{i,t}^C\right)$	0.0200	0.0218

Note: All calculations are based on data that do not control for educational attainment of the individuals. The data with education controls give better performance of the fitted moments with the actual values.

4.1 Choice of Model Specification

One must notice that in our model specification the incomes of two successive generations are linked through the levels of their permanent incomes, and not through the shocks to those permanent incomes. In other words, the shocks to permanent income of a parent do not get transmitted to the child. Also due to identification issues in such a model with intergenerational transmission through the levels of variables instead of shocks to those variables, we cannot separate out permanent and transitory shocks to income — instead we can at most only have a persistent shock. Since there is no a priori reason to believe in such specification, we try, as an alternative specification, to model the income and consumption processes in a way such that we can separately identify permanent and transitory shocks to income, and also the permanent shocks get imperfectly transferred to the child's income process. This model is a direct extension of the one in Blundell et al. (2008) in the sense that we just replicate their model separately for the parents and the children, but linking them through the income processes. Details of the model are presented in Appendix F. Nevertheless, the GMM estimates from this alternative model suggest that the intergenerational persistence of the permanent income shocks are statistically insignificant.

5 Counterfactual Analysis

Once we have estimated the parameters of our statistical model, we are in a position to simulate data using those estimated parameters of the data-generating process. This simulated data can then be used to analyze the counterfactual question as to what amounts of different favourable shocks does a child hailing from a family of *poor* parents require in order to become *rich* himself. For answering this question in a tractable manner, we simulate the earnings, consumption, and the different shocks for both parents and children, and place them in four quartiles for each generation. Then we calculate the average value for the variables and the shocks for each of these four quartiles, and calculate the difference of the average value of each quartile from the highest quartile. These differences inform us about the amount of favourable shock of a particular type that is required to equate the earnings and consumption of the average child from a certain quartile to those

of his counterpart in the topmost quartile. Based off the GMM estimates¹⁷ above, we can make the following statements:

(i) If a child with earnings equal to the average earnings of the first quartile of the earnings distribution requires 100 units of favourable shocks to move to the average earnings of the top quartile then about 26.51 units come from the parental factor (higher parental permanent earnings), 65.21 units come from the idiosyncratic permanent earnings component of the child, and 8.28 units come from the per-period shock to the earnings process.

(ii) If a child with consumption equal to the average consumption of the first quartile of the consumption distribution requires 100 units of favourable shocks to move to the average consumption of the top quartile then about 12.15 units come from the parental factors (higher parental permanent earnings and higher per-period parental 'asset'), 87.41 units come from the idiosyncratic permanent earnings component of the child, and 0.44 units come from the per-period shocks to the earnings and 'asset' processes.¹⁸

Having evaluated the importance of parental factors in the residual earnings and consumption data, our next agenda is to arrive at the total parental influence — both in the observables that we have controlled for, and the residual data. To achieve that we need two things — first, the intergenerational persistence in each of the observable variables, and second, some weighting scheme for these variables to ascertain the total persistence in observables.

The persistence in each observable variable is calculated by averaging a dummy variable which took the value one when parents and children had the same value for the variable. For example, if a given year both a parent and his adult child are married then the dummy variable will be equal to unity, otherwise if any one is unmarried then the dummy variable is zero. basically, we are taking the average number of parent-child pairs having the same value of an observable as the mean persistence level in that observable.

The contribution of the observables are calculated using the R^2 contribution of each of

¹⁷We use the estimates that used panel data and did not control for education.

¹⁸Note that the reduction in the importance of the per-period shocks in explaining the cross-sectional variation in child consumption is consistent with the standard permanent income hypothesis.

them in the regression of earnings and consumption on those factors. These weights are then re-weighted by the R^2 contribution of all observables. However, the contributions of the residuals are calculated using the relative importance of each factor in explaining the difference between the mean earnings and consumption of the bottom and top quartiles (as discussed above). These residual weights are again weighted by the residual sum of squares of the controlling regression. Such reweighting is required so that the total observed variance of earnings and consumption adds up.

Table 3: Explanatory Power of Observable Characteristics and Residual Shocks

Factors Explaining Total Variation		Earnings	Consumption
Observables	Cohort and Year Fixed Effects	3.26%	9.05%
	Family Size Dummies	2.15%	8.85%
	Marital Status Dummies	6.48%	1.98%
	State of Residence Dummies	2.59%	12.72%
	Dummies for Number of Children	0.67%	9.69%
	Employment Status Dummies	10.53%	7.41%
	Race Dummy	1.54%	3.03%
Residuals	Parental Factor	19.29%	5.74%
	Permanent Child Factor	47.45%	41.32%
	Shock to Child	6.03%	0.21%

We calculate the average sample persistence of the following observables: family size, marital size, state of residence, number of children, employment status and race. We then calculate the total persistence in the observables by weighting the individual persistences by the contribution of the corresponding variable in the total variance (as calculated in Table 3 above). Then we add this persistence in observables to the contribution of the residual parental factor to get a measure of the total importance of parental factors (both observed and unobserved) in explaining the total variance of earnings and consumption. We find the numbers to be 29.22 percent for consumption and 33.39 percent for earnings.

6 Conclusion

Intergenerational mobility in the United States, both in terms of total family income and labour earnings of the household head, has decreased between 1977 and 2010, with most

of the reduction occurring before 1985 and after the early 2000s. The period between 1985 and 2000 is, instead characterized by a fairly stable intergenerational elasticity of both income and earnings. The intergenerational mobility in consumption appears to have changed proportionally even less, remaining fairly stable with the exception of the last 5 years of the sample period.

Crucially, our new reduced-form evidence highlights how the intergenerational elasticity of consumption closely tracks its (household) income counterpart, and is generally higher than the intergenerational elasticity of earnings (a cross-generation elasticity around 0.45 for consumption and total income versus 0.3 for earnings, on average).

These facts motivate our model-based estimation exercise, in which we posit that adult children's consumption choices are a function of their earnings and unobserved (asset) income. Both observed and unobserved components of income are linked across-generations (parent-to-child) and are subject to both transitory and permanent shocks that introduce individual-specific deviations from the predetermined family component.

After deriving the optimal consumption choices associated to different realization of income, we obtain a set of variance-covariance restrictions that describe the co-movement of earnings (income) and consumption. We exploit these moment restrictions to estimate the cross-generational law of motions of earnings (income) and consumption, and to perform a counterfactual decomposition of observed inequality among adult children. We find that parental factors (both observable and unobservable) explain roughly 29 percent of the consumption variance and 33 percent of the earnings' variance among adult children. Among the observable characteristics, parental education accounts for a large part (more than one-third) of the intergenerational elasticity in permanent (earnings) income.

References

- [1] Aaronson, D. and Mazumder, B., "Intergenerational Economic Mobility in the United States, 1940 to 2000", *The Journal of Human Resources*, XLIII(1):139-172, 2008.
- [2] Altonji, J.G., Hayashi, F. and Kotlikoff, L.J., "Parental Altruism and Inter Vivos Transfers: Theory and Evidence", *Journal of Political Economy*, 105(6):1121-1166, 1997.


- [3] Altonji, J.G. and Segal, L., "Small Sample Bias in GMM Estimation of Covariance Structures", *Journal of Business and Economic Statistics*, 14(3):353-366, 1996.
- [4] Andreski, P., Li, G., Samancioglu, M.Z. and Schoeni, R., "Estimates of Annual Consumption Expenditures and Its Major Components in the PSID in Comparison to the CE", *The American Economic Review*, 104(5):132-135, 2014.
- [5] Attanasio, O. and Pistaferri, L., "Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure", *The American Economic Review*, 104(5):122-126, 2014.
- [6] Blundell, R., Pistaferri, L. and Preston, I., "Consumption Inequality and Partial Insurance", *The American Economic Review*, 98(5):1887-1921, 2008.
- [7] Bound, J., Brown, C., Duncan, G.J. and Rodgers, W.L., "Evidence on the Validity of Cross-sectional and Longitudinal Labor Market Data", *Journal of Labor Economics*, 12(3):345-368, 1994.
- [8] Brown, Charles, "Notes on the 'SEO' or 'Census' Component of the PSID", 1996. <http://psidonline.isr.umich.edu/Publications/Papers/SEO.pdf>
- [9] Caballé, Jordi, "Intergenerational mobility: measurement and the role of borrowing constraints and inherited tastes", *SERIEs Journal of the Spanish Economic Association*, doi: 10.1007/s13209-016-0149-2, 2016.
- [10] Charles, K. K., Danziger, S., Li, G. and Schoeni, R., "The Intergenerational Correlation of Consumption Expenditures", *The American Economic Review*, 104(5):136-140, 2014.
- [11] Charles, K. K. and Hurst, E., "The Correlation Wealth Across Generations", *Journal of Political Economy*, 111(6):1155-1182, 2003.
- [12] Cox, D., Ng, S. and Waldkirch, A., "Intergenerational Linkages in Consumption Behavior", *The Journal of Human Resources*, XXXIX(2):355-381, 2004.
- [13] Flavin, M. and Yamashita, T., "Owner-Occupied Housing and the Composition of the Household Portfolio", *The American Economic Review*, 92(1):345-362, 2002.

- [14] Haider, S. and Solon, G., "Life-Cycle Variation in the Association between Current and Lifetime Earnings", *The American Economic Review*, 96(4):1308-1320, 2006.
- [15] Heathcote, J., Perri, F. and Violante, G.L., "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006", *Review of Economic Dynamics*, 13(1):15-51, 2010.
- [16] Heckman, J. and Landerso, R., "The Scandinavian Fantasy: The Sources of Intergenerational Mobility in Denmark and the U.S.", *NBER Working Paper 22465*, 2016.
- [17] Hertz, T., "Trends in the Intergenerational Elasticity of Family Income in the United States", *Industrial Relations*, 46(1):22-50, 2007.
- [18] Hertz, T., "Understanding Mobility in America", *Center for American Progress*, April 2006.
- [19] Lee, C. and Solon, G., "Trends in Intergenerational Income Mobility", *The Review of Economics and Statistics*, 2009.
- [20] Lefgren, L., Lindquist, M. J. and Sims, D., "Rich Dad, Smart Dad: Decomposing the Intergenerational Transmission of Income", *Journal of Political Economy*, 120(2):268-303, 2012.
- [21] Li, G., Schoeni, R.F., Danziger, S. and Charles, K.K., "New Expenditure Data in the PSID: Comparisons with the CE", *Monthly Labor Review*, 133(2):29-39, 2010.
- [22] Lopoo, L. and Mayer, S., "Has the Intergenerational Transmission of Economic Status Changed?", *Journal of Human Resources*, 40(1):169-185, 2005.
- [23] Moffitt, R.A. and Gottschalk, P., "Trends in the Covariance Structure of earnings in the U.S.: 1969-1987", *Journal of Economic Inequality*, 9(3): 439-459, 2011.
- [24] Peters, H. Elizabeth, "Patterns of Intergenerational Mobility in Income and Earnings", *The Review of Economics and Statistics*, 74(3):456-466, 1992.
- [25] Saez, Emmanuel, "Income and Wealth Inequality: Evidence and Policy Implications", *Neubauer Collegium Lecture, University of Chicago*, October 2014.

- [26] Shorrocks, A.F., "The Measurement of Mobility", *Econometrica*, 46(5):1013-1024, 1978.
- [27] Solon, Gary, "Intergenerational Income Mobility in the United States", *The American Economic Review*, 82(3):393-408, 1992.
- [28] Zimmerman, David, "Regression toward Mediocrity in Economic Stature", *The American Economic Review*, 82(3):409-29, 1992.

Appendix A

Figure A1. Goodness of Prediction of the Imputation Exercise



Note: In Panels A and B, the imputed series are normalized to equal the actual series in 1998 (the year from which both series are observed) to ease comparison. In Panel C, the normalization is done to the 2004 value.

Appendix B

Table B1: Estimates of Intergenerational Elasticities by Year (1977-2010)

Year	β_{ee}	β_{yy}	β_{cc}	β_{ce}	β_{cy}
1977	-	0.26	0.37***	0.24***	0.35***
1978	0.15	0.34*	0.45***	0.16***	0.34***
1979	0.17	0.28**	0.45***	0.16**	0.33***
1980	0.26**	0.39***	0.38**	0.14**	0.26***
1981	0.36***	0.43***	0.48***	0.17***	0.31***
1982	0.17**	0.44***	0.46***	0.11	0.24***
1983	0.27**	0.41***	0.41***	0.12*	0.25***
1984	0.19**	0.43***	0.41***	0.14**	0.29***
1985	0.19**	0.38***	0.42***	0.12*	0.24***
1986	0.32***	0.44***	0.45***	0.16***	0.28***
1987	0.24***	0.38***	-	-	-
1988	0.29***	0.41***	-	-	-
1989	0.27***	0.45***	0.45***	0.18***	0.29***
1990	0.30***	0.40***	0.45***	0.17***	0.25***
1991	0.33***	0.45***	0.43***	0.17***	0.25***
1992	0.31***	0.42***	0.47***	0.19***	0.29***
1993	0.29***	0.45***	0.44***	0.18***	0.26***
1994	0.34***	0.41***	0.46***	0.17***	0.29***
1995	0.32***	0.40***	0.45***	0.17***	0.27***
1996	0.26***	0.40***	0.44***	0.16***	0.25***
1998	0.29***	0.38***	0.44***	0.15***	0.23***
2000	0.32***	0.41***	0.45***	0.17***	0.26***
2002	0.30***	0.42***	0.49***	0.17***	0.26***
2004	0.29***	0.51***	0.50***	0.16***	0.23***
2006	0.31***	0.49***	0.52***	0.18***	0.26***
2008	0.36***	0.54***	0.52***	0.18***	0.24***
2010	0.39***	0.46***	0.53***	0.17***	0.23***

Note: ***, ** and * denote statistical significance at 1%, 5% and 10% levels respectively. Standard errors (not reported here to improve readability) are clustered at the level of the unique parent identity number. Missing estimate for β_{ee} in 1977 is due to dropping of explanatory variable due to multicollinearity; the other missing estimates are due to absence of consumption data for 1987 and 1988.

Appendix C

C.1 CRRA utility function

In the main article, we assumed a quadratic utility function for all individuals, and such utility function can be motivated as a second order approximation of preferences around the equilibrium. Alternatively, we can assume a more general utility function, and still arrive at the same log-consumption equations as equations (4.5) and (4.6) after some approximation of the Euler equation.

In the CRRA case, the Euler equation is given by $C_{i,t}^{-\tau} = \beta(1+r) E_t \left(C_{i,t+1}^{-\tau} \right)$, where $\tau > 0$ is the parameter capturing the degree of risk aversion as also the intertemporal elasticity of substitution. Maintaining the assumption $\beta(1+r) = 1$, we get from the Euler equation $E_t \left[\left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{-\tau} \right] = 1$. We define the function $h(g_c) = (1+g_c)^{-\tau}$, where $g_c = \frac{C_{i,t+1}}{C_{i,t}} - 1$ such that $E_t[h(g_c)] = 1$. A first order Taylor series expansion of $h(g_c)$ around $g_c^* = 0$ yields $h(g_c) \approx 1 - \tau g_c$. Taking expectations on both sides of this approximate equation, we get $E_t(g_c) = 0$, implying $C_{i,t} = E_t(C_{i,t+1})$. This is exactly same as the Euler equation that one has from quadratic utility function without any approximation. Now, since we did not derive explicitly the consumption expression from this Euler equation in the paper, we provide the derivation here. Iterating forward the per-period budget constraint $A_{i,t+1} = (1+r)(A_{i,t} + Y_{i,t} - C_{i,t})$ by one period and combining it with the Euler equation $C_{i,t} = E_t(C_{i,t+1})$, we get,

$$\begin{aligned}
 \left(1 + \frac{1}{1+r}\right) C_{i,t} &= A_{i,t} - \left(\frac{1}{1+r}\right)^2 A_{i,t+2} + \left[Y_{i,t} + \frac{1}{1+r} Y_{i,t+1}\right] \\
 &\vdots \\
 \Rightarrow \left[1 + \frac{1}{1+r} + \left(\frac{1}{1+r}\right)^2 + \dots \infty\right] C_{i,t} &= A_{i,t} - \lim_{k \rightarrow \infty} \left(\frac{1}{1+r}\right)^k A_{i,t+k} + \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j E_t(Y_{i,t+j}) \\
 \Rightarrow \left[\frac{1+r}{r}\right] C_{i,t} &= A_{i,t} + \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j E_t(Y_{i,t+j}) \\
 \Rightarrow C_{i,t} &= \frac{r}{1+r} \left[A_{i,t} + \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j E_t(Y_{i,t+j}) \right]
 \end{aligned}$$

Note that in the above derivation we have assumed the No-Ponzi condition that prevents an individual from continuously borrowing and rolling over his debt to future periods, $\lim_{k \rightarrow \infty} \left(\frac{1}{1+r} \right)^k A_{i,t+k} = 0$.

C.2 Assumption of infinitely lived agents

In the paper we assume that the individuals are infinitely lived. This implies, while optimally choosing consumption, the agents do not take into account the fact that they die, because any asset that is left after their demise will be inherited by their children, and they are as much concerned about the consumption of their children as they are about their own. Another way to explain the assumption of infinitely lived agents is that when the individuals choose their consumption (at time t), they are assumed to be very far from their demise (at time T), such that $(T - t)$ is large enough. This assumption is not altogether without justification because we take the income and consumption data for individuals only till 65 years of age (the approximate age for superannuation), while the age of demise of the individual can be assumed to be sufficiently higher than 65 years. Thus, the choice of the current time point t for any individual is maximum at age 65, which is sufficiently distant from the expected time T of demise, say 80 years or more.

Having discussed the validity of the assumption of infinitely lived agents, the crucial question that remains is what do we gain in terms of our estimation strategy by making this assumption. To see that we derive the expression for log consumption in the finite horizon case, assuming that the agent dies at the end of period T with zero asset position ($A_{i,T+1} = 0$):¹⁹

$$c_{i,t} = \begin{cases} \frac{1}{1-(1+r)^{-(T-t)-1}} a_{i,t} + \bar{y}_i + f(\alpha; r) \frac{(1+r)^{-\alpha(T-t)+1} (1+r)^{-(T-t)}}{(1+r)^{-(1+r)^{-(T-t)}}} \eta_{i,t} & \text{for } r > 0 \\ \frac{1+\bar{a}_{i,t}}{T-t+1} + \bar{y}_i + \frac{1-\alpha^{T-t+1}}{(1-\alpha)(T-t+1)} \eta_{i,t} & \text{for } r = 0 \end{cases}$$

where symbols have the same definitions as in the main text and α is the AR(1) persistence parameter in the transitory shock process. In our baseline model we assume $\alpha = 0$.

Clearly, the presence of the asset term in the consumption expression for $r = 0$ defeats

¹⁹This assumption is the finite horizon counterpart of the No-Ponzi condition $\lim_{k \rightarrow \infty} \left(\frac{1}{1+r} \right)^k A_{i,t+k} = 0$ in the infinite horizon case.

the purpose of the assumption $r = 0$ for explicitly incorporating measurement error because we cannot get rid of the three parameters $(\sigma_{a^P}^2, \sigma_{a^P, \bar{y}_i^P}$ and $\sigma_{\psi^C}^2)$ like in the infinite horizon case. So, the infinite horizon assumption is essentially crucial for the identification in our measurement error model.

As for the usual $r > 0$ case, we can re-define $a_{i,t}$ as $\frac{1}{1-(1+r)^{-(T-t)-1}} a_{i,t}$ because we do not actually observe $a_{i,t}$ in the data. However, the only difference in estimation would arise from the presence of the extra multiplicative term $\frac{(1+r)-\alpha^{(T-t)+1}(1+r)^{-(T-t)}}{(1+r)-(1+r)^{-(T-t)}}$ with $[f(\alpha; r)] \eta_{i,t}$. Even with a finite $(T - t)$, we can ignore this multiplicative term so long α is close enough to unity, that is, the AR(1) shock to permanent income is very persistent.

Appendix D

We present here the full version of the model with AR(1) persistence in the transitory shocks, where α^P and α^C denote the intertemporal persistence parameters. Also, $v_{i,t}^P$ and $v_{i,t}^C$ can be thought as the innovations to the AR(1) processes because note that in the absence of the autoregressive persistence the innovation is nothing but the i.i.d. transitory shock in our baseline model. The identification argument can be simplified for our baseline model without such persistence. Also, the variances of the first differences of earnings (as shown in equations D.3 and D.4 below) are dropped for the baseline model.

Moment conditions

$$Var(y_{i,t}^P) = \sigma_{\bar{y}_i^P}^2 + \frac{\sigma_{v^P}^2}{1 - (\alpha^P)^2} \quad (D.1)$$

$$Var(y_{i,t}^C) = \rho^2 \sigma_{\bar{y}_i^P}^2 + \sigma_{\epsilon^C}^2 + \frac{\sigma_{v^C}^2}{1 - (\alpha^C)^2} \quad (D.2)$$

$$Var(\Delta y_{i,t}^P) = \frac{2\sigma_{v^P}^2}{1 + \alpha^P} \quad (D.3)$$

$$Var(\Delta y_{i,t}^C) = \frac{2\sigma_{v^C}^2}{1 + \alpha^C} \quad (D.4)$$

$$Var(c_{i,t}^P) = \sigma_{a^P}^2 + \sigma_{\bar{y}_i^P}^2 + 2\sigma_{a^P, \bar{y}_i^P} + \left[f(\alpha^P; r) \right]^2 \frac{\sigma_{v^P}^2}{1 - (\alpha^P)^2} \quad (D.5)$$

$$Var\left(c_{i,t}^C\right) = \phi^2\sigma_{a^P}^2 + \rho^2\sigma_{\bar{y}_i^P}^2 + 2\rho\phi\sigma_{a^P,\bar{y}_i^P} + \sigma_{\epsilon^C}^2 + \sigma_{\psi^C}^2 + \left[f\left(\alpha^C;r\right)\right]^2 \frac{\sigma_{v^C}^2}{1 - (\alpha^C)^2} \quad (D.6)$$

$$Cov\left(y_{i,t}^P, y_{i,t}^C\right) = \rho\sigma_{\bar{y}_i^P}^2 \quad (D.7)$$

$$Cov\left(y_{i,t}^P, c_{i,t}^C\right) = \rho\sigma_{\bar{y}_i^P}^2 + \phi\sigma_{a^P,\bar{y}_i^P} \quad (D.8)$$

$$Cov\left(y_{i,t}^P, c_{i,t}^P\right) = \sigma_{\bar{y}_i^P}^2 + \sigma_{a^P,\bar{y}_i^P} + \left[f\left(\alpha^P;r\right)\right] \frac{\sigma_{v^P}^2}{1 - (\alpha^P)^2} \quad (D.9)$$

$$Cov\left(y_{i,t}^C, c_{i,t}^C\right) = \rho^2\sigma_{\bar{y}_i^P}^2 + \rho\phi\sigma_{a^P,\bar{y}_i^P} + \sigma_{\epsilon^C}^2 + \left[f\left(\alpha^C;r\right)\right] \frac{\sigma_{v^C}^2}{1 - (\alpha^C)^2} \quad (D.10)$$

$$Cov\left(y_{i,t}^C, c_{i,t}^P\right) = \rho\sigma_{\bar{y}_i^P}^2 + \rho\sigma_{a^P,\bar{y}_i^P} \quad (D.11)$$

$$Cov\left(c_{i,t}^P, c_{i,t}^C\right) = \phi\sigma_{a^P}^2 + \rho\sigma_{\bar{y}_i^P}^2 + (\phi + \rho)\sigma_{a^P,\bar{y}_i^P} \quad (D.12)$$

Identification

Subtract equation (D.7) from equations (D.8) and (D.11) to get two differences, say d_1 and d_2 respectively. Then, $\phi = \frac{d_1}{d_2}\rho$.

From equation (D.12),

$$\begin{aligned} \phi\sigma_{a^P}^2 &= Cov\left(c_{i,t}^P, c_{i,t}^C\right) - \left[d_1 + d_2 + Cov\left(y_{i,t}^P, y_{i,t}^C\right)\right] \equiv N_1 \\ \implies \rho\sigma_{a^P}^2 &= \frac{d_2}{d_1}N_1 \equiv N_2 \end{aligned}$$

From equation (D.5),

$$Var\left(c_{i,t}^P\right) = \frac{1}{\rho} \left[N_2 + 2d_2 + Cov\left(y_{i,t}^P, y_{i,t}^C\right) \right] + \left[f\left(\alpha^P;r\right) \right]^2 \frac{\sigma_{v^P}^2}{1 - (\alpha^P)^2} \quad (D.13)$$

Combining equations (D.1) and (D.7) we have,

$$Var\left(y_{i,t}^P\right) = \frac{1}{\rho} Cov\left(y_{i,t}^P, y_{i,t}^C\right) + \frac{\sigma_{v^P}^2}{1 - (\alpha^P)^2} \quad (D.14)$$

Equations (D.3), (D.13) and (D.14) constitute a system of three non-linear equations in three unknowns, namely, ρ , α^P and $\sigma_{v^P}^2$, and hence these three parameters are all identified. This consequently identifies $\sigma_{\bar{y}_i^P}^2$ from equation (D.7), $\sigma_{a^P, \bar{y}_i^P}$ from equation (D.11), ϕ from equation (D.8) and $\sigma_{a^P}^2$ from equation (D.12). Now, equations (D.2), (D.10) and (D.4) can be simultaneously solved to identify α^C , $\sigma_{\epsilon^C}^2$ and $\sigma_{v^C}^2$. Finally, $\sigma_{\psi^C}^2$ can be identified from equation (D.6). Note that equation (D.9) is not being used in the identification argument above, implying that the system of equations (D.1) through (D.12) over-identify the set of eleven parameters.

Appendix E

Statistical Model of Heathcote, Perri, Violante (2010)

Heathcote, Perri and Violante (2010) (henceforth HPV) analyzes the time series of residual (after controlling for age and education in a standard Mincerian wage equation) log hourly wage. We are replicating their methodology for residual labour earnings instead of wages. Let $y_{i,c,t}$ be the residual log earnings for individual i of cohort c at date t that follows a permanent-transitory model of the form

$$y_{i,c,t} = z_{i,c,t} + v_{i,c,t} \quad (\text{E.1})$$

$$z_{i,c,t} = z_{i,c,t-1} + \eta_{i,c,t} \quad (\text{E.2})$$

where $v_{i,c,t}$ and $\eta_{i,c,t}$ are innovations that are independent across individuals and over time, identically distributed across individuals, and orthogonal to each other with time-varying variances but constant-across-cohorts variances $\sigma_{v,t}^2$ and $\sigma_{\eta,t}^2$. HPV then estimates these variance parameters in two different methods: first using the second moments of the data in levels, and then in time-differences.

In levels, the within-cohort moment restrictions used are as follows:

$$\text{var}_c(y_{i,c,t}) - \text{cov}_c(y_{i,c,t+2}, y_{i,c,t}) = \sigma_{v,t}^2 \quad (\text{E.3})$$

$$\text{var}_c(y_{i,c,t}) - \text{cov}_c(y_{i,c,t-2}, y_{i,c,t}) = \sigma_{v,t}^2 + \sigma_{\eta,t}^2 + \sigma_{\eta,t-1}^2 \quad (\text{E.4})$$

The moment conditions given by (E.3) identifies $\sigma_{v,t}^2$ for $t = 1967, \dots, 1994, 1996, 1998, \dots, 2008$. Then (E.4) identifies $(\sigma_{\eta,t}^2 + \sigma_{\eta,t-1}^2)$ for $t = 1969, \dots, 1994, 1996, 1998, \dots, 2008$. Keeping in mind that the PSID data has turned biennial from year 1998, HPV decided to take time-differences of the wage data at two-year intervals for the entire sample to maintain parity, $\Delta y_{i,c,t} \equiv y_{i,c,t} - y_{i,c,t-2}$. This implies the following moment conditions in differences:


$$\text{cov}_c (\Delta y_{i,c,t+2}, \Delta y_{i,c,t}) = -\sigma_{v,t}^2 \quad (\text{E.5})$$


$$\text{var}_c (\Delta y_{i,c,t}) = \sigma_{\eta,t}^2 + \sigma_{\eta,t-1}^2 + \sigma_{v,t}^2 + \sigma_{v,t-2}^2 \quad (\text{E.6})$$

The first set of moments (E.5) identifies $\sigma_{v,t}^2$ for $t = 1969, \dots, 1994, 1996, 1998, \dots, 2008$. Given those estimates, (E.6) identifies $(\sigma_{\eta,t}^2 + \sigma_{\eta,t-1}^2)$ for $t = 1971, \dots, 1994, 1996, 1998, \dots, 2008$.

HPV notes that under the assumption that the considered model is indeed the true one, none of the moments depend on cohort, c , thereby allowing estimating the variances at date t by averaging across all cohorts in the sample at t . Figure E1 below plots the time series of the variances of temporary and permanent shocks. Also, since the year-specific permanent shock variance cannot be identified, the plotted series is simply $\frac{\sigma_{\eta,t}^2 + \sigma_{\eta,t-1}^2}{2}$ for each date t .

Figure E1. Variance of Transitory and Permanent Shocks to Earnings





Appendix F

Disclaimer: This appendix should be read in disjunction with the rest of the paper as there are some conflicting notations. Also, this model is written with a more generalized notation structure to accommodate time-varying parameters.

Alternative Model

The Parent

At any year t , the parent i will have an income $Y_{i,t}^P$ such that the log income $y_{i,t}^P$ can be expressed as a sum of a permanent component, $p_{i,t}^P$ and a white noise process $u_{i,t}^P$, that is, $y_{i,t}^P = p_{i,t}^P + u_{i,t}^P$. The permanent component follows a random walk, $p_{i,t}^P = p_{i,t-1}^P + v_{i,t}^P$, with $v_{i,t}^P$ being a mean zero i.i.d. process. Now, taking the first difference of parental income, we get,

$$\Delta y_{i,t}^P = v_{i,t}^P + \Delta u_{i,t}^P \quad (\text{F.1})$$

The unexplained time-change in log consumption is defined as

$$\Delta c_{i,t}^P = \phi_t v_t^P + \psi_t u_{i,t}^P + \zeta_{i,t}^P \quad (\text{F.2})$$

where the loading factors ϕ_t and ψ_t are the partial insurance parameters. The three processes $v_{i,t}^P$, $u_{i,t}^P$ and $\zeta_{i,t}^P$ are assumed to be uncorrelated.

The Child

The child will have the same basic income process as the parents, except that their permanent income is going to be determined by their parental permanent income, that is, $y_{i,t}^C = p_{i,t}^C + u_{i,t}^C$, where once again $p_{i,t}^C = p_{i,t-1}^C + v_{i,t}^C$, with the additional structure that the permanent shock to the child's income is correlated with the parental permanent shock in the following way: $v_{i,t}^C = \rho v_{i,t}^P + \epsilon_{i,t}^C$. The parameter ρ captures the intergenerational elasticity in the shock to permanent income. Once again, the growth rate in consumption for the child is given as $\Delta c_{i,t}^C = \phi_t v_t^C + \psi_{i,t} u_{i,t}^C + \zeta_{i,t}^C$, where it is assumed $\zeta_{i,t}^C = \gamma_t \zeta_{i,t}^P + \chi_{i,t}^C$. Thus, γ captures the intergenerational intra-family linkage in consumption growth that is not explained through income shocks. Combining all these structures, we can write the income and consumption growth rates for the child as follows:

$$\Delta y_{i,t}^C = \rho_t v_{i,t}^P + \epsilon_{i,t}^C + \Delta u_{i,t}^C \quad (\text{F.3})$$

$$\Delta c_{i,t}^C = \phi_t \rho_t v_{i,t}^P + \phi_t \epsilon_{i,t}^C + \psi_{i,t} u_{i,t}^C + \gamma_t \zeta_{i,t}^P + \chi_{i,t}^C \quad (\text{F.4})$$

Moment Conditions

$$\text{Var} \left(\Delta y_t^P \right) = \sigma_{v_t^P}^2 + \sigma_{\Delta u_t^P}^2 \quad (\text{F.5})$$

$$\text{Var} \left(\Delta y_t^C \right) = \rho_t^2 \sigma_{v_t^P}^2 + \sigma_{\epsilon_t^C}^2 + \sigma_{\Delta u_t^C}^2 \quad (\text{F.6})$$

$$\text{Var} \left(\Delta c_t^P \right) = \phi_t^2 \sigma_{v_t^P}^2 + \psi_t^2 \sigma_{u_t^P}^2 + \sigma_{\zeta_t^P}^2 \quad (\text{F.7})$$

$$Var(\Delta c_t^C) = \phi_t^2 \rho_t^2 \sigma_{v_t^P}^2 + \psi_t^2 \sigma_{u_t^C}^2 + \phi_t^2 \sigma_{\epsilon_t^C}^2 + \gamma_t^2 \sigma_{\zeta_t^P}^2 + \sigma_{\chi_t^C}^2 \quad (F.8)$$

$$Cov(\Delta y_t^P, \Delta y_{t+1}^P) = -\sigma_{u_t^P}^2 \quad (F.9)$$

$$Cov(\Delta y_t^C, \Delta y_{t+1}^C) = -\sigma_{u_t^C}^2 \quad (F.10)$$

$$Cov(\Delta y_t^P, \Delta y_t^C) = \rho_t \sigma_{v_t^P}^2 \quad (F.11)$$

$$Cov(\Delta c_t^P, \Delta c_t^C) = \phi_t^2 \rho_t \sigma_{v_t^P}^2 + \gamma_t \sigma_{\zeta_t^P}^2 \quad (F.12)$$

$$Cov(\Delta c_t^P, \Delta y_t^P) = \phi_t \sigma_{v_t^P}^2 + \psi_t \sigma_{u_t^P}^2 \quad (F.13)$$

$$Cov(\Delta c_t^C, \Delta y_t^C) = \phi_t \rho_t^2 \sigma_{v_t^P}^2 + \phi_t \sigma_{\epsilon_t^C}^2 + \psi_t \sigma_{u_t^C}^2 \quad (F.14)$$

$$Cov(\Delta c_t^P, \Delta y_{t+1}^P) = -\psi_t \sigma_{u_t^P}^2 \quad (F.15)$$

$$Cov(\Delta c_t^C, \Delta y_{t+1}^C) = -\psi_t \sigma_{u_t^C}^2 \quad (F.16)$$

Identification

Equations (F.9) and (F.10) directly identify $\sigma_{u_t^P}^2$ and $\sigma_{u_t^C}^2$ respectively. This automatically pins down $\sigma_{\Delta u_t^P}^2$ and $\sigma_{\Delta u_t^C}^2$ in equations (F.5) and (F.6). This implies that $\sigma_{v_t^P}^2$ gets identified from equation (F.5). Consequently, ρ_t is identified from equation (F.11), $\sigma_{\epsilon_t^C}^2$ from (F.6), ψ_t from either equation (F.15) or (F.16), and ϕ_t from equation (F.13). This leaves $\sigma_{\zeta_t^P}^2$ to be identified from (F.7), γ_t from (F.12), and finally, $\sigma_{\chi_t^C}^2$ from equation (F.8). Thus, we have over-identification with ten parameters and twelve moment conditions.