

INSURANCE IN EXTENDED FAMILY NETWORKS

By

Orazio Attanasio, Costas Meghir, and Corina Mommaerts

March 2015

Revised June 2018

COWLES FOUNDATION DISCUSSION PAPER NO. 1996R



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Insurance in extended family networks

Orazio Attanasio (UCL, IFS and NBER)

Costas Meghir (Yale, IFS, IZA, CEPR and NBER)

Corina Mommaerts (University of Wisconsin - Madison)*

June 13, 2018

Abstract

We investigate partial insurance and group risk sharing in extended family networks. Our approach is based on decomposing income shocks into group aggregate and idiosyncratic components, allowing us to measure the extent to which each component is insured. We apply our framework to extended family networks in the United States by exploiting the unique intergenerational structure of the Panel Study of Income Dynamics. We find that over 60% of shocks to household income are potentially insurable within extended family networks. However, we find little evidence that the extended family provides insurance for such idiosyncratic shocks.

*We are grateful to four anonymous referees and the editor of the Journal for numerous constructive comments. We thank Joseph Altonji, Richard Blundell, Dirk Krueger, Ross Milton, Emily Nix, Alessandra Voena, and participants at the NBER Summer Institute 2013 Macro Public Finance Workshop, the ARIA 2014 Annual Meeting, the 2017 Risk Theory Society meeting, and the 2018 Arizona State University Empirical Microeconomics conference. We thank Luigi Pistaferri for generously providing imputed consumption data. Orazio Attanasio thanks the ESRC Professorial Fellowship ES/K010700/1 and the ESRC CPP (RES-544-28-0001) for financial support. Costas Meghir thanks the ISPS and the Cowles foundation at Yale for financial assistance. Corina Mommaerts thanks the NSF Graduate Research Fellowship for support. The usual disclaimers apply.

Orazio Attanasio: o.attanasio@ucl.ac.uk,

Costas Meghir: c.meghir@yale.edu,

Corina Mommaerts: cmommaerts@wisc.edu.

1 Introduction

Much research has been devoted to the intertemporal allocation of resources by households. The ability of individuals and households to absorb income and resource shocks has substantial implications for their welfare, and limits to this ability could constitute an important motivation for policy interventions. How much smoothing a household can achieve depends crucially on the instruments at their disposal and the markets they have access to. With complete markets, households have access to a full set of state-contingent assets that allow them to completely diversify idiosyncratic risk; at the other extreme of self-insurance, households can only save and borrow through a simple asset, which limits their smoothing capabilities such that certain income shocks are reflected into consumption changes. In reality, households participate in a variety of markets, have access to many sources of income and transfers, and interact with other households both formally and informally. However, they also face a number of barriers, such as incomplete information about shocks or imperfect enforceability of contracts, which prevent full risk sharing. It is thus unsurprising that most empirical tests soundly reject complete markets and also find smoothing beyond what is feasible through self-insurance.¹

The absence of complete markets does not preclude the existence of partial insurance among groups of households that may be in a position to alleviate frictions that prevent full insurance. The extent to which such groups provide insurance is an interesting question both because it documents the importance of a potential source of risk sharing and because it could point to the relevance of certain frictions. In this paper, we develop an empirical framework of partial insurance that builds risk sharing within small groups into a model of self insurance, and apply it to extended families in the United States. In particular, we seek to understand (a) the potential amount of insurance that the extended family can offer, and (b) the extent to which this insurance actually takes place. We focus on the extended family because some of the barriers that prevent full insurance, such as information and enforceability problems, may be less relevant among households that are more tightly connected. We take into account that extended families are small and less diversified than the overall economy, which may limit the scope for insurance among its members. Ultimately, comparing the smoothing provided within the extended family against the smoothing within the

¹For example, see Attanasio and Davis (1996) for the rejection of complete markets, and Campbell and Deaton (1989) and Blundell, Pistaferri, and Preston (2008) for excess smoothness.

economy at large can be informative about the mechanisms underlying the overall amount of consumption smoothing observed in the data.

Our framework is based on a joint model of the stochastic process of income and the corresponding behavior of consumption for an extended family group. Income evolves as a permanent-transitory process, as in previous studies (see MaCurdy (1983), Abowd and Card (1989), Meghir and Pistaferri (2004) and Attanasio and Borella (2014)). However, we extend this structure to distinguish between family-aggregate and purely idiosyncratic components of income. This distinction is meaningful for risk sharing, because the group can insure components of shocks that are idiosyncratic to its individual members but is not able to do so for components that are aggregate for the group, thus affecting all members.

As in Blundell, Pistaferri, and Preston (2008) we approximate consumption growth with a linear function of income shocks,² and we extend this relationship to allow for the distinction between idiosyncratic and group-aggregate income components. By considering the difference between uninsurable “family-aggregate” and potentially insurable idiosyncratic shocks, we can detect and quantify the amount of extended family insurance that takes place by estimating the extent to which each of these shocks are passed through to consumption. We can also distinguish extended families risk sharing from self-insurance and insurance from other outside sources.

To estimate our model and quantify the extent of insurance in extended-family networks, we require longitudinal data on income and consumption, as well as familial links between households. The main United States data that contains (most of) this information is the Panel Study of Income Dynamics (PSID), which has the unique feature of following the households of children of originally sampled households once they have split away from their families. In addition, it contains detailed income data, food consumption, and since 1999 much richer consumption information. We use the 1979-2011 waves, exploiting the long time series structure and observing extended families over a long period of time. So as to be able to use a more comprehensive consumption measure dating back before 1999, we follow Blundell, Pistaferri, and Preston (2008) by using consumption imputed from the 1980-2008 Consumer Expenditure Survey (CEX).

We document three main empirical findings. First, our decomposition shows that over 60% of income shocks are potentially insurable by extended-family risk sharing networks. Thus, although

²See Hall (1988), Hall and Mishkin (1982), and Blundell and Preston (1998), amongst others.

they are small in number compared to other natural groupings of households (e.g., coworkers, religious groups, or neighborhoods), extended families have a sizable potential to share risk among its households. Second, despite this potential, we find no evidence of any insurance within the family network. Why this is the case is harder to establish. It may well be that the institutional framework of public welfare programs and other sources of insurance is viewed as providing adequate insurance, or it may be that frictions, either in the form of stigma, transaction costs, or imperfect information and moral hazard, may hamper the ability of extended families to share risk. Our third set of findings investigates this further by exploring the possibility that the family responds more in settings in which shocks are easier to observe, including when households live near one another or when shocks are particularly large. We find that no evidence of the latter, but some suggestive evidence that extended families whose households live near one another provide some – though not full – insurance, possibly because they are better able to solve information problems.

An earlier literature also focused specifically on risk sharing within the extended family in the United States. In a series of papers, Altonji, Hayashi, and Kotlikoff (1992) and Hayashi, Altonji, and Kotlikoff (1996), consider whether extended families can be viewed as collective units sharing resources and risk efficiently and reject this hypothesis.³ Hayashi, Altonji, and Kotlikoff (1996) conclude that “Future research should be directed to estimating the extent of consumption insurance over and above self-insurance” (p. 288). We follow up on this suggestion and extend this work in a number of ways. By estimating separate household and extended family income processes, we are able to understand better the relative importance of extended family versus (idiosyncratic) household level shocks and therefore the extent to which extended family insurance is feasible. Moreover, we go beyond testing (and rejecting) full risk-sharing by (a) detecting and quantifying the *amount* of extended family insurance that takes place and (b) distinguishing it from self-insurance and insurance from other outside sources. Finally, we exploit broader measures of consumption than did the earlier papers, which relied exclusively on food consumption.

Our main finding of no extended family insurance is perhaps surprising in light of related work in the United States as well as developing country contexts. In the United States, for example, parents offer coresidence to their adult children as insurance against labor market risk (Rosenzweig and Wolpin, 1993; Kaplan, 2012) and save in a precautionary manner in response to the occupations

³Choi, McGarry, and Schoeni (2016) update these results using more recent data, and come to a similar conclusion.

of their children (Boar, 2018). In other parts of the world, extended family networks are often found to provide an important – yet incomplete – insurance role (e.g., Angelucci, De Giorgi, and Rasul (2017) in Mexico and Fafchamps and Lund (2003) in the Philippines), though not always: Kinnan and Townsend (2012), for example, find that kinship networks facilitate investment financing but not consumption smoothing.⁴ While we cannot pin down a definitive reason for our contrasting results, it may be the case that commitment issues or information frictions are more severe in the US, perhaps because the relatively easy option to move offers more outside options (Morten, 2017) or because extended families are more geographically dispersed and/or can rely on other mechanisms for insurance, such as public insurance programs that can crowd out private arrangements. Thus, our findings have implications for the modeling of extended families and are an important step towards the identification of the relevant frictions that prevent risk sharing beyond what exists *across* families.

More generally, our framework and empirical findings contribute to the broader literature on risk sharing and partial insurance. On the theory side, an important literature has developed that formally derives the conditions under which partial insurance can occur. Several models have been developed that characterize constrained efficient equilibria, where the constraints arise from several sources, such as the imperfect enforceability of contracts (such as in Thomas and Worrall (1988), Ligon, Thomas, and Worrall (2002), Kocherlakota (1996)) or imperfect information and moral hazard (such as Cole and Kocherlakota (2001), Golosov, Tsyvinski, and Werning (2007) and, more recently, Karaivanov and Townsend (2014)). Dubois, Jullien, and Magnac (2008) also consider how insurance possibilities vary with the presence of formal contracting, showing that they effectively expand the set of incentive compatible allocations by acting as collateral. An important step in this research agenda is to relate the amount of risk sharing and the deviations from perfect risk sharing observed in the data to the implications of specific models of frictions. Attanasio and Pavoni (2011), for instance, map the parameters that measure how much of income shocks are reflected into consumption changes to the severity of a moral hazard friction. Our paper explores an alternative context where many of these frictions could be mitigated and through which partial

⁴A parallel literature on the economic arrangements within *nuclear* families typically finds that spouses can sustain some – but again, not full – insurance (see Chiappori (1988) and follow-up work in the United States and Thomas (1990) for an example in a developing country context). This paper complements this literature in that it examines a broader notion of families (the *extended* family) and finds that models of nuclear families do not describe well the behavior of extended families.

insurance can arise, namely the extended family.

The paper proceeds in Section 2 with a discussion of alternative approaches using data on direct transfers. Section 3 presents our model of partial insurance and Section 4 discusses identification using covariance moments. We describe the PSID and CEX data and general method of moments estimation procedure in Section 5, and Section 6 reports our main results. Section 7 provides several robustness checks, including various subsample analyses and modeling extensions, and Section 8 concludes.

2 Evidence of extended family insurance using direct transfers

We begin our analysis by showing patterns of direct transfers between households of various relationships, including parents, children, siblings, other relatives, and non-relatives. The tests we propose in the next section do not use direct information on transfers and focus, instead, on the relationship between the distribution of consumption and income. Nevertheless, information on direct transfers is useful to assess the importance of these informal mechanisms for risk sharing. Beyond showing that transfers are most commonly exchanged between extended family members, we view the extended family as a natural grouping with which to share risk because of the presence of altruism, the frequency of contact and the relative proximity between family members.⁵ Of course, other relationships (such as friends) could exhibit some of these features, but the combination of these characteristics and the patterns in the transfer data suggest that risk-sharing among extended family members may be first-order.

Indeed, we are not the first to investigate the insurance relationship between extended family members. One strand of literature models specific in-kind transfers, such as those of goods, housing (i.e. shared residence) or time help. Kaplan (2012) and Rosenzweig and Wolpin (1993), for instance, model the decision of adult children to co-reside with their parents as insurance against income risk, and find it to be an important source of insurance. Transfers of time in the form of babysitting or caregiving may also be an important source of insurance: Blau and Currie (2006), for

⁵For instance, during our sample period, two-thirds of adult children spoke on the phone with their mother at least once a week, and three-fourths live within 100 miles of their father among those who do not already reside in the same household (from the 1986 General Social Survey and 1988 PSID). In addition, in the early 1970s, 45% of PSID respondents reported that a relative lived within walking distance, and today the median American lives only 18 miles from their mother (from the 1968-1972 PSID and the 2008 HRS).

example, find that three-fourths of child care provided to working mothers by relatives is unpaid. Another literature looks directly at cash transfers. Edwards (2015) finds that an unemployment spell increases the likelihood of receiving a cash transfer from a family member in the PSID. McGarry (2016) uses 17 years of data from the Health and Retirement Study to examine the dynamic aspects of transfer behavior from parents to children. She finds that around 12-15% of children receive a transfer greater than \$500 from their parents in any given year, and that the probability of receiving a transfer correlates strongly with changes in a child's income.

We can perform an analysis similar to that in McGarry (2016) using our sample from the PSID (see Section 5.1 for a description of the main data and sample selection). In 1988 and 2013, the PSID collected supplementary data on monetary and time transfers between parents and their children. Using transfer data from 2013, Table 1 presents annual monetary and time transfers given from parents to children in the top panel and the transfers received from children in bottom panel. From the top panel, we see that 37 percent of adult children received transfers in the form of time (around 270 hours a year, on average) and 32 percent received monetary transfers (around \$3400 a year on average) in the previous year. In the other direction, only 20 percent of parents received time transfers (around 170 hours a year) and 7 percent receive money transfers (\$380 a year).

Table 2 presents a similar analysis using supplementary data from 1988 with an expanded universe of transfer recipients (parents, children, siblings, other relatives, and non-relatives) and again finds that most transfers are between parents and children, which underlies our motivation to consider the extended family as a possible insurance network.

Additionally, we are able to link the 1988 transfer data to income changes between 1987 and 1988. As described in more detail in Section 3.1, we isolate *unexplained* income, defined as the residual in a regression of log (pre-transfer) income on a set of demographics, to capture unexpected income changes. In Table 3, we run probit regressions of the probability of receiving a transfer on the quartile of a household's unexplained income change from 1987 to 1988. By splitting the sample into quartiles of the income "shock" distribution, we get a rough sense for whether households that receive the most negative income shocks (those in the first quartile) are more likely to receive transfers from family and friends relative to those who receive the most positive income shocks (those in the fourth quartile). Panel A reports marginal effects on the probability of receiving a

Table 1: Parent-child transfers (2013 PSID)

	Any amount (%)	Mean (conditional)	25th percentile (conditional)	Median (conditional)	75th percentile (conditional)
<i>Transfers to children</i>					
Hours	0.368	272	20	60	240
Money	0.321	3390	500	1100	3500
<i>Transfers from children</i>					
Hours	0.204	168	15	50	208
Money	0.068	382	100	300	500

Note: The sample includes extended families in the 2013 PSID, where we define extended families as a cohabiting couple under 65 and their adult children over 25. The data come from parent reports of transfers in 2012. “Transfers Given” refer to transfers given from parents to children, and “Transfers Received” refer to transfers received by parents from children. Only monetary transfers over \$100 are ascertained. Column (1) reports the percent of parent households reporting non-zero amounts of time or money, and columns (2)-(5) report amounts of time or money, conditional on a non-zero amount. Column (2) reports the conditional mean, and columns (3)-(5) report the conditional 25th percentile, median, and 75th percentile of the conditional distribution.

monetary transfer. From column 1, we see that households in the bottom quartile of income shocks are 6 percent more likely to receive a transfer than households in the top quartile of income shocks, suggesting that transfers may play an insurance role. The subsequent columns present estimates for different subsamples and suggest that most of this effect is driven by transfers from parents to children. Panel B repeats this analysis for time transfers and shows that time transfers are not significantly correlated with unexplained income changes.

The simple correlations presented in this section, therefore, suggest that: (1) parents and adult children may be the appropriate risk sharing network, (2) monetary transfers from parents to children are associated with income changes and (3) time transfers do not appear to be related to income changes. One limitation of these results, however, is that they do not take into account the circumstances of the potential provider of transfers. For instance, one would expect the transfer behavior of a potential donor who themselves had a large negative shock to be very different from the transfer behavior of a potential donor who received a large positive shock. In addition, these transfer results, which ignore the variation in consumption and its relation to income shocks, do not capture the *degree* to which households are insured against such shocks, both because the magnitudes of the effects in Table 3 are not interpretable in a risk sharing framework and because it is unclear how these transfers interact with other sources of insurance available to households. In the next section, we discuss a model of partial insurance and extended family risk sharing that explicitly takes these caveats into account and provides a structural interpretation for our results.

Table 2: Family and friends transfers (1988 PSID)

	Children		Parents	
	Any amount (%)	Amount (conditional)	Any amount (%)	Amount (conditional)
<i>Transfers given (money)</i>				
Total money given	0.123	1547 (2003)	0.142	3224 (5796)
To parents	0.021	780 (1217)	0.036	2416 (6816)
To children	0.049	2642 (2128)	0.090	3847 (5795)
To siblings	0.027	477 (406)	0.006	750 (354)
To other relatives	0.010	464 (386)	0.009	1767 (1935)
To non-relatives	0.021	545 (819)	0.006	158 (81)
<i>Transfers received (money)</i>				
Total money received	0.291	2231 (6796)	0.054	6947 (23350)
From parents	0.264	1913 (6236)	0.042	1554 (2749)
From children	0.000	NA (NA)	0.006	645 (502)
From siblings	0.016	2150 (4633)	0.003	2000 (NA)
From other relatives	0.019	5048 (12605)	0.000	NA (NA)
From non-relatives	0.019	831 (922)	0.000	NA (NA)
<i>Transfers given (time)</i>				
Total hours given	0.397	337 (670)	0.380	501 (929)
To parents	0.332	301 (659)	0.238	418 (864)
To children	0.001	408 (NA)	0.127	520 (738)
To siblings	0.071	106 (134)	0.030	352 (661)
To other relatives	0.034	118 (125)	0.030	371 (791)
To non-relatives	0.078	266 (598)	0.039	77 (93)
<i>Transfers received (time)</i>				
Total hours received	0.435	424 (721)	0.114	149 (230)
From parents	0.387	398 (687)	0.036	229 (389)
From children	0.000	NA (NA)	0.045	83 (55)
From siblings	0.073	237 (466)	0.012	104 (104)
From other relatives	0.013	75 (60)	0.009	183 (144)
From non-relatives	0.084	147 (196)	0.033	64 (55)

Note: The sample includes extended families in the 2013 PSID, where we define extended families as a cohabiting couple under 65 and their adult children over 25. The data come from child (columns (1) and (2)) and parent (columns (3) and (4)) reports of transfers in 1987. Only monetary transfers over \$100 are ascertained. Columns (1) and (3) report the percent of households reporting non-zero amounts of time or money, and columns (2) and (4) report mean amounts of time or money, conditional on a non-zero amount. Total transfers include family and non-family transfers. Standard deviations are in parentheses.

Table 3: Receipt of money and time transfers on income change quartile (1988 PSID)

Transfer from:	Full sample	Children			Parents
	Family & friends	Family & friends	Parents	Non-parents	Family & Friends
<u>Panel A: Money transfers</u>					
Income change quartile					
1 (negative)	0.062* (0.035)	0.109** (0.048)	0.130*** (0.047)	-0.020 (0.024)	-0.018 (0.030)
2	0.024 (0.038)	0.031 (0.048)	0.036 (0.048)	0.008 (0.022)	-0.076* (0.042)
3	-0.021 (0.035)	-0.021 (0.049)	-0.014 (0.048)	-0.000 (0.024)	-0.017 (0.030)
4 (positive) - omitted					
N	1033	701	701	701	332
<u>Panel B: Time transfers</u>					
Income change quartile					
1 (negative)	0.006 (0.043)	0.041 (0.057)	0.029 (0.055)	0.046 (0.037)	-0.056 (0.047)
2	-0.031 (0.044)	-0.067 (0.055)	-0.067 (0.054)	0.021 (0.038)	-0.036 (0.051)
3	-0.043 (0.042)	-0.050 (0.053)	-0.044 (0.053)	-0.007 (0.037)	-0.030 (0.046)
4 (positive) - omitted					
N	1033	701	701	701	332

Note: The table reports marginal effects coefficients from a probit regression of transfer receipt on income change quartile (highest quartile omitted). The sample includes extended families in the 1988 PSID, where we define extended families as a cohabiting couple under 65 and their adult children over 25. The data on transfers received come from all (column (1)), child (columns (2)-(4)) and parent (column (5)) reports of transfers in 1987. Only monetary transfers over \$100 are ascertained. Income is defined as the residual of a regression of log (post tax and public transfer) income on a set of demographics (see Section 3.1 for the full set of demographic controls), and income change is the difference between residuals in 1987 and 1986. Standard errors clustered by extended family are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

3 Risk sharing: A theoretical framework

In this section, we specify a model in which households choose consumption to maximize an intertemporal utility function given an exogenous income process and a budget constraint that reflects the insurance possibilities they have access to. Households are seen as part of a group, such as the extended family, and the income processes will be written to reflect this. That is, we decompose the household income process into a group component and a purely idiosyncratic one. This decomposition is useful as we want to consider explicitly the risk sharing possibility within the group. One could additionally decompose household income into additional components (say, an economy wide component, a sector component and so on). These decompositions would matter to the extent we want to consider insurance possibilities within those other groups.

We consider different market environments, ranging from complete markets with economy-wide perfect risk sharing, to an environment where households can perfectly share risk within a smaller group such as the extended family, to one where they only have access to ‘self-insurance’ in the form of individual savings (and possibly borrowing). The consideration of these different cases and some approximations of the consumption function allow us to consider intermediate cases in which households are able to insure a fraction of idiosyncratic shocks. Throughout this section, we assume that the only source of uncertainty is exogenous, post-tax and government transfer household income and that preferences over consumption are separable from leisure. We also abstract from labor supply decisions. While this is a simplification, we view insurance in our model as that provided above and beyond insurance that is incorporated in income (e.g. added worker effects, implicit worker-firm contracts, government transfers). Additionally, the time transfer results in the previous section suggest that labor – at least in the form of direct transfers – does not play a major extended family insurance role.⁶

⁶See Blundell, Pistaferri, and Saporta-Eksten (2016) and Attanasio, Low, and Sanchez-Marcos (2005) for models of consumption insurance that incorporate household labor supply decisions. Endogenous labor supply may matter in two ways: first, under nonseparability between consumption and leisure keeping the marginal utility constant may require lower (higher) transfers of income both because of substitutability (complementarity) between consumption and leisure and than in the separable case, because of added worker effects. Since such nonseparability would apply only within the nuclear family this mainly relates to how much insurance can be achieved in autarky. Second, some of the income variability that we attribute to unexpected shocks may be endogenous labor supply responses. This may lead to the impression of the existence of more insurance than reality. While these are all interesting issues, they are unlikely to affect much the quantification of insurance from the extended family network because they are likely to act proportionately on the relevant parameters.

3.1 Preferences and income processes

We begin by considering preferences and income processes. We assume that, at time t , each household values sequences of future consumption flows according to the expected utility they provide. Utility, in turn, is given by an intertemporally separable utility function that depends on household consumption at different points in time. We assume that the future is discounted geometrically and that utility is a concave function with standard regularity conditions. Therefore, sequences of consumption from time t to time T , $\mathbf{C}_{i,t} = \{C_{i,t}, C_{i,t+1}, \dots, C_{i,T}\}$, are valued by household i as $V_i(\mathbf{C}_{i,t})$:

$$V_i(\mathbf{C}_{i,t}) = E_t \sum_{s=t}^T \beta^{s-t} U(C_{i,s})$$

Notice that in addition to the standard restrictions used in the literature (such as that of intertemporal separability), we assume that utility for household i depends only on their consumption and is not affected by the consumption of other households.

The household is entitled to streams of uncertain income that are seen as exogenous stochastic processes $Y_{i,t}$.⁷ Following earlier empirical results, we model household income as a permanent-transitory process (MaCurdy (1983), Abowd and Card (1989) and Meghir and Pistaferri (2004)). This is because it provides a good and parsimonious fit to the stochastic structure of income and at the same time allows an important distinction between the ability to insure events that only have a temporary effect on income relative to events that may have a much more persistent effect, and hence a much larger impact on resources. The stochastic structure of income is made up of three components: (1) a deterministic component which we model as a function of demographics $z_{i,t}$,⁸ (2) a permanent component $P_{i,t}$, and (3) a transitory component $\nu_{i,t}$. In addition, measured income is affected by a multiplicative measurement error $m_{i,t}^y$.

$$\log Y_{i,t} = z_{i,t}\varphi_t + P_{i,t} + \nu_{i,t} + m_{i,t}^y$$

The permanent component follows a random walk in which the innovations, $u_{i,t}$, are mean zero and

⁷We are therefore assuming that labor supply is exogenous. The assumption that labor supply is separable from consumption and that wage effects are positive may underestimate the amount of insurance, but given that hours elasticities are relatively small and we see no time transfers in response to income shocks in the previous section, we do not expect a large bias.

⁸Specifically, in the estimation we control for year, age, education, race, family size, number of kids, region, urbanicity, and interactions of year with education, race, region, and urbanicity.

serially uncorrelated:

$$P_{i,t} = P_{i,t-1} + u_{i,t}$$

The transitory component follows an MA(q) process in which the innovations $e_{i,t}$ are also mean zero and serially uncorrelated:

$$\nu_{i,t} = e_{i,t} + \sum_{k=1}^q \theta_k e_{i,t-k}$$

In the estimation section, we determine that the transitory component follows an MA(1) process ($\theta_1 = \theta$ and $\theta_k = 0$ for all $k > 1$), so we henceforth write it as such. If we define $\log y_{i,t} \equiv \log Y_{i,t} - z_{i,t}\varphi_t$, the growth in the deviation of log income from its deterministic component is given by:

$$\Delta \log y_{i,t} = u_{i,t} + \Delta(e_{i,t} + \theta e_{i,t-1}) + \Delta m_{i,t}^y \quad (1)$$

In the rest of this section, we use this equation as the starting point from which households share risk.

3.2 Risk sharing arrangements

The second block of our conceptual framework is the definition of risk sharing groups. In this subsection and the next, we analyse two different risk sharing set-ups: the entire economy as a potential risk sharing group and a smaller group such as the extended family. Although the allocations that would prevail under full risk sharing in the economy at large are first best (under some assumptions), such allocations might not be attainable because of the presence of a number of frictions, might those be informational frictions or enforceability problems. In such a situation, it is interesting to consider smaller risk sharing arrangements, like the extended family, which might be better equipped to deal with certain types of frictions.

The simplest way to describe the properties of full risk sharing within a group G is to consider the problem of a social planner that maximizes the weighted average of the group members' utilities, subject to an aggregate budget constraint.⁹ Our approach focuses on intertemporal allocations *within group* G and it is completely agnostic about what happens *across groups*. Specifically, at

⁹For notational simplicity, we write the problem without aggregate savings. The conditions we would use would not be different in the presence of aggregate savings.

time t the social planner maximizes:

$$\max_{\{\mathbf{C}_{i,t}\}} E_t[\sum_{i \in G} \lambda_i V(\mathbf{C}_{i,t})] \quad \text{s.t.} \quad \sum_{i \in G} Y_{i,\tau}(\mathbf{s}_\tau) = \sum_{i \in G} C_{i,\tau}(\mathbf{s}_\tau) \quad \forall \tau \geq t \quad (2)$$

in which λ_i is the weight given by the social planner to household i and \mathbf{s}_t describes the state of the economy, which evolves as a Markov process such that the probability of state s_τ at time τ given the current state $s(t)$ at time t is $\mathbf{P}_{s(t)}^{s(\tau)}$. Given that all income realization for all households are fully contractible, the first order condition for consumption for each state of the world s_τ at time τ of household i is:

$$\lambda_i \beta U'(C_{i,\tau}) = \mu(\mathbf{s}_\tau) / \mathbf{P}_{s(t)}^{s(\tau)} \quad (3)$$

where $\mu(\mathbf{s}_\tau)$ is the multiplier associated with the aggregate budget constraint and $U'(C_{i,\tau})$ is the marginal utility of consumption for household i . Notice that, as all states of the world are fully contractible, marginal utility of consumption depends on the state of the world but not idiosyncratic differences between households in the group. If we take the ratio of this equation at two different points in time, τ and τ' , we obtain:

$$\frac{U'(C_{i,\tau})}{U'(C_{i,\tau'})} = \frac{\mu(\mathbf{s}_\tau) / \mathbf{P}_{s(t)}^{s(\tau)}}{\mu(\mathbf{s}_{\tau'}) / \mathbf{P}_{s(t)}^{s(\tau')}} = \nu(\tau, \tau') \quad (4)$$

Notice that in equation (4), the right-hand side does not depend on i , implying that the change in the marginal utility of consumption is the same across all households in the sharing group. Assuming power utility and a multiplicative measurement error in consumption, one can take the log of equation (4) considered at two adjacent time periods and obtain:

$$\Delta \log(c_{i,t}) = \psi_t + \varepsilon_{i,t} \quad (5)$$

Cochrane (1991) and Townsend (1994) test such an equation by adding to it a realization of idiosyncratic income and testing the hypothesis that the coefficient on such a variable is zero. The idea behind such a test is that under perfect risk sharing, household consumption adjusts in such a way that changes in marginal utilities (approximated by the log-changes in consumption under CRRA utility) is the same across households in the risk sharing group. Therefore, any shock to

household income should not enter significantly in such an equation. The test is silent and agnostic about the specific decentralization through which first best allocations are achieved or about the specific assets and contracts (formal and informal) that households might be using. Furthermore, under perfect risk sharing there is no distinction between (idiosyncratic) permanent and transitory shocks.

At the other extreme, one can consider an economy where households in group G have no risk sharing possibilities and they can only smooth income shocks using a single asset that pays an interest R_t , which can be either constant or variable. This market structure – the “Bewley” model – implies a very simple individual budget constraint. In such an environment, one is within the realm of a standard life cycle model in which the distinction between (idiosyncratic) transitory and permanent shocks matters: transitory shocks are almost fully smoothed out and permanent ones are, instead, almost completely reflected in consumption.¹⁰

As is well known, a closed form solution that expresses consumption as a function of the state variables to the problem (and innovations to the income process) can only be obtained under special circumstances, such as quadratic utility and constant interest rates. However, a number of contributions, including Blundell and Preston (1998) and Blundell, Pistaferri, and Preston (2008) use log-linear approximations to express innovations to consumption as a function of innovations to income. That is, they derive an equation of the following form:

$$\Delta \log c_{i,t} = \delta u_{i,t} + \gamma(1 + \theta)e_{i,t} + \Delta m_{i,t}^c + \xi_{i,t} \quad (6)$$

where δ and γ measure the degree to which permanent and transitory shocks, respectively, pass through to consumption. In addition, there may be preference shocks ($\xi_{i,t}$) and classical measurement error ($m_{i,t}^c$).

In a model such as Bewley’s, the values of δ and γ should be dictated solely by the ability to smooth shocks through self-insurance. Under CRRA preferences, self-insurance is attained through the potential to borrow from future income streams as well as precautionary savings. With an infinite horizon, $\delta = 1$ and $\gamma = 0$. With finite lives, Blundell, Pistaferri, and Preston (2008) show that an approximation of the Euler equation yields $\delta \simeq \pi_{i,t}$ and $\gamma \simeq \alpha_{i,t}\pi_{i,t}$ where $\pi_{i,t}$ is the

¹⁰The two ‘almost’ qualifiers in the previous sentence derive from the fact that the time horizon of the household problem is finite. The closer a household is to T , the more ‘permanent’ are ‘transitory’ shocks.

percentage of future income in current wealth (in other words, the percentage of lifetime wealth that is tied up in future income) and $\alpha_{i,t}$ is an annuitization factor. Equation (6) can also be empirically evaluated to estimate the fraction of permanent and transitory idiosyncratic income shocks that are transmitted to consumption in the data. The size of the coefficients can therefore be informative of the market structure that is relevant in a given context. Using this model, work by Blundell, Low, and Preston (2013) and Blundell, Pistaferri, and Preston (2008) find that $\pi_{i,t} = 0.8$ while $\delta = 0.64$. Under CRRA preferences, since self-insurance implies $\delta = \pi_{i,t}$, the empirical finding that $\delta < \pi_{i,t}$ can be interpreted as evidence of insurance above and beyond self-insurance. Next we turn to a model of extended family risk sharing that may help provide an explanation for this additional insurance.

3.3 Incorporating extended family risk sharing

We now consider explicitly risk sharing within a smaller group – specifically, the extended family – in comparison to economy-wide risk sharing.¹¹ The extended family might be particularly interesting as a risk sharing institution because it may be able to deal more effectively with the frictions that may underlie the failure of insurance in larger groups: (i) it might face less severe information constraints in the sense that shocks to the various family members may be better observable (i.e., avoiding moral hazard issues), and (ii) it may be easier to enforce commitment, which is important for implementing transfers. One can then relate the estimates of the ‘risk sharing’ parameters in equation (6), where one considers overall risk sharing, with those that obtained by explicitly considering the difference in risk sharing between extended family members and the whole economy.

To start, we can express the income process in equation (1) in terms of deviations of the household idiosyncratic component from the extended family aggregate component. In particular, we define $u_{j,t}^F$ as the aggregate permanent shock to extended family resources for extended family j , or the average of the shocks received by members of extended family j : $u_{j,t}^F = \frac{1}{n_j} \sum_{i=1}^{n_j} u_{i,j,t}$. We denote with $u_{i,j,t}^I$ the idiosyncratic permanent shock to household i in extended family j , or the deviation of the shock for household i from the family-aggregate shock $u_{j,t}^F$. We therefore decompose $u_{i,t}$ as the sum of $u_{j,t}^F$ and $u_{i,j,t}^I$. Analogously, we consider extended family aggregate temporary

¹¹To be clear, by extended family we mean multiple households that share familial ties, *not* a nuclear family within a single household.

shocks $e_{j,t}^F$ and the decomposition $e_{i,t} = e_{j,t}^F + e_{i,j,t}^I$. Notice that, by definition, the sum of the idiosyncratic shocks across extended family members is zero for both permanent and transitory shocks: $\sum_{i=1}^{n_j} u_{i,j,t}^I = \sum_{i=1}^{n_j} e_{i,j,t}^I = 0$. Rewriting equation (1), the growth in log income is:

$$\Delta \log y_{i,j,t} = u_{j,t}^F + u_{i,j,t}^I + \Delta(e_{j,t}^F + \theta e_{j,t-1}^F) + \Delta(e_{i,j,t}^I + \theta e_{i,j,t-1}^I) + \Delta m_{i,j,t}^y. \quad (7)$$

The interpretation of the aggregate components deserves attention at this point. If the overall shocks are i.i.d. across households, and abstracting from economy wide aggregate shocks, the aggregate-group shock should be zero for large groups; the fact that it may not be is a result of the finite (and small) number of members of extended families. In this case, the variance of the aggregate income shock is the variance of the household shocks divided by the number of group members. In the relatively small extended family groups we consider, these variances can be substantial. However, we do not take this statistical interpretation of the aggregate family shock literally, because it may well be that the stochastic processes of related extended family members are correlated as a result of common regional, occupational and educational environments for example.¹² In this case there is an additional component to the aggregate variance and the θ coefficients may differ between aggregate and idiosyncratic transitory shocks. For now, we remain agnostic about the origins of the family-aggregate shock and for parsimony we keep the restriction of a common θ , but we explore this restriction further in Section 7. The estimation of the permanent components, however, remain the same; it is only the interpretation that is modified.

The decomposition of income shocks into idiosyncratic and family-aggregate components allows us to quantify the percentage of shocks that *could* be insured by the extended family, which effectively defines the risk sharing opportunity of the extended family. Idiosyncratic shocks are, by definition, household-level deviations from the family-average shock, and hence the extended family network can redistribute funds between households to smooth these shocks. On the other hand, family-aggregate shocks, which arise from the fact that extended families are small and thus induce small-sample correlation in shocks, cannot be smoothed by extended family networks. Therefore, the pass-through of idiosyncratic income shocks to consumption may differ from the pass-through

¹²Indeed in the context of India, Rosenzweig and Stark (1989) argue that parents marry their daughters to males in other villages so that the extended family has sufficiently diversified risk, in this case associated to weather-related income shocks.

of family-aggregate shocks.

To study these differences in pass-through rates, we rewrite equation (6), the growth in log consumption, as:

$$\begin{aligned}\Delta \log c_{i,j,t} &= \delta_I u_{i,j,t}^I + \delta_F u_{j,t}^F + \gamma_I(1 + \theta)e_{i,j,t}^I + \gamma_F(1 + \theta)e_{j,t}^F + \Delta m_{i,j,t}^c + \xi_{i,j,t} \\ &= \delta_F u_{i,j,t} + (\delta_I - \delta_F)u_{i,j,t}^I + \gamma_F(1 + \theta)e_{i,j,t} + (\gamma_I - \gamma_F)(1 + \theta)e_{i,j,t}^I + \Delta m_{i,j,t}^c + \xi_{i,j,t}\end{aligned}\tag{8}$$

where δ_F measures the degree to which family-aggregate permanent shocks pass through to consumption and δ_I measures the degree to which idiosyncratic permanent shocks pass through to consumption. Similarly, γ_F and γ_I measure the sensitivity of consumption to transitory shocks that are family-aggregate and idiosyncratic, respectively.¹³

Equation (8) identifies the portion of extended family aggregate and purely idiosyncratic permanent and temporary shocks that are reflected into changes to individual consumption. It nests neatly the extreme cases of Bewley and perfect risk sharing models. For models with partial risk sharing, it measures the fraction of each shock that is reflected in consumption and can be informative of the type of frictions that prevent full risk sharing. For these types of models, a structural interpretation of the estimated coefficients can be difficult. For instance, in a model with imperfectly enforceable contracts, changes in consumption and the level of enforceable insurance will depend in a non-linear fashion on the entire distribution of shocks in the insurance group, rather than on the simple aggregate and individual decomposition used here. In some situations, such as the one considered by Attanasio and Pavoni (2011), it is possible to relate the estimated parameter to a structural parameter. In many other cases, such a mapping is not possible. Nevertheless, by identifying and estimating all of the parameters of this equation, we can consider simultaneously risk sharing within and across extended families.

It may be useful to recast our discussion of how the predictions of the two extreme models (the Bewley model and the model of perfect risk sharing) manifest themselves in the insurance parameters of equation (8). We focus the discussion on the partial insurance parameters for permanent

¹³In this specification, we treat the income process of all members of the extended family symmetrically (later in the paper we show that variances of the income process do not change with age). We also treat the transmission parameters symmetrically across different extended family members, which would be a natural assumption if extended families engage in risk sharing or if they have similar alternative sources of insurance. Note, however, that consumption *levels* are not necessarily symmetric; instead, this model only assumes that changes in marginal utility are symmetric.

shocks, δ_I and δ_F , as their permanence necessarily has larger welfare implications than transitory shocks, but the logic follows for transitory shocks as well.

Bewley model. Under autarky, insurance parameters are dictated solely by the ability of households to smooth consumption through self-insurance using the income stream of their household. In other words, the distinction between family-aggregate and idiosyncratic shocks is meaningless and has no bearing on consumption: both get transmitted into consumption to the same extent. It follows that, in this environment, $\delta_I = \delta_F$. In addition, as discussed above, partial insurance coefficients are a function of assets and age as a result of precautionary savings and the ability to borrow from future income.

Perfect extended family risk sharing. Under perfect extended family risk sharing, the distribution of income between extended family members has no effect on the distribution of consumption between extended family members (Hayashi, Altonji, and Kotlikoff (1996)). Thus, controlling for shocks to the extended family's aggregate resources, a shock to a household should have no effect on a household's consumption. This restriction is equivalent to $\delta_I = 0$ in our framework. In addition, because the distribution of income does not determine the distribution of consumption, the shock to aggregate resources should affect each household similarly (in terms of consumption growth). In our framework, this additional restriction corresponds to $\delta_F^A = \delta_F^B$ for any households A, B in extended family j . Overall, perfect extended family risk sharing predicts that $0 = \delta_I \leq \delta_F$.

Partial extended family risk sharing. Although, under partial extended family risk sharing, it may be difficult to give a structural interpretation of the estimated parameters, we expect $0 < \delta_I < \delta_F$. This would reflect an attenuation of the effect of insurable shocks, but less than complete insurance. We can separately estimate these parameters, and their difference quantifies the amount of insurance that the extended family offers over and above insurance offered by other channels, including public programs and the broader network of friends. This is shown more explicitly in the second line of equation (8), where the transmission of the idiosyncratic component of the shock $u_{i,j,t}^I$ is attenuated by $\delta_F - \delta_I$. Therefore, in what follows, we report our estimates of $\delta_F - \delta_I$ and interpret it as the fraction of idiosyncratic shocks that are insured through the extended family.

4 Identification

The model presented above can be seen as a stochastic factor model consisting of two equations, one for income growth and one for consumption growth. As is characteristic of factor models, each equation depends on several unobserved factors, some of which are common across individuals in a family and some of which are mutually independent.

To understand insurance within a group, it is useful to recover both the covariance structure of the income factors (which determines the extent of uncertainty facing households and the extent to which shocks are insurable within the extended family) as well as the coefficients associated with these factors in the consumption growth equation because they reflect the amount of insurance that occurs within and between extended families.

The specific set of parameters we wish to estimate are (a) the transmission parameters δ_I , δ_F , γ_I , and γ_F , (b) permanent income variances $\text{var}(u_t^F)$ and $\text{var}(u_t^I)$ and transitory income variances $\text{var}(e_t^F)$ and $\text{var}(e_t^I)$, (c) measurement error variances for consumption $\text{var}(m_{c,t})$ and income $\text{var}(m_y)$, and (d) consumption preference shock variances $\text{var}(\xi)$. We allow all variances to vary over time except the consumption preference shock variance and the income measurement error variance.¹⁴ To identify parameters, we exploit covariance structures across both time and within-family dimensions.

In addition to the factor structures of individual income and consumption growth in equations (7) and (8), we define family-level (j) income and consumption growth below as the averages over individuals (i) within an extended family:

$$\overline{\Delta \log y_{j,t}} \equiv \frac{1}{n_j} \sum_{i=1}^{n_j} \Delta \log y_{i,j,t} = u_{j,t}^F + \Delta(e_{j,t}^F + \theta e_{j,t-1}^F) + \frac{1}{n_j} \sum_{i=1}^{n_j} \Delta m_{i,j,t}^y \quad (9)$$

$$\overline{\Delta \log c_{j,t}} \equiv \frac{1}{n_j} \sum_{i=1}^{n_j} \Delta \log c_{i,j,t} = \delta_F u_{j,t}^F + \gamma_F(1 + \theta)e_{j,t}^F + \frac{1}{n_j} \sum_{i=1}^{n_j} \Delta m_{i,j,t}^c + \frac{1}{n_j} \sum_{i=1}^{n_j} \xi_{i,j,t} \quad (10)$$

Using these four equations (i.e. equations defining $\Delta \log y_{i,j,t}$, $\Delta \log c_{i,j,t}$, $\overline{\Delta \log y_{j,t}}$, and $\overline{\Delta \log c_{j,t}}$), we construct a matrix of the time series auto-covariances as well as cross-covariances between income and consumption for each of the household-level and family-level equations. Each covariance

¹⁴We could easily extend this to identify time-varying consumption preference shock variances and income measurement error variances, but since we do not allow them to vary over time in estimation (due to data concerns), we do not demonstrate this here.

consists of an observable component (i.e., the left-hand side of each of the four equations) and a function of the parameters of interest (i.e., the right-hand side of each of the four equations).

The structure we use implies a number of natural restrictions as well as more substantive assumptions on these covariances. At the heart of identification are two economic assumptions: (a) no borrowing restrictions, i.e. individuals can borrow and lend based on one riskless asset such that at the end of life they are left with no liability, and (b) the shocks to income are unanticipated, so there is no foresight.¹⁵ In addition, as described in Section 3, we represent log (residual) income as a random walk plus an MA(1) error. Had it been empirically warranted, more general structures can be allowed (e.g., an MA(k) error). We also assume that all shocks and measurement errors are independent of one another. A final set of restrictions follow from the definition of extended family level shocks: by construction, family-level shocks and idiosyncratic shocks for each household within the extended family are uncorrelated. Moreover, the idiosyncratic shocks of each household within an extended family are correlated in a specific way implied by the fact that they must sum to zero. These restrictions are completely innocuous as they follow from the definition.

With these covariances, we are able to identify each parameter of interest.¹⁶ Appendix A provides details of the exact covariance structures, but the basic identification argument is as follows. As is the case in prior studies, first- and second-order time-series covariances identify parameters related to transitory shocks and measurement error, while cross-sectional covariances identify parameters related to permanent shocks and taste shocks. Building on this, family-level covariances identify family-level parameters, while household-level covariances identify the remaining household-level parameters. This is because of the feature that all idiosyncratic (household-level) parameters drop out of family-level covariances since $\sum_{i=1}^{n_j} u_{i,j,t}^I = 0$ and $\sum_{i=1}^{n_j} e_{i,j,t}^I = 0$.

One particularly noteworthy feature of our model of risk-sharing is that it allows for identification of the variance of measurement error in income. This is not possible in most models that include both measurement error and transitory shocks because the variance of measurement error

¹⁵As discussed in Kaplan and Violante (2010), the approach taken by Blundell, Pistaferri, and Preston (2008), on which we build, is valid under perfect risk sharing and in a Bewley economy. It also works exactly in some partial risk sharing contexts and approximately well in others (Attanasio and Pavoni, 2011, for example). A violation of these assumptions would lead to the orthogonality conditions we use to be violated. However, Blundell, Pistaferri, and Preston (2008) show that at least the latter assumption of unanticipated shocks is not violated.

¹⁶Our model only requires the estimation of the covariance structure of consumption and income. As a result we need no assumptions on higher order moments or indeed on the full distribution of the shocks, which can be completely general, so long as the second order moments exist.

enters one-for-one with the variance of the transitory shock. Our model, however, puts additional structure on the transitory shock when we consider the role of the extended family. Intuitively, each of $\text{var}(e_t^I)$, $\text{var}(e_t^F)$, and $\text{var}(m_y)$ enter one-for-one in household-level covariances, but since $\sum_{i=1}^{n_j} e_{i,j,t}^I = 0$, family-level covariances do not contain $\text{var}(e_t^I)$, while $\text{var}(e_t^F)$ enters in full and $\text{var}(m_y)$ enters as an average (that is, multiplied by $1/n_j$) since family-level statistics are defined as averages over households in the extended family. Thus, family-level covariances allow for the separate identification of measurement error from transitory shocks.

5 Data and estimation

Our main data is the 1979-2011 waves of the Panel Study of Income Dynamics (PSID), which includes information on income, consumption, and demographics. We supplement this data with imputed non-durable consumption estimated using the 1980-2010 Consumer Expenditure Survey (CEX). Our method of imputation is identical to that of Blundell, Pistaferri, and Preston (2008).

5.1 PSID sample

We use the 1979-2011 PSID, a longitudinal study of US households that began in 1968 with a sample of around 5,000 households.¹⁷ Each survey contains questions about income, expenditures, and demographic and family information. In addition to the questions asked each survey wave, the PSID collect more detailed information on various themes. For example, in Section 2 we use supplementary data on time and monetary transfers between family and friends, which was collected in 1988 and 2013. In every year following 1968, both these original households and households that were formed after 1968 by children of the original households (“split-offs”) were followed yearly until 1997 and then every other year thereafter. Because the PSID follows split-offs, we are able to track and link parents and children of original households even after the children form their own households. These parent-child linkages are the basis of our extended family risk sharing network definition.

We define an “extended family” in our sample as a cohabiting couple and their adult children

¹⁷Of these, around 3,000 were representative of the US population (the core sample), and around 2,000 were low-income households (the Census Bureau’s SEO sample). In addition, in 1997 around 500 immigrant households were added in an effort to keep the sample representative of the US population.

who have broken off from the parent household to form their own households. To be clear, by “extended family” (or sometimes “family” – we use the terms interchangeably), we mean multiple separate households that share familial ties, *not* a nuclear family within a single household. We only include adult children who are at least 25 years old (to avoid large education changes) and only once they have split from the original household unit.¹⁸ We then follow the parent household and the children households until the parents divorce, reach age 65 (to avoid retirement issues) or die. Finally, we drop income and consumption outliers by trimming the top and bottom 1% of the income and consumption distributions as well as observations for which income or consumption is not measured in consecutive years (since estimation relies on *changes* in yearly income and consumption). This results in a final sample of 1,411 unique extended families consisting of a parent household and their adult children households, as shown in Table 4. Parents and children are on average 56 and 30 years old, respectively. Households within these extended families consist of three individuals on average, and child households have at least one young child on average. Income and consumption are slightly lower for child households, but not drastically so. Average extended family size is three households, i.e. a parent household and two child households.

Our measure of income is net annual household income from the prior year (see Appendix D for a more detailed description of our income and consumption measures). Specifically, we begin with the sum of taxable (wage and salary income, asset income, and business profits) and transfer (annuity, child support, retirement, SSI, TANF, UI, VA pension, welfare, workers compensation, and food stamps) income of members of the household, net of taxes and, importantly, not including help from family and friends. This measure allows us to control for the insurance that is offered by the tax and benefit system, but *not* for any insurance provided by informal transfers. Finally, we deflate income by the CPI-Urban deflator (with a base year of 1983-1984).

We use three measures of non-durable consumption. The first two stem from weekly food expenditures, which we aggregate to the yearly level.¹⁹ Food consumption is defined as expenditures on food eaten inside and outside the home as well as food stamps. We use this measure directly and also to construct a measure of total consumption, which is imputed from food consumption using

¹⁸We treat splitting off as a terminal state. If an adult child later moves back in with their parents, the PSID still classifies them as a separate household. Hence we still have separate income and consumption information that allows us to treat them as a separate household in our analysis.

¹⁹The reference period for food expenditure is not clear, so we assume that it is the same as the reference period for income, as previous work has done.

Table 4: Summary statistics

	Parents	Children
Age	56.02 (4.94)	30.16 (3.95)
Household size	2.79 (1.33)	3.09 (1.45)
Number of children in household	0.29 (0.76)	1.30 (1.21)
White	0.77 (0.42)	0.78 (0.41)
Married	0.98 (0.13)	0.72 (0.45)
No high school degree	0.31 (0.46)	0.06 (0.24)
High school graduate	0.34 (0.48)	0.38 (0.48)
Lives in big city	0.33 (0.47)	0.35 (0.48)
Annual income	33885 (18706)	27996 (15029)
Annual imputed consumption	17731 (11282)	14855 (9746)
Annual food consumption	4126 (1820)	3746 (1738)
Annual reported consumption	8871 (3655)	8415 (3909)
Extended family size		3.28 (1.24)
Number of unique extended families		1411
Number of family-year observations		8786
Years of data		25 (1980-2010)

Note: The sample includes extended families in the 1979-2011 PSID, where we define extended families as a cohabiting couple under 65 and their adult children over 25. Income is defined as total income net of taxes and family transfers. Imputed consumption is derived from a food demand system estimated on the CEX (see 5.2). Reported consumption is only available from 1999-2009 and is defined as the sum of expenditures on food, transportation, utilities, and home insurance. Income, imputed consumption, and reported consumption are deflated by the CPI-Urban deflator and food consumption is deflated by the food CPI deflator. Standard deviations are in parentheses.

estimates from a food demand system estimated on the CEX (see Section 5.2). Starting in 1999, the PSID began collecting information on additional consumption expenditures, so we construct a “reported consumption” measure that including gasoline, other transportation expenses, utilities, and home insurance.²⁰ Reported and imputed consumption are deflated by the CPI-Urban deflator, and food consumption is deflated by the food CPI deflator.

It is worth noting that these measures of consumption are based on expenditures, which is not equivalent to consumption per se. They do not include any in-kind transfers, such as food gifts or dinner invitations, or substituting expenditure with time (Aguiar and Hurst, 2005). To the extent that this an important transfer mechanism we may be missing some of the insurance that may take place.

5.2 Non-durable consumption from the CEX

To create a measure of non-durable consumption that spans the 1980-2010 period of our sample, we impute consumption using estimates of a food demand equation (see Blundell, Pistaferri, and Preston (2008) and Blundell, Pistaferri, and Preston (2006) for more technical details) from the 1980-2010 CEX. The CEX is a short rotating panel survey of US households with detailed information on hundreds of expenditure categories as well as demographic and earnings information. The consumption definition we use is the sum of food (at home and outside the home), alcohol, tobacco, and expenditure on other non-durable goods such as services, utilities, transportation expenses, personal care, clothing, and footwear. For each year of data, we estimate a regression of log food on the number of children, age, self-employment status, education, log consumption, and log consumption interacted with education. The estimated coefficients are then used to impute a measure of consumption in the PSID.

5.3 Estimation procedure

Estimation of the income and transmission parameters of the model involves three steps. First, we construct a measure of the change in unexplained log income and log consumption for households and calculate data covariances. Second, we estimate the parameters of the income process, and

²⁰The PSID also collected data on health, childcare, education, and rent, but we exclude these categories to be more consistent with our imputed consumption definition.

third, we estimate the consumption transmission equation parameters.

In the first step, we remove the impact of deterministic and aggregate effects on log income and log consumption by taking the residuals $\log y_{i,t}$ and $\log c_{i,t}$ of a regression of each variable on dummies of household size, number of children under 18, age, education, race, region, whether the household lives in a large city, whether the household is part of the SEO or immigrant sample, and interactions of year with education, race, region, urbanicity, and SEO and immigrant sample status. We take differences over time of each variable to obtain log unexplained income growth and log unexplained consumption growth. To obtain the average growth by extended family, we take the average log unexplained growth of all households in the extended family: $\overline{\Delta \log y_{j,t}} = \frac{1}{n_j} \sum_j \Delta \log y_{i,j,t}$ and $\overline{\Delta \log c_{j,t}} = \frac{1}{n_j} \sum_j \Delta \log c_{i,j,t}$, where n_j is the number of households in extended family j . Finally we take covariances of these measures of income and consumption growth.

In the second and third steps, we estimate the parameters of the model with the moments described in Section 4, using minimum distance.²¹ We use diagonally-weighted minimum distance, which imposes that the weighting matrix is a diagonal matrix whose diagonal entries are the diagonal of the inverse of the variance-covariance matrix. This is similar to the optimal weighting matrix except the off-diagonal elements are zero and hence we avoid problems related to optimal minimum distance (Altonji and Segal, 1996). In the second step, we exclusively use income moments to estimate the parameters of the income process. In the third step, we take the estimates of the income parameters as given and use consumption moments as well as consumption crossed with income moments to estimate parameters of the consumption model. We compute standard errors using block bootstrap over all three steps of the estimation procedure, clustering at the extended family level (Hall and Horowitz, 1996; Horowitz, 2001). This method allows us to account for arbitrary serial correlation between extended family members and for the fact that the second and third steps took the estimates from the previous steps as given.

6 Results

Our goal is to quantify the extent of insurance provided by a household's extended family risk sharing network. To arrive at this, we first characterize the structure of the income and consumption

²¹Estimating family-average moments is complicated by the fact that n_j is not constant within our sample. See Appendix B for a short discussion of this issue.

processes, and then estimate the relevant parameters.

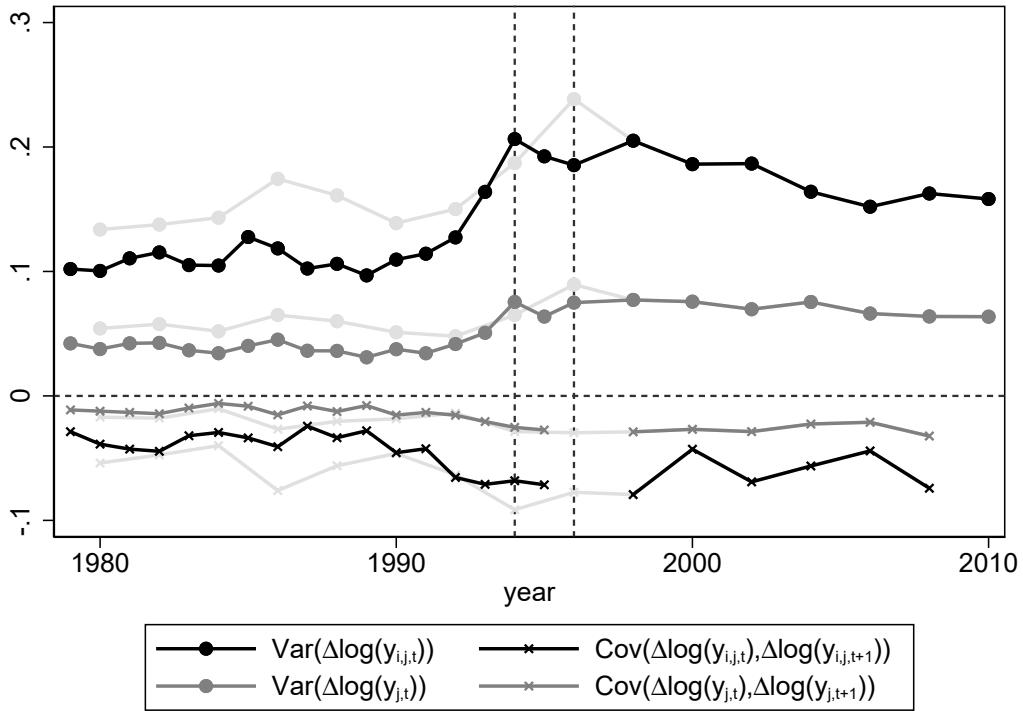
6.1 Characterization of the income and consumption processes

Figure 1 plots variances and first-order autocovariances of residual log income growth by year at both the household level and the family-average level (see Appendix Table 1 for point estimates and standard errors). The family-average variances are much smaller than the household variances. This result is expected, since variances of household level income reflect both purely idiosyncratic and family-average shock variances, while family-average variances reflect only the latter. The size of the difference, however, gives an idea of how much of the idiosyncratic income variance could potentially be insured by the extended family. A second feature to note about the variances is that variances after 1996 are considerably larger than before 1996. This difference can be partially explained by the fact that the PSID started collecting data biennially beginning in 1996: the lightest gray lines prior to 1996 depict what the variance estimates would be if only even years of data were available, and show that they are almost always higher than the variances derived from single-year differences. We take into account this change in survey structure in the estimation below (details can also be found in Appendix C).

In terms of the first-order autocovariances, both the household and extended family level autocovariances are negative (consistent with previous literature), and the household level autocovariance is considerably larger in absolute value than the extended family average autocovariance, similarly to the variance patterns. In addition, the ratio of the first order autocovariances to the variances for the individual and family level shocks is roughly of the same order of magnitude, which helps corroborate our assumption that the MA(1) parameter θ is consistent over time. We additionally use second-order autocorrelations, at least prior to 1996, to help characterize and identify parameters of the income process (see Appendix Table 1). These autocorrelations are even smaller, in absolute value, than the first order autocorrelations and are often not significantly different from zero. When they are, they are typically negative, which also justifies our modeling specification as these moments should be small and negative.


Figure 2 repeats the exercise in Figure 1 but with consumption (and similarly, see Appendix Table 2 for point estimates and standard errors). As with income, there are large differences between household and family-average variances, and we again note the large differences between

Figure 1: Log income growth variances and autocovariances



Note: This figure reports variances of log income growth (circle markers) and first-order autocovariances of log income growth (x markers) at the individual household level (black) and the family-average level (gray), where individual households are denoted by i and family averages are denoted by j , as defined in Section 3.3. The vertical lines at 1994 and 1996 indicate PSID changes in income coding and survey frequency, respectively. The lightest gray lines denote the variances and covariances derived from long-differences if only even years were available. The sample includes “families” in the 1979-2011 PSID, where we define “families” as a cohabiting couple under 65 and their adult children over 25. Income is defined as total income net of taxes and family transfers.

Figure 2: Log consumption growth variances and autocovariances, imputed from CEX



Note: This figure reports variances of log consumption growth (circle markers) and first-order autocovariances of log consumption growth (x markers) at the individual household level (black) and family-average level (gray) where individual households are denoted by i and family averages denoted by j , as defined in Section 3.3. The vertical line at 1996 indicates a change in PSID survey frequency. The lightest gray lines denote the variances and covariances derived from long-differences if only even years were available. Gaps in the late 1980’s are due to missing food consumption questions in 1988 and 1989. The sample includes “families” in the 1979-2011 PSID, where we define “families” as a cohabiting couple under 65 and their adult children over 25. Consumption is imputed and is derived from a food demand system estimated on the CEX (see 5.2).

variances prior to and after 1996. Unlike income, the first-order autocovariances fully consist of measurement error in consumption (for single-year differences, at least; see equations (23) and (28) in the appendix), which includes both recall error and error from the imputation procedure. Finally, higher-order autocovariances are insignificant (not shown), which – like income – is consistent with the implications of the model.

6.2 Parameter estimation

To the best of our knowledge, this section presents the first characterisation of the income process that decomposes the shocks a household receives in those of the extended family and those that are idiosyncratic to the individual household. In Table 5, we report average estimates of the income process.²² As mentioned above, we assume that the persistence parameter of transitory shocks θ is the same for the idiosyncratic and family-average components of income, while the variances are allowed to differ.

The first thing to note is that both permanent (columns 1 and 2) and transitory (columns 4 and 5) shocks are important sources of risk. More importantly for the purposes of this paper is to compare the magnitude of the idiosyncratic component to the family-aggregate component of shocks, which effectively determines the amount of opportunity there is for the extended family to share risk. At one extreme, if individuals' shocks ($u_{i,t}$ and $e_{i,t}$, i.e., the sum of the idiosyncratic and family-aggregate components) are perfectly positively correlated among extended family members, the extended family is an ineffective group within which to share risk, because each member is faced with the same shock (in other words, the idiosyncratic component is zero). At the other extreme, if shocks are perfectly negatively correlated among extended family members, the family can completely smooth fluctuations between extended family members (in other words, the family-aggregate component is zero).

We find that the idiosyncratic component of income accounts for a substantial proportion of the variance of the overall permanent and transitory shocks. For both types of shocks, on average, the idiosyncratic component makes up around 60% of the overall variance of the shock, as shown in columns (3) and (6). This means that over 60% of shocks are potentially insurable by extended family risk sharing networks. The question, of course, is the extent to which it does.

²²For the full set of time-varying income parameter estimates, see Appendix Table 3.

Table 5: Income process estimates

Variance of permanent shocks			Variance of transitory shocks		
$\text{var}(u_{i,j,t}^I)$ (idiosyncratic)	$\text{var}(u_{i,j,t}^F)$ (family)	% insurable by family	$\text{var}(e_{i,j,t}^I)$ (idiosyncratic)	$\text{var}(e_{i,j,t}^F)$ (family)	% insurable by family
0.017 (0.003)	0.011 (0.002)	0.578 (0.054)	0.027 (0.008)	0.015 (0.004)	0.644 (0.043)
θ (serial correlation of transitory shock)			0.163 (0.142)		
$\text{var}(m_y)$ (income measurement error variance)			0.015 (0.011)		

Note: This table reports pooled estimates of the income process (Appendix Table 3 reports the year-by-year estimates) where idiosyncratic household and family-aggregate components are defined as in Section 3.3. Columns 1-3 report estimates for permanent shocks and columns 5-7 report estimates for transitory shocks. Columns 3 and 7 report the percent of the total income shock that is potentially insurable by the extended family. The sample includes “families” in the 1979-2011 PSID, where we define “families” as a cohabiting couple under 65 and their adult children over 25. Income is defined as total income net of taxes and family transfers. Standard errors in parentheses are based on 100 block bootstrap replications, clustered by extended family.

There are two other important aspects to note before moving onto the insurance estimates. First, the estimates of the transitory shock variances are purged of measurement error. As shown in Appendix A, we are able to separately identify the variance of measurement error from transitory shocks. Our estimate, $\text{var}(m_y)=0.015$, is remarkably close to the estimate of 0.0138 used in Meghir and Pistaferri (2004) and subsequent papers using similar estimation strategies, which is based on Bound and Krueger (1991)’s validation study of CPS data. However, our estimate is noisy, as is the MA parameter on the transitory shock. If the true θ is in fact zero, then the variance of the measurement error of income is not identified using income data alone. We run robustness checks in which we use an external estimate of the variance of measurement error, as well as setting $\theta = 0$. Results (not shown, but available on request) are almost identical. Second, we estimated the overall income processes separately for younger households (heads aged 25-44) and older households (heads aged 45-64), and did not find any difference in the variance of their shocks.²³ Thus it does not appear that older households in our sample have systematically more stable income processes than younger households.

Before discussing the extended family insurance estimates, Table 6 presents partial insurance estimates based on a model in which the household may be a member of some unspecified insurance network, or indeed just self insured based on their own assets. In other words, we ignore the division

²³Pooling all years together, the variance of the permanent (transitory) shocks for younger households is 0.033 (0.035) while those of the older households is 0.032 (0.040).

of income into an extended family and household component. This is essentially the Blundell, Pistaferri, and Preston (2008) model except that our income measure includes fluctuations that may occur due to changes in employment, while Blundell, Pistaferri, and Preston (2008) effectively only allow for shocks to wages.²⁴

Table 6: Partial insurance estimates

	δ (Permanent)	γ (Transitory)
<i>All years (1980-2008)</i>		
Imputed consumption	0.499 (0.093)	-0.006 (0.052)
Food consumption	0.313 (0.058)	-0.026 (0.034)
<i>Recent years (1998-2008)</i>		
Reported consumption	0.283 (0.085)	0.042 (0.065)
Imputed consumption	0.347 (0.141)	0.073 (0.110)
Food consumption	0.234 (0.088)	-0.006 (0.064)

Note: This table reports partial insurance estimates against permanent shocks (column (1)) and transitory shocks (column (2)) based on equation (6) for various measures of consumption. Slope heterogeneity and time-varying measurement error variances also estimated but not reported. The sample includes “families” in the 1979-2011 PSID, where we define “families” as a cohabiting couple under 65 and their adult children over 25. Imputed consumption is derived from a food demand system estimated on the CEX (see Section 5.2). Reported consumption is only available from 1996-2009 and is defined as the sum of expenditures on food, transportation, utilities, and home insurance. Imputed consumption and reported consumption are deflated by the CPI-Urban deflator and food consumption is deflated by the food CPI deflator. Standard errors based on 100 block bootstrap replications (clustered at the extended family level) are in parentheses.

To perform this exercise, we sum the idiosyncratic and family-aggregate components of income estimated above and estimate the consumption parameters as in equation (6). Consistent with prior work, we cannot reject the hypothesis that transitory shocks are completely smoothed and not reflected in consumption (column 2). Of course this could happen either through intra-household and intra-family transfers or through self-insurance (savings). In contrast, households do not perfectly smooth permanent shocks (column 1). Over the full time span of our data (1980-2008), a 10% decrease in permanent income is associated with a 5.0% decrease in non-durable consumption, a 3.1% decrease in food consumption. For reported non-durable consumption, which is only available since 1998, this income decrease is associated with a 2.8% decrease in consumption.²⁵ This evidence on

²⁴There is a logic to controlling for changes in employment since this is in part an endogenous decision. However our model does not allow for labor supply and it is also hard to control for fluctuations in employment and at the same time control for its endogeneity, something that is ignored in Blundell, Pistaferri, and Preston (2008).

²⁵Note that when we estimate δ using imputed consumption over the same time period as reported consumption,

the pass-through of permanent shocks is also consistent with prior work and unsurprising given that the lifetime effect of a permanent shock is much larger than a transitory shock. However, the fact that 50% of the decrease in income does *not* translate into non-durable consumption indicates that there is a substantial amount of insurance against shocks to disposable income above and beyond the government tax and transfer system, which we explicitly take into account. As we argue in this paper, the observed level of insurance may be due to a number of different channels. It may be because of self-insurance through assets, or because of insurance within various networks. We now turn our focus to whether it is because of the extended family network.

Table 7 reports estimates of the partial within-family insurance coefficients. The first group of estimates show transmission parameters of permanent shocks, and the second group of estimates show transmission parameters of transitory shocks. For both groups, the second column (δ_F and γ_F) reports the transmission of shocks that are aggregate to the extended family and hence not insurable through extended family networks, while the first column (δ_I and γ_I) reports the transmission of the idiosyncratic component of shocks which are therefore insurable through extended family networks. The difference between these columns ($\delta_F - \delta_I$ and $\gamma_F - \gamma_I$), reported in column 3, quantifies the level of extended family insurance. As found in the overall insurance estimates of Table 6, transitory shocks are fully insured. This is even true for shocks that are aggregate to the extended family, suggesting that households are able to insure these shocks through other means, such as savings.

For permanent shocks, despite the finding of partial insurance in Table 6, we find no evidence that this partial insurance is driven by insurance within the extended family. For the full year sample, the difference between the transmission of extended family-aggregate and idiosyncratic shocks (column 3) is very close to zero. This difference is larger when the sample is constrained to more recent years (for which we additionally have reported consumption), but these effects are still relatively small. This suggests that at least overall the extended family in the US is not a source of partial insurance.

While we find no evidence of insurance within the extended family here, the results in Section 2 show that transfers between extended family members occur, and are responsive to shocks. These two results are not necessarily inconsistent. First, the mere existence of transfers does not imply

the value of δ decreases by one-third. This may explain part of the difference in the estimates using the two measures of consumption.

Table 7: Extended family insurance estimates

	Permanent			Transitory		
	δ_I (idiosyncratic)	δ_F (family)	$\delta_F - \delta_I$	γ_I (idiosyncratic)	γ_F (family)	$\gamma_F - \gamma_I$
<i>All years (1980-2008)</i>						
Imputed	0.494 (0.139)	0.516 (0.103)	0.022 (0.159)	-0.022 (0.070)	0.020 (0.071)	0.042 (0.094)
Food	0.323 (0.084)	0.286 (0.061)	-0.037 (0.093)	-0.031 (0.043)	-0.012 (0.046)	0.019 (0.059)
<i>Recent years (1998-2008)</i>						
Reported	0.215 (0.165)	0.355 (0.143)	0.139 (0.229)	0.081 (0.125)	0.038 (0.077)	-0.043 (0.142)
Imputed	0.278 (0.249)	0.449 (0.256)	0.170 (0.368)	0.183 (0.232)	0.008 (0.150)	-0.175 (0.281)
Food	0.215 (0.126)	0.255 (0.148)	0.040 (0.185)	0.023 (0.108)	-0.021 (0.090)	-0.044 (0.141)

Note: This table reports partial insurance estimates against permanent shocks (columns (1)-(3)) and transitory shocks (columns (4)-(6)) based on equation (8) for various measures of consumption. Columns (3) and (6) report the level of extended family insurance. Slope heterogeneity and time-varying measurement error variances also estimated but not reported. The sample includes “families” in the 1979-2011 PSID, where we define “families” as a cohabiting couple under 65 and their adult children over 25. Imputed consumption is derived from a food demand system estimated on the CEX (see Section 5.2). Reported consumption is only available from 1996-2009 and is defined as the sum of expenditures on food, transportation, utilities, and home insurance. Imputed consumption and reported consumption are deflated by the CPI-Urban deflator and food consumption is deflated by the food CPI deflator. Standard errors based on 100 block bootstrap replications (clustered at the extended family level) are in parentheses.

an insurance relationship: households may transfer money for other reasons, such as gifts (e.g., birthdays) or social norms (e.g., parents often pay for higher education and weddings). This is consistent with the relatively larger transfers from parents to children and smaller transfers to other family members in Tables 1 and 2, and other evidence of parental transfers for particular life events (McGarry, 2016). Second, our evidence on transfers in response to income shocks (Table 3) shows that parents provide significantly more monetary transfers to their children when their children have more negative income shocks, but it is difficult to interpret this result from a risk-sharing perspective because (a) risk-sharing suggests that a similar result would appear for child transfers to parents who receive negative income shocks, (b) the results are relatively small (children in the bottom 25% of the income shock distribution are only 13 percentage points more likely to receive a transfer than children in the top 25%) and only speak to the extensive margin of transfers, and (c) perhaps most importantly, these regressions do not condition on the extended family budget constraint. In other words, if other extended family members also had negative shocks, we would not necessarily expect to see transfers.

Why is there no detectable insurance within the extended family? In the next section we rerun our analysis for various subsamples to understand whether it is due to moral hazard problems, certain types of shocks, and whether certain characteristics of extended families are better able to insure its members. We also examine the robustness of our results to measurement problems and alternative definitions of extended family shocks.

7 Robustness

In this section we conduct additional robustness exercises to our main result of a lack of within-extended family insurance. We first run our estimation procedure on additional sub-samples to explore potential non-linearities in insurance. We then analytically derive the impact of mismeasuring the proper extended family. Finally, we theoretically and empirically explore the role of correlated shocks across extended family members. Our main findings remain true, but the findings of this section suggest interesting extensions when applying our framework to other risk-sharing settings.

7.1 Additional sample breakdowns

The estimates in the previous section imply that over the whole population of extended families in our sample, there is no evidence of extended family insurance. However, this overall result may mask heterogeneity across sub-populations, which may also provide insight into the underlying mechanisms for why risk-sharing between extended family members succeeds or fails. Table 8 shows estimates for various sample breakdowns, concentrating on the imputed consumption measure over all years of data.

One reason risk sharing might break down is moral hazard or incomplete (and asymmetric) information: in an environment in which households have an incentive to shirk or hide assets, perfect risk-sharing is not always attainable. One interpretation of the results in the previous section is that the extended family in the US is not sufficiently linked to be able to improve monitoring of actions and verification of shocks. While our framework cannot directly test this theory, row 1 in Table 8 shows estimates for extended families in which all members live in the same state. This restriction captures the idea that extended families that live near one another may be more likely to monitor

Table 8: Extended family insurance estimates, various subsamples

	Permanent			Transitory		
	δ_I (idiosyncratic)	δ_F (family)	$\delta_F - \delta_I$	γ_I (idiosyncratic)	γ_F (family)	$\gamma_F - \gamma_I$
Same state	0.416 (0.171)	0.619 (0.132)	0.203 (0.215)	0.015 (0.087)	-0.060 (0.100)	-0.076 (0.121)
Small extended family	0.582 (0.243)	0.610 (0.171)	0.029 (0.303)	0.030 (0.299)	0.013 (0.259)	-0.017 (0.534)
Negative employment shock	0.259 (0.192)	0.151 (0.080)	-0.107 (0.185)	0.111 (0.162)	-0.044 (0.064)	-0.155 (0.172)
Large negative income shock	0.508 (0.244)	0.320 (0.119)	-0.188 (0.288)	-0.168 (0.117)	-0.093 (0.110)	0.076 (0.173)
High wealth	0.365 (0.138)	0.361 (0.112)	-0.004 (0.179)	-0.062 (0.364)	0.088 (0.153)	0.151 (0.445)
Low wealth	0.603 (0.185)	0.471 (0.142)	-0.131 (0.234)	-0.087 (1.260)	0.049 (0.103)	0.135 (1.268)
No parental coresidence	0.560 (0.143)	0.590 (0.102)	0.030 (0.158)	-0.053 (0.075)	-0.018 (0.099)	0.035 (0.106)
Kids unmarried	0.123 (0.299)	0.413 (0.188)	0.290 (0.348)	0.157 (1.292)	-0.671 (0.905)	-0.828 (1.782)

Note: This table reports partial insurance estimates against permanent shocks (columns (1)-(3)) and transitory shocks (columns (4)-(6)) based on equation (8) for various measures of consumption and various subsamples. Columns (3) and (6) report the level of extended family insurance. Slope heterogeneity and time-varying measurement error variances also estimated but not reported. Subsample analyses (run separately) use imputed consumption and set $\text{var}(m_y) = 0.014$. The sample includes “families” in the 1979-2011 PSID, in which we define “families” as a cohabiting couple under 65 and their adult children over 25. Imputed consumption is derived from a food demand system estimated on the CEX (see 5.2) and is deflated by the CPI-Urban deflator. ‘Same state’ is defined as all family members living in the same state; ‘small extended family’ is families consisting of the parent and one child household; ‘negative employment shock’ is households that transitioned from employment to non-employment (for the consumption sample only); ‘large negative income shock’ is households whose residual log income change was in the bottom quartile (for the consumption sample only); ‘high wealth’ and ‘low wealth’ are families whose parent household average assets over the sample are in the top and bottom 40% of the asset distribution of parents, respectively; ‘no parental coresidence’ is families in which the child household never moves back into the parent household; ‘kids unmarried’ are families in which all child households are unmarried. Standard errors based on 100 block bootstrap replications (clustered at the extended family level) are in parentheses.

the behavior and events of each other than families who are more geographically disperse. Focusing on the permanent shocks, the results of this exercise show suggestive evidence that families who live near one another provide some insurance to one another: while extended family insurance is $\delta_F - \delta_I = 0.02$ for all extended families, it is $\delta_F - \delta_I = 0.20$ for geographically close extended families. This number is imprecisely estimated, however, so while highly suggestive, this evidence cannot be taken as conclusive because we cannot reject the hypothesis that the coefficients are the same. A larger data set would be useful to pursue this channel further, as would more detailed information on geographic distance. Of course, there is also the possibility that distance itself is endogenous in the sense that families with weaker links are more likely to live far apart from each other. An additional sample that may also be better than average at curbing moral hazard problems is small extended families since there are fewer people to “monitor.” Row 2 shows the results for the sample of families with two households (i.e., the parent household and one child household), and finds no evidence of extended family insurance for this sample.

A second theory of extended family interactions is that extended family members provide transfers when large income shocks occur (such as a job loss), but less so for small income shocks. Such large shocks may be easier to verify. Moreover, this could be the result of an underlying model of transfer frictions, in which families must pay a monetary or utility cost to give or receive transfers. To examine this non-linearity hypothesis further, we restrict the consumption sample in two different ways: first, to households that transitioned from employment to non-employment in row 3, and second, to households whose residual income change was in the bottom 25th percentile in row 4.²⁶ For both of these samples, there is no evidence of extended family insurance; in fact, the within-family insurance for the large negative income shocks run is even wrong-signed. Interestingly, however, negative employment shocks are overall better insured (only around 20% pass-through) than the overall sample of shocks, and also better insured than unspecified large negative shocks. One interpretation of these findings is that taggable shocks such as unemployment are more generally insurable than other shocks, potentially due to their visibility and verifiability.

Another important mechanism for insuring shocks is self-insurance. To explore how self-insurance interacts with extended family insurance, rows 5 and 6 split the sample into families

²⁶The income estimates for these samples remain the full sample estimates; it is just the consumption parameters that are estimated on these restricted samples.

whose parents' average assets over the sample period (defined as asset income divided by the interest rate, plus housing assets) are in the top and bottom 40% of the asset distribution of parents, respectively. Families with larger asset holdings are better able to self-insure, which may translate into a higher willingness to insure other family members. We find that generally the transmission of permanent shocks is lower for high-asset families than low-asset families, which is consistent with a simple consumption-savings model in which families with higher assets are better able to self-insure. On the other hand, we are not able to reject that the difference between the two sets of estimates is zero. This points in the direction of self-insurance as the main mechanism, rather than intra-family transfers.

The final two rows explore robustness to the way we measure consumption and the risk-sharing network. One concern is that we are incorrectly measuring extended family insurance provided through alternative means, such as coresidence. If some families provide insurance through monetary transfers and other families provide insurance through in-kind transfers, then we may be mis-measuring the extent of monetary insurance by including the families that provide in-kind transfers. To explore this, Row 7 excludes extended families in which the child households move back in with the parent household at some point during the sample period. For this sample, which may be more likely to engage in monetary transfers, we again see no evidence of extended family insurance. Finally, the last row explores the extent to which we are missing other members of the risk-sharing network by only including extended families in which the children are unmarried, because married children may additionally have access to insurance from the spouse's extended family. This is a very small sample of less than 2,600 households, so the estimates are not precise, but this specification – at least for permanent shocks – results in some amount of extended family insurance. Given these results, the next subsection provides an analytical framework for understanding the possible mis-measurement of risk-sharing groups.

7.2 Mismeasuring the extended family

In all of our analyses, we assume that we have defined, and observed, the appropriate members of the risk-sharing group. How do the insurance results change if we mismeasure who is in the group? To answer this, we analytically derive the bias of the insurance coefficients under three separate scenarios: (1) the measured group is missing one member of the true group, (2) the measured group

includes an extra member that is not in the true group, and (3) the measured group has swapped an incorrect member in place of a true member.

Consider a simple model of extended family risk-sharing of permanent shocks:

$$\Delta \log y_{i,j,t} = u_{i,j,t} = u_{j,t}^F + u_{i,j,t}^I \quad (11)$$

$$\Delta \log c_{i,j,t} = \delta_F u_{j,t}^F + \delta_I u_{i,j,t}^I \quad (12)$$

in which $u_{i,j,t}$ are drawn from an IID distribution with mean zero and variance σ_u^2 . This is equivalent to Equations (7) and (8) without transitory shocks, measurement error, or consumption heterogeneity. While the following analysis could be easily extended to include transitory shocks, we focus on permanent shocks for simplicity and because they are more pertinent for risk-sharing since transitory shocks are well-insured even without risk-sharing. Identification of the insurance parameters δ_F and δ_I follows from:

$$\delta_F^2 = \frac{\text{var}(\overline{\Delta \log c})}{\text{var}(\overline{\Delta \log y})} \quad (13)$$

$$\delta_I^2 = \frac{\text{var}(\Delta \log c) - \text{var}(\overline{\Delta \log c})}{\text{var}(\Delta \log y) - \text{var}(\overline{\Delta \log y})} \quad (14)$$

where the $\overline{\Delta \log y}$ and $\overline{\Delta \log c}$ signify averages over the group (with the subscripts suppressed).

Table 9: Effect of mismeasuring the group on insurance parameters


Group type	δ_F^2	δ_I^2
True group	δ_F^2	δ_I^2
Missing member	$\frac{N-1}{N} \delta_F^2 + \frac{1}{N} \delta_I^2$	δ_I^2
Extra member	$\frac{1+N^2}{N(1+N)} \delta_F^2 + \frac{N-1}{N(1+N)} \delta_I^2$	$\frac{2}{N(N+1)} \delta_F^2 + \frac{(2+N)(N-1)}{N(N+1)} \delta_I^2$
Wrong member	$\frac{(N-1)^2+1}{N^2} \delta_F^2 + \frac{2(N-1)}{N^2} \delta_I^2$	$\frac{2}{N^2} \delta_F^2 + \frac{(N^2-2)}{N^2} \delta_I^2$

Suppose that the true group consists of N members. Table 9 shows the values of δ_F^2 and δ_I^2 when the measured group contains $N - 1$ members (row 2), $N + 1$ members (row 3) and N members but one of them is not part of the true group (row 4).²⁷ Overall, this table shows that when the risk-sharing group is mismeasured, the measured transmission parameters are a convex combination

²⁷Appendix Table 4 shows the corresponding measured income and consumption variances that generate these insurance values.

of the two true transmission parameters, and the amount of bias depends on the size N of the risk-sharing group. In the context of extended families, arguably the most relevant case is when a member of the group is missing, e.g. a parent's sibling is also a member of the true risk-sharing group. In this case, interestingly, our method recovers the true transmission of idiosyncratic shocks. However, the main object of interest is the difference between the transmission of family-aggregate and idiosyncratic shocks, $\delta_F - \delta_I$. Figure 3 quantifies the error in the measured amount of extended family insurance, or $(\delta_F - \delta_I)^{\text{meas}} - (\delta_F - \delta_I)^{\text{true}}$ for different levels of the true transmission of the idiosyncratic shock δ_I^{true} (x-axis), different levels of the true transmission of the family-aggregate shock δ_F^{true} (100% transmission in solid and 50% transmission in dashes), and different size groups (shading).

Figure 3: Effect of mismeasuring the group on the mismeasurement of extended family insurance



Note: The figures correspond to the difference between the extent of measured extended family insurance ($\delta_F^{\text{meas}} - \delta_I^{\text{meas}}$) minus the extent of true extended family insurance ($\delta_F^{\text{true}} - \delta_I^{\text{true}}$), for three cases: (A) the measured group is missing one true member, (B) the measured group includes an extra non-member, and (C) the measured group swaps out a true member for a non-member. For each case, the bias is plotted by the true transmission of idiosyncratic shocks (δ_I) on the x-axis and two values of the true transmission of family-aggregate shocks: $\delta_F = 1.0$ (solid) and $\delta_F = 0.5$ (dashed). For each of those two cases, the bias is plotted for group sizes of $N = 3$ (black), $N = 5$ (medium gray), and $N = 25$ (light gray).

The first thing to note from Figure 3 is that the measured amount of extended family insurance is weakly lower than the true amount of extended family insurance. If the true extended family insurance is zero, then there is no bias, but the bias increases in absolute terms as the true extended family insurance increases (i.e., moving left along the x-axis or jumping between the solid and dashed lines). The amount of this bias is also affected by the size of the network: as N increases (moving from dark to lightly shaded lines), the bias decreases. Concentrating on Panel A, in which the

measured group is missing a true group member, the worst case scenario represented is a 20% bias in the extended family insurance estimate for a group size of $N = 3$ and 100% true extended family insurance. Thus, for a true group size of $N = 3$, at worst our method could underestimate the amount of insurance by 20%. However, under this scenario, because we recover the true δ_I , which is closer to 0.5 in our data than 0.0, the extended family insurance bias is significantly smaller. In addition, because many families in our data have more than three households, the bias should be even smaller. Thus, while mismeasurement of groups can cause bias in our estimates of extended family insurance, we do not believe the bias is very large nor does it change our main conclusions.

7.3 Extended family shocks

In our main analysis, we have been agnostic about the origins of family-aggregate shocks, which could be due to small groups and/or because the income processes are genuinely correlated. We have, however, imposed the restriction that the MA coefficient (θ) is the same for idiosyncratic and family-aggregate shocks. Here we look closer at the implications of correlated shocks and consider two cases:

1. Benchmark case: $\Delta \log y_{i,j,t} = v_{i,j,t}$ in which v consist of entirely independent shocks across individuals and follow the usual permanent-transitory decomposition. As defined above, family-level shocks are $\overline{\Delta \log y_{i,j,t}} = \overline{v_{i,j,t}}$.
2. ‘‘Genuine’’ extended family shock case: $\Delta \log y_{i,j,t} = v_{i,j,t} + a_{j,t}$ in which v consist of entirely independent shocks across individuals while a consist of aggregate shocks to extended families that are independent across extended families but perfectly correlated within extended families. Both v and a follow the usual permanent-transitory decomposition. Family-level shocks now consist of two different processes, $\overline{\Delta \log y_{i,j,t}} = \overline{v_{i,j,t}} + a_{j,t}$.

The important distinction between these cases is that the distribution of ‘‘ v ’’ shocks and ‘‘ a ’’ shocks may be different, particularly the persistence of transitory shocks. Labelling family-aggregate components of shocks that arise from the ‘‘ v ’’ distribution with a V -superscript and those that arise from the ‘‘ a ’’ distribution with an A -superscript, the expanded income process is:

$$\Delta \log y_{i,j,t} = u_{j,t}^A + u_{j,t}^V + u_{i,j,t}^I + \Delta(e_{j,t}^A + \theta_A e_{j,t-1}^A) + \Delta(e_{j,t}^V + \theta e_{j,t-1}^V) + \Delta(e_{i,j,t}^I + \theta e_{i,j,t-1}^I) + \Delta m_{i,j,t}^y \quad (15)$$

Estimates from this model yield small but significant “genuine” components: 22% (25%) of family-aggregate permanent (transitory) shocks derive from a correlated shock distribution.²⁸ However, we cannot reject that the difference in persistence between the uncorrelated and correlated shock distributions is zero, and thus unsurprisingly the insurance estimates are very similar to our main estimates (estimates available upon request). In addition, this specification raises further caveats, including the fact that the observability of “genuine” and other shocks may vary, making it harder to interpret the extended family insurance estimates. Hence, our preferred specification remains the parsimonious specification.

8 Discussion and conclusions

Income shocks to households are substantial, even accounting for the redistributive effects of taxation and welfare programs. It is well understood that such fluctuations can have very large welfare effects (see for example Low, Meghir, and Pistaferri (2010)) and yet on average only about half of these fluctuations are insured.

A growing theoretical literature explores classes of underlying models that can generate this “partial” insurance. This literature typically examines risk-sharing arrangements within a well-defined group and under particular frictions, such as commitment problems (Kocherlakota, 1996; Ligon, Thomas, and Worrall, 2002) or moral hazard (Attanasio and Pavoni, 2011). Since the extended family may be able to solve (at least in part) these issues that may otherwise plague groups of households, one may expect the extended family network to provide some degree of insurance. Empirically, Hayashi, Altonji, and Kotlikoff (1996) and others have shown that complete extended family insurance does not occur, but they cannot say whether some, or any, occurs. On the other hand, other work has found evidence that extended family networks provide insurance through particular (in-kind) channels, such as coresidence (Kaplan, 2012).

In this paper, we follow up and extend the literature on partial insurance in a number of ways. In the first part of the paper, we present a framework of partial insurance and group risk sharing that decomposes income shocks into shocks that are aggregate to the group – and thus uninsurable by the group – and shocks that are idiosyncratic to particular group members, and

²⁸Identification of this additional income process requires additional cross-individual covariances; contact the authors for more details.

therefore insurable by its other members. This allows us to study the differential impact of group-aggregate and idiosyncratic shocks on consumption. Using covariance restrictions on household and group level income and consumption processes, we show how one can fully identify – including measurement error parameters – the income and insurance parameters. In the second part, we apply our framework to extended family networks in the United States. Exploiting the intergenerational structure of the Panel Study of Income Dynamics, our estimates of the income process imply that a substantial portion (around 60%) of income shocks are idiosyncratic within extended family networks. We argue that this implies a large potential for the extended family risk sharing network to have a non-trivial impact on the transmission of income shocks into consumption. However, our estimates of the transmission parameters show that no such insurance occurs on average. Thus, despite the fact that public welfare programs do not fully insure income risk (for example, the average replacement rate of unemployment insurance benefits in the US is around 60%), private networks do not replace them.

These findings beg the question of *why* extended families do not provide insurance for its members. Are families not as able to solve commitment and moral hazard problems as one might think? Are there social or cultural barriers to sharing risk? Developing insights into the frictions that prevent families from sharing risk is an interesting avenue for future research.

In conclusion, while we apply this framework to study extended families in the United States, we believe that it is more generally applicable to any group of households whom we suspect may collectively share risk, such as rural villagers in developing countries or other close-knit communities. By considering the difference between group-aggregate and idiosyncratic shocks to a network, the framework presented in this paper allows us to detect whether risk sharing takes place and to distinguish it from alternative theoretical benchmarks like complete markets and self-insurance.

References

- Abowd, J. M. and Card, D. (1989). ‘On the Covariance Structure of Earnings and Hours Changes.’ Econometrica, vol. 57(2):pages 411–445.
- Aguiar, M. and Hurst, E. (2005). ‘Consumption versus Expenditure.’ Journal of Political Economy, vol. 113(5):pages 919–948.
- Altonji, J., Hayashi, F., and Kotlikoff, L. (1992). ‘Is the Extended Family Altruistically Linked? Direct Tests Using Micro Data.’ American Economic Review, vol. 82(5):pages 1177–98.
- Altonji, J. G. and Segal, L. M. (1996). ‘Small-Sample Bias in GMM Estimation of Covariance Structures.’ Journal of Business & Economic Statistics, vol. 14(3):pages 353–366.
- Angelucci, M., De Giorgi, G., and Rasul, I. (2017). ‘Consumption and investment in resource pooling family networks.’ The Economic Journal.
- Attanasio, O. and Davis, S. J. (1996). ‘Relative Wage Movements and the Distribution of Consumption.’ Journal of Political Economy, vol. 104(6):pages 1227–62.
- Attanasio, O. P. and Borella, M. (2014). ‘Modeling Movements in Individual Consumption: A Time-Series Analysis of Grouped Data.’ International Economic Review, vol. 55(4):pages 959–991.
- Attanasio, O. P., Low, H., and Sanchez-Marcos, V. (2005). ‘Female Labor Supply As Insurance Against Idiosyncratic Risk.’ Journal of the European Economic Association, vol. 3(2-3):pages 755–764.
- Attanasio, O. P. and Pavoni, N. (2011). ‘Risk Sharing in Private Information Models With Asset Accumulation: Explaining the Excess Smoothness of Consumption.’ Econometrica, vol. 79(4):pages 1027–1068.
- Blau, D. and Currie, J. (2006). ‘Pre-school, day care, and after-school care: Who’s minding the kids?’ Handbook of the Economics of Education, vol. 2:pages 1163–1278.
- Blundell, R., Low, H., and Preston, I. (2013). ‘Decomposing changes in income risk using consumption data.’ Quantitative Economics, vol. 4(1):pages 1–37.
- Blundell, R., Pistaferri, L., and Preston, I. (2006). ‘Imputing consumption in the PSID using food demand estimates from the CEX.’ Working paper.
- Blundell, R., Pistaferri, L., and Preston, I. (2008). ‘Consumption Inequality and Partial Insurance.’ American Economic Review, vol. 98(5):pages 1887–1921.
- Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). ‘Consumption inequality and family labor supply.’ American Economic Review, vol. 106(2):pages 387–435.
- Blundell, R. and Preston, I. (1998). ‘Consumption inequality and income uncertainty.’ Quarterly Journal of Economics, vol. 113(2):pages 603–640.
- Boar, C. (2018). ‘Dynastic Precautionary Savings.’ Mimeo, Princeton University.
- Bound, J. and Krueger, A. B. (1991). ‘The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?’ Journal of Labor Economics, vol. 9(1):pages 1–24.

- Campbell, J. J. and Deaton, A. (1989). ‘Why Is Consumption So Smooth?’ The Review of Economic Studies, vol. 56(3):pages 357–374.
- Chiappori, P.-A. (1988). ‘Rational household labor supply.’ Econometrica: Journal of the Econometric Society, pages 63–90.
- Choi, H., McGarry, K., and Schoeni, R. F. (2016). ‘Consumption and the extended family.’ Economics Letters, vol. 140:pages 34–38.
- Cochrane, J. H. (1991). ‘A Simple Test of Consumption Insurance.’ Journal of Political Economy, vol. 99(5):pages 957–76.
- Cole, H. L. and Kocherlakota, N. R. (2001). ‘Efficient allocations with hidden income and hidden storage.’ The Review of Economic Studies, vol. 68(3):pages 523–542.
- Dubois, P., Jullien, B., and Magnac, T. (2008). ‘Formal And Informal Risk Sharing In LDCs: Theory And Empirical Evidence.’ Econometrica, vol. 76(4):pages 679–725.
- Edwards, K. (2015). ‘Who Helps the Unemployed? Young Workers’ Receipt of Private Cash Transfers.’ Working paper.
- Fafchamps, M. and Lund, S. (2003). ‘Risk-sharing networks in rural Philippines.’ Journal of development Economics, vol. 71(2):pages 261–287.
- Golosov, M., Tsyvinski, A., and Werning, I. (2007). ‘New dynamic public finance: a user’s guide.’ In NBER Macroeconomics Annual 2006, Volume 21, pages 317–388. MIT Press.
- Hall, P. and Horowitz, J. L. (1996). ‘Bootstrap critical values for tests based on generalized-method-of-moments estimators.’ Econometrica, vol. 64(4):pages 891–916.
- Hall, R. E. (1988). ‘Intertemporal Substitution in Consumption.’ Journal of Political Economy, vol. 96(2):pages 339–357.
- Hall, R. E. and Mishkin, F. S. (1982). ‘The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households.’ Econometrica, vol. 50(2):pages pp. 461–481.
- Hayashi, F., Altonji, J., and Kotlikoff, L. (1996). ‘Risk-Sharing between and within Families.’ Econometrica, vol. 64(2):pages pp. 261–294.
- Horowitz, J. L. (2001). ‘The bootstrap.’ Handbook of Econometrics, vol. 5:pages 3159–3228.
- Kaplan, G. (2012). ‘Moving back home: Insurance against labor market risk.’ Journal of Political Economy, vol. 120(2):pages 446–512.
- Kaplan, G. and Violante, G. L. (2010). ‘How Much Consumption Insurance Beyond Self-Insurance?’ American Economic Journal: Macroeconomics, vol. 2(3):pages 53–87.
- Karaivanov, A. and Townsend, R. M. (2014). ‘Dynamic Financial Constraints: Distinguishing Mechanism Design From Exogenously Incomplete Regimes.’ Econometrica, vol. 82(3):pages 887–959.
- Kinnan, C. and Townsend, R. (2012). ‘Kinship and financial networks, formal financial access, and risk reduction.’ American Economic Review, vol. 102(3):pages 289–93.

- Kocherlakota, N. R. (1996). ‘Implications of efficient risk sharing without commitment.’ The Review of Economic Studies, vol. 63(4):pages 595–609.
- Ligon, E., Thomas, J., and Worrall, T. S. (2002). ‘Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies.’ The Review of Economic Studies, vol. 69(1):pages 209–44.
- Low, H., Meghir, C., and Pistaferri, L. (2010). ‘Wage Risk and Employment Risk over the Life Cycle.’ American Economic Review, vol. 100:pages pp. 1432–1467.
- MaCurdy, T. E. (1983). ‘A simple scheme for estimating an intertemporal model of labor supply and consumption in the presence of taxes and uncertainty.’ International Economic Review, vol. 24(2):pages 265–289.
- McGarry, K. (2016). ‘Dynamic aspects of family transfers.’ Journal of Public Economics, vol. 137:pages 1 – 13.
- Meghir, C. and Pistaferri, L. (2004). ‘Income Variance Dynamics and Heterogeneity.’ Econometrica, vol. 72(1):pages pp. 1–32.
- Meyer, B. D. and Mok, W. K. (2006). ‘Calculating the Federal Income Tax Burden after 1993 Using TAXSIM—a Technical Appendix.’ Working paper.
- Meyer, B. D., Mok, W. K., and Sullivan, J. X. (2009). ‘The under-reporting of transfers in household surveys: Its nature and consequences.’ Tech. rep., National Bureau of Economic Research.
- Morten, M. (2017). ‘Temporary Migration and Endogenous Risk Sharing in Village India.’ Journal of Political Economy, vol. Forthcoming.
- Rosenzweig, M. and Stark, O. (1989). ‘Consumption Smoothing, Migration, and Marriage: Evidence From Rural India.’ Journal of Political Economy, vol. 97:pages 905–926.
- Rosenzweig, M. R. and Wolpin, K. I. (1993). ‘Intergenerational support and the life-cycle incomes of young men and their parents: Human capital investments, coresidence, and intergenerational financial transfers.’ Journal of Labor Economics, vol. 11(1):pages 84–112.
- Thomas, D. (1990). ‘Intra-household resource allocation: An inferential approach.’ Journal of human resources, pages 635–664.
- Thomas, J. and Worrall, T. S. (1988). ‘Self-enforcing Wage Contracts.’ The Review of Economic Studies, vol. 55(4):pages 541–54.
- Townsend, R. M. (1994). ‘Risk and Insurance in Village India.’ Econometrica, vol. 62(3):pages 539–91.

Appendix A: Identification

To estimate the income and consumption parameters of our model, we minimize the distance between data covariances and a combination of parameters specified by our model. We use the following moments for identification and estimation, where the left-hand side is data and the right-hand side consists of the parameters of the model.

1. Income

Household autocovariances

$$\text{cov}(\Delta y_{i,j,t}, \Delta y_{i,j,t}) = \text{var}(u_t^I) + \text{var}(e_t^I) + (\theta - 1)^2 \text{var}(e_{t-1}^I) + \theta^2 \text{var}(e_{t-2}^I) \quad (16)$$

$$+ \text{var}(u_t^F) + \text{var}(e_t^F) + (\theta - 1)^2 \text{var}(e_{t-1}^F) + \theta^2 \text{var}(e_{t-2}^F) + 2\text{var}(m_y)$$

$$\text{cov}(\Delta y_{i,j,t}, \Delta y_{i,j,t+1}) = (\theta - 1)\text{var}(e_t^I) - \theta(\theta - 1)\text{var}(e_{t-1}^I) \quad (17)$$

$$+ (\theta - 1)\text{var}(e_t^F) - \theta(\theta - 1)\text{var}(e_{t-1}^F) - \text{var}(m_y)$$

$$\text{cov}(\Delta y_{i,j,t}, \Delta y_{i,j,t+2}) = -\theta\text{var}(e_t^I) - \theta\text{var}(e_t^F) \quad (18)$$

Extended family-average autocovariances

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t}}) = \text{var}(u_t^F) + \text{var}(e_t^F) + (\theta - 1)^2 \text{var}(e_{t-1}^F) + \theta^2 \text{var}(e_{t-2}^F) + \frac{2}{n_j} \text{var}(m_y) \quad (19)$$

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t+1}}) = (\theta - 1)\text{var}(e_t^F) - \theta(\theta - 1)\text{var}(e_{t-1}^F) - \frac{1}{n_j} \text{var}(m_y) \quad (20)$$

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t+2}}) = -\theta\text{var}(e_t^F) \quad (21)$$

Identification of the income parameters follows from a combination of these 6 moments:

1. From (21) we know $\text{var}(e_t^F) = -\frac{\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t+2}})}{\theta}$ and plugging that into (18) we know $\text{var}(e_t^I) = \frac{\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t+2}})}{\theta} - \frac{\text{cov}(\Delta y_{j,t}, \Delta y_{j,t+2})}{\theta}$
2. Subtracting n_j times (20) from (17) and plugging the formulas for $\text{var}(e_t^F)$ and $\text{var}(e_t^I)$ into this identifies θ from the following implicit equation:

$$\begin{aligned} \text{cov}(\Delta y_{i,j,t}, \Delta y_{i,j,t+1}) - n_j \text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t+1}}) &= (\theta - 1)\text{var}(e_t^I) - \theta(\theta - 1)\text{var}(e_{t-1}^I) \\ &\quad - (n_j - 1)(\theta - 1)\text{var}(e_t^F) + (n_j - 1)\theta(\theta - 1)\text{var}(e_{t-1}^F) \end{aligned}$$

1. With θ identified, (21) identifies $\text{var}(e_t^F)$ and consequently (18) identifies $\text{var}(e_t^I)$

2. (20) or (17) then identifies $\text{var}(m_y)$.
3. (19) identifies $\text{var}(u_t^F)$.
4. (16) identifies $\text{var}(u_t^I)$.

Measurement error in income, which is notoriously difficult to identify separately from transitory shocks, can be isolated in this model. In most income models, measurement error is not separately identified from transitory shocks because the variance of measurement error co-moves one-for-one with the variance of the transitory shock. Our model puts structure on the transitory shocks that inadvertently allows us to separately identify these shocks from measurement error (as well as from each other). In the household autocovariances (equations (16) through (18)), measurement error co-moves one-for-one with transitory shocks (similarly to most income models). In family-average autocovariances (equations (19) through (21)), it does not. The reason for this is two-fold. First, because the sum of all transitory idiosyncratic shocks of family members is zero by definition, transitory idiosyncratic shocks do not move one-for-one with measurement error. Second, each family member receives the same transitory family-aggregate shock but receives different measurement error realizations. Thus, since the variance of an average is equal to $1/n_j$ of the variance of a variable, the variance of measurement error enters $1/n_j$ -to-one with the variance of the transitory family-aggregate shock, allowing us to separately identify the two parameters.²⁹

2. Consumption

Identification of consumption parameters δ_I , δ_F , γ_I , γ_F , $\text{var}(\xi)$, and $\text{var}(m_{c,t})$ follows from a combination of consumption autocovariances and income-consumption covariances.

²⁹Note that we are under-identified if the transitory component does not have persistence ($\theta = 0$) and we restrict identification to solely income covariances. However, all income parameters would still be identified if we included consumption covariances (contact authors for the identification proof).

Household autocovariances

$$\begin{aligned} \text{cov}(\Delta c_{i,j,t}, \Delta c_{i,j,t}) &= \delta_I^2 \text{var}(u_t^I) + \gamma_I^2 (1 + \theta)^2 \text{var}(e_t^I) + \delta_F^2 \text{var}(u_t^F) + \gamma_F^2 (1 + \theta)^2 \text{var}(e_t^F) \\ &\quad + \text{var}(m_{c,t}) + \text{var}(m_{c,t-1}) + \text{var}(\xi) \end{aligned} \quad (22)$$

$$\text{cov}(\Delta c_{i,j,t}, \Delta c_{i,j,t+1}) = -\text{var}(m_{c,t}) \quad (23)$$

$$\text{cov}(\Delta c_{i,j,t}, \Delta y_{i,j,t}) = \delta_I \text{var}(u_t^I) + \gamma_I (1 + \theta) \text{var}(e_t^I) + \delta_F \text{var}(u_t^F) + \gamma_F (1 + \theta) \text{var}(e_t^F) \quad (24)$$

$$\text{cov}(\Delta c_{i,j,t}, \Delta y_{i,j,t+1}) = \gamma_I (1 + \theta) (\theta - 1) \text{var}(e_t^I) + \gamma_F (1 + \theta) (\theta - 1) \text{var}(e_t^F) \quad (25)$$

$$\text{cov}(\Delta c_{i,j,t}, \Delta y_{i,j,t+2}) = -\gamma_I \theta (1 + \theta) \text{var}(e_t^I) - \gamma_F (1 + \theta) \theta \text{var}(e_t^F) \quad (26)$$

Family-average autocovariances

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta c_{j,t}}) = \delta_F^2 \text{var}(u_t^F) + \gamma_F^2 (1 + \theta)^2 \text{var}(e_t^F) + \frac{1}{n_i} (\text{var}(m_{c,t}) + \text{var}(m_{c,t-1}) + \text{var}(\xi)) \quad (27)$$

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta c_{j,t+1}}) = -\frac{1}{n_i} \text{var}(m_{c,t}) \quad (28)$$

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta y_{j,t}}) = \delta_F \text{var}(u_t^F) + \gamma_F (1 + \theta) \text{var}(e_t^F) \quad (29)$$

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta y_{j,t+1}}) = \gamma_F (1 + \theta) (\theta - 1) \text{var}(e_t^F) \quad (30)$$

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta y_{j,t+2}}) = -\gamma_F \theta (1 + \theta) \text{var}(e_t^F) \quad (31)$$

Identification of consumption parameters is as follows:³⁰

1. (23) or (28) identifies $\text{var}(m_{c,t})$.
2. (30) or (31) identifies γ_F .
3. (25) or (26) identifies γ_I .
4. (29) identifies δ_F .
5. (24) identifies δ_I .
6. Either of (22) or (27) identifies $\text{var}(\xi)$.

Appendix B: Estimating family-average moments

Family-average moments contain a multiplicative factor of $1/n_j$ that is not constant between families because families are not all the same size. As an example, take the variance of family-average

³⁰Note that we are over-identified. It is theoretically possible to add more richness, such as allowing for the possibility of correlation between income and consumption measurement errors. We abstract from this addition due to data concerns. Contact authors for an identification proof with correlated measurement error.

income:

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t}}) = \text{var}(u_t^F) + \text{var}(e_t^F) + (\theta - 1)^2 \text{var}(e_{t-1}^F) + \theta^2 \text{var}(e_{t-2}^F) + \frac{2}{n_j} \text{var}(m_y)$$

From this equation it is easy to see that the variance of family-average income varies by extended family size. This means that the distribution of income is a mixture distribution in which the components are defined by the size of the extended family. In our estimation procedure, we must modify the moments to account for this.

Let the mean of family-average income equal μ and let w_s be the proportion of families of size s . Then we know that the variance of family-average income is equal to

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t}}) = \sum_s w_s [(\mu_s - \mu)^2 + \sigma_s^2]$$

where μ_s and σ_s^2 are the mean and variance of family-average income for families of size s . Hence the simple modification for estimation is:

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t}}) - \sum_s w_s (\mu_s - \mu)^2 = \text{var}(u_t^F) + \text{var}(e_t^F) + (\theta - 1)^2 \text{var}(e_{t-1}^F) + \theta^2 \text{var}(e_{t-2}^F) + \sum_s w_s \cdot \frac{2}{s} \text{var}(m_y)$$

Analogous logic follows for covariances of both income and consumption.

Appendix C: Long-difference moments

In 1999, the PSID switched from interviewing households every year to every other year. This will affect our model for two reasons. First, we measure income and consumption at the yearly level, so we are missing every other year of data from 1999 onwards. Second, our model is dynamic and income is persistent, so income shocks that occur in off-survey years have repercussions to both income and consumption in future years. To account for this, we define “long-differences” (with the notation Δ_2) as the two-year difference in income and consumption and rewrite covariances that factors in this change.

Income long-difference

$$\Delta_2 \log y_{i,j,t} \equiv \log y_{i,j,t} - \log y_{i,j,t-2} = u_{j,t}^F + u_{j,t-1}^F + u_{i,j,t}^I + u_{i,j,t-1}^I + \Delta_2(e_{j,t}^F + \theta e_{j,t-1}^F) + \Delta_2(e_{i,j,t}^I + \theta e_{i,j,t-1}^I) + \Delta_2 m_{i,j,t}^y$$

We will not be able to separately identify $u_{i,j,t}^S$ from $u_{i,j,t-1}^S$ for $S \in F, I$, so we will assign $\text{var}(u_{t-1}^S) = \text{var}(u_t^S)$. Same goes for $\text{var}(e_{t-1}^S)$.

$$\begin{aligned} \text{cov}(\Delta_2 y_{i,j,t}, \Delta_2 y_{i,j,t}) &= 2\text{var}(u_t^I) + (1 + \theta^2)\text{var}(e_t^I) + (1 + \theta^2)\text{var}(e_{t-2}^I) \\ &\quad + 2\text{var}(u_t^F) + (1 + \theta^2)\text{var}(e_t^F) + (1 + \theta^2)\text{var}(e_{t-2}^F) + 2\text{var}(m_y) \end{aligned} \quad (32)$$

$$\text{cov}(\Delta_2 y_{i,j,t}, \Delta_2 y_{i,j,t+2}) = -(1 + \theta^2)\text{var}(e_t^I) - (1 + \theta^2)\text{var}(e_t^F) - \text{var}(m_y) \quad (33)$$

$$\text{cov}(\overline{\Delta_2 y_{j,t}}, \overline{\Delta_2 y_{j,t}}) = 2\text{var}(u_t^F) + (1 + \theta^2)\text{var}(e_t^F) + (1 + \theta^2)\text{var}(e_{t-2}^F) + \frac{2}{n_j} \text{var}(m_y) \quad (34)$$

$$\text{cov}(\overline{\Delta_2 y_{j,t}}, \overline{\Delta_2 y_{j,t+2}}) = -(1 + \theta^2)\text{var}(e_t^F) - \frac{1}{n_j} \text{var}(m_y) \quad (35)$$

We cannot identify the MA(1) parameter (θ) or income measurement error ($\text{var}(m_y)$) using only long-difference moments. In estimation, since we do not allow θ or $\text{var}(m_y)$ to vary over time, identification of these parameters comes from the short-difference covariances. Identification of the income shock variances is then straightforward: (35) identifies $\text{var}(e_t^F)$, (34) identifies $\text{var}(u_t^F)$, (33) identifies $\text{var}(e_t^I)$, and (32) identifies $\text{var}(u_t^I)$.

Consumption long-difference

$$\begin{aligned} \Delta_2 \log c_{i,j,t} \equiv \log c_{i,j,t} - \log c_{i,j,t-2} &= \delta_I(u_{i,j,t}^I + u_{i,j,t-1}^I) + \gamma_I(1 + \theta)(e_{i,j,t}^I + e_{i,j,t-1}^I) \\ &\quad + \delta_F(u_{j,t}^F + u_{j,t-1}^F) + \gamma_F(1 + \theta)(e_{j,t}^F + e_{j,t-1}^F) + \Delta_2 m_{i,j,t}^c + \xi_{i,j,t} + \xi_{i,j,t-1} \end{aligned}$$

Analogously to the income shock variances, we cannot separately identify $\xi_{i,j,t}$ from $\xi_{i,j,t-1}$, hence we assign $\text{var}(\xi_{i,j,t-1}) = \text{var}(\xi_{i,j,t})$

$$\begin{aligned} \text{cov}(\Delta_2 c_{i,j,t}, \Delta_2 c_{i,j,t}) &= 2\delta_I^2 \text{var}(u_t^I) + 2\delta_F^2 \text{var}(u_t^F) + 2\gamma_I^2(1+\theta)^2 \text{var}(e_t^I) + 2\gamma_F^2(1+\theta)^2 \text{var}(e_t^F) \\ &\quad + \text{var}(m_{c,t}) + \text{var}(m_{c,t-2}) + 2\text{var}(\xi) \end{aligned} \quad (36)$$

$$\text{cov}(\Delta_2 c_{i,j,t}, \Delta_2 c_{i,j,t+2}) = -\text{var}(m_{c,t}) \quad (37)$$

$$\text{cov}(\Delta_2 y_{i,t}, \Delta_2 c_{i,t}) = 2\delta_I \text{var}(u_t^I) + 2\delta_F \text{var}(u_t^F) + \gamma_I(1+\theta)^2 \text{var}(e_t^I) + \gamma_F(1+\theta)^2 \text{var}(e_t^F) \quad (38)$$

$$\text{cov}(\Delta_2 y_{i,t+2}, \Delta_2 c_{i,t}) = -\gamma_I(1+\theta)^2 \text{var}(e_t^I) - \gamma_F(1+\theta)^2 \text{var}(e_t^F) \quad (39)$$

$$\text{cov}(\overline{\Delta_2 c_{i,t}}, \overline{\Delta_2 c_{i,t}}) = 2\delta_F^2 \text{var}(u_t^F) + 2\gamma_F^2(1+\theta)^2 \text{var}(e_t^F) + \frac{1}{n_i} (\text{var}(m_{c,t}) + \text{var}(m_{c,t-2}) + 2\text{var}(\xi)) \quad (40)$$

$$\text{cov}(\overline{\Delta_2 c_{i,t}}, \overline{\Delta_2 c_{i,t+2}}) = -\frac{1}{n_i} \text{var}(m_{c,t}) \quad (41)$$

$$\text{cov}(\overline{\Delta_2 y_{i,t}}, \overline{\Delta_2 c_{i,t}}) = 2\delta_F \text{var}(u_t^F) + \gamma_F(1+\theta)^2 \text{var}(e_t^F) \quad (42)$$

$$\text{cov}(\overline{\Delta_2 y_{i,t+2}}, \overline{\Delta_2 c_{i,t}}) = -\gamma_F(1+\theta)^2 \text{var}(e_t^F) \quad (43)$$

Identification

Identification of consumption parameters using long-difference covariances is analogous to the short difference covariances: (37) or (41) identifies $\text{var}(m_{c,t})$, (43) identifies γ_F , (42) identifies δ_F , (39) identifies γ_I , (38) identifies δ_I , and finally (36) or (40) identifies $\text{var}(\xi)$.

Appendix D: Details on the PSID income and expenditure data

As summarized in Section 5.1, our main data source is the Panel Study of Income Dynamics (PSID) from 1979 through 2011. In this appendix we provide more details about our income and consumption measures and how they have evolved over time (see Blundell, Pistaferri, and Preston (2008) for details on the CEX data).

Income Our measure of household income is annual post tax-and-transfer income, excluding transfers from family and friends. A household consists of a head and potentially a spouse and other family members that live in the household. To construct this measure, we begin with the PSID's summary variable for total (nuclear) family money income from the previous tax year, which consists of head and spouse taxable income, head and spouse transfer income, and other (nuclear) family unit members' taxable income and transfer income. This summary variable was bottom coded at \$1 prior to 1994, so we continue to bottom code this variable in later years for consistency (this applies to only 12 observations in our sample). When the PSID became biennial in 1999, additional income questions were asked about the gap years, but we do not incorporate this data since we do not have consumption in gap years and other work has shown that recall bias is large for these years.

Taxable income consists of income from wages and salary, assets, and business and farm profits, and was collected for the head, spouse, and other (nuclear) family members. Transfer income consists of alimony, annuity income, child support, help from relatives and non-relatives, Supplemental Security income, Social Security old-age and disability income, unemployment insurance benefits, worker's compensation, welfare/TANF, retirement income, VA pension income, and other transfer income. Total transfer income is collected for the head, spouse, and other (nuclear) family members, but measures of each benefit are not always collected separately in all years (for a breakdown of transfer questions by year and household member see Meyer, Mok, and Sullivan (2009)). From this measure of total income, we subtract alimony and child support *given* to others since alimony and child support *received* is included in our measure of income, and we subtract transfers from relatives and non-relatives. Note that questions about alimony and child support were only explicitly asked beginning in the 1985 survey and only for the head and spouse, and transfers from relatives

were not collected for other (nuclear) family members between 1994 and 2003 while transfers from non-relatives were collected beginning in 1994 and only for the head and spouse. We also add the monetary value of food stamps to our measure of household income.

We then subtract federal taxes from this income measure to obtain a post tax and (public) transfers measure of income. Prior to 1991, the PSID provided computed federal taxes. From 1991 forward, we imputed federal taxes using the NBER's TAXSIM program, following the assumptions used by Meyer and Mok (2006) to link information in the PSID to the inputs needed by TAXSIM. For the 45 observations in which TAXSIM does not provide a federal tax burden, we impute their federal tax burden using a quintic in income, a quadratic in (nuclear) family size, number of children, whether they are self-employed, and asset income. We drop income observations in which income is truncated (132 observations) or outlier federal tax burdens (25 observations), and trim the bottom and top 1% of the income distribution. Income is deflated by the CPI-Urban deflator.

Consumption We use three measures of non-durable consumption. The first two stem from weekly food expenditures, which we aggregate to the yearly level. The reference period for food expenditure is not clear, so we assume that it is the same as the reference period for income, as previous work has done. Food consumption is defined as expenditures on food eaten inside and outside the home as well as the dollar value of food stamps. Prior to 1993, non-food stamp expenditure was collected separately for food consumed inside and outside the house, but afterwards it separated out food delivered. We use this measure directly and also to construct a measure of total consumption, which is imputed from food consumption using estimates from a food demand system estimated on the CEX (see Blundell, Pistaferri, and Preston (2008) for more details).

Starting in 1999, the PSID began collecting information on additional consumption expenditures, including transportation expenses (gasoline, vehicle insurance, vehicle repairs, parking, bus fares, taxi fares, and other transportation expenses), childcare, utilities (electricity, heating, water, miscellaneous utilities), and home insurance. The PSID also collected data on health, education, and rent, but we exclude these categories to be more consistent with our imputed consumption definition, which is intended to capture non-durable consumption. In constructing this "reported consumption" measure, missing categories are set to zero, unless all categories are missing, in which case reported consumption is set to zero. Similar to income, for all consumption measures we trim

the bottom and top 1% of their respective distributions. Reported and imputed consumption are deflated by the CPI-Urban deflator, and food consumption is deflated by the food CPI deflator.

Appendix Tables

Appendix Table 1: Autocovariance Matrix of Income Growth

Year	Individual Autocovariances			Family-average Autocovariances		
	$\text{var}(\Delta y_t)$	$\text{cov}(\Delta y_{t+1}, \Delta y_t)$	$\text{cov}(\Delta y_{t+2}, \Delta y_t)$	$\text{var}(\overline{\Delta y_t})$	$\text{cov}(\overline{\Delta y_{t+1}}, \overline{\Delta y_t})$	$\text{cov}(\overline{\Delta y_{t+2}}, \overline{\Delta y_t})$
1979	0.102 (0.007)	-0.029 (0.006)	-0.006 (0.006)	0.042 (0.003)	-0.011 (0.002)	0.001 (0.002)
1980	0.100 (0.010)	-0.039 (0.008)	-0.002 (0.009)	0.038 (0.003)	-0.012 (0.002)	-0.005 (0.002)
1981	0.111 (0.009)	-0.043 (0.008)	0.007 (0.007)	0.042 (0.003)	-0.013 (0.002)	0.002 (0.002)
1982	0.115 (0.009)	-0.044 (0.008)	-0.002 (0.005)	0.043 (0.003)	-0.014 (0.002)	-0.002 (0.002)
1983	0.105 (0.010)	-0.032 (0.006)	-0.010 (0.008)	0.037 (0.003)	-0.010 (0.002)	-0.007 (0.002)
1984	0.105 (0.009)	-0.029 (0.007)	0.001 (0.006)	0.034 (0.002)	-0.006 (0.002)	0.001 (0.001)
1985	0.128 (0.012)	-0.034 (0.007)	-0.010 (0.006)	0.040 (0.003)	-0.008 (0.002)	-0.001 (0.002)
1986	0.119 (0.009)	-0.041 (0.006)	-0.014 (0.007)	0.045 (0.003)	-0.015 (0.002)	-0.004 (0.002)
1987	0.102 (0.008)	-0.024 (0.006)	-0.009 (0.005)	0.036 (0.002)	-0.008 (0.002)	-0.003 (0.002)
1988	0.106 (0.009)	-0.034 (0.007)	-0.011 (0.005)	0.036 (0.003)	-0.013 (0.002)	-0.002 (0.002)
1989	0.097 (0.009)	-0.028 (0.006)	-0.004 (0.005)	0.031 (0.003)	-0.008 (0.002)	-0.001 (0.002)
1990	0.110 (0.009)	-0.046 (0.008)	0.011 (0.007)	0.038 (0.003)	-0.015 (0.002)	0.006 (0.003)
1991	0.114 (0.010)	-0.042 (0.009)	0.004 (0.007)	0.034 (0.003)	-0.013 (0.002)	-0.005 (0.002)
1992	0.127 (0.011)	-0.066 (0.009)	-0.016 (0.009)	0.042 (0.003)	-0.016 (0.002)	-0.008 (0.003)
1993	0.164 (0.014)	-0.071 (0.012)	-0.000 (0.008)	0.051 (0.004)	-0.021 (0.003)	0.002 (0.004)
1994	0.206 (0.023)	-0.068 (0.013)	-0.018 (0.013)	0.076 (0.007)	-0.025 (0.004)	-0.012 (0.004)
1995	0.193 (0.024)	-0.071 (0.019)	NA	0.064 (0.006)	-0.027 (0.007)	NA
1996	0.185 (0.020)	NA	NA	0.075 (0.006)	NA	NA
1998	0.205 (0.022)	-0.079 (0.016)	NA	0.077 (0.008)	-0.029 (0.004)	NA
2000	0.186 (0.019)	-0.043 (0.013)	NA	0.076 (0.007)	-0.027 (0.006)	NA
2002	0.187 (0.028)	-0.069 (0.018)	NA	0.070 (0.007)	-0.029 (0.006)	NA
2004	0.164 (0.018)	-0.056 (0.010)	NA	0.075 (0.007)	-0.023 (0.004)	NA
2006	0.152 (0.013)	-0.044 (0.009)	NA	0.066 (0.005)	-0.021 (0.004)	NA
2008	0.163 (0.014)	-0.074 (0.010)	NA	0.064 (0.004)	-0.032 (0.004)	NA
2010	0.158 (0.013)	NA	NA	0.064 (0.005)	NA	NA

Note: This table reports variances (columns (1) and (4)), first-order autocovariances (columns (2) and (5)) and second-order autocovariances (columns (3) and (6)) of the income process where idiosyncratic household and family-aggregate components are defined as in Section 3.3. The sample includes “families” in the 1979-2011 PSID, where we define “families” as a cohabiting couple under 65 and their adult children over 25. Income is defined as total income net of taxes and family transfers. Standard errors in parentheses are based on 100 block bootstrap replications, clustered by extended family.

Appendix Table 2: Autocovariance Matrix of Consumption Growth

Year	Autocovariances		Family-average Autocovariances	
	$\text{var}(\Delta c_t)$	$\text{cov}(\Delta c_{t+1}, \Delta c_t)$	$\text{var}(\overline{\Delta c_t})$	$\text{cov}(\overline{\Delta c_{t+1}}, \overline{\Delta c_t})$
1981	0.345 (0.028)	-0.132 (0.019)	0.118 (0.009)	-0.042 (0.005)
1982	0.310 (0.027)	-0.132 (0.021)	0.111 (0.008)	-0.035 (0.006)
1983	0.315 (0.026)	-0.115 (0.015)	0.105 (0.007)	-0.033 (0.004)
1984	0.293 (0.022)	-0.132 (0.016)	0.106 (0.007)	-0.047 (0.006)
1985	0.338 (0.028)	-0.142 (0.026)	0.114 (0.008)	-0.039 (0.007)
1986	0.327 (0.027)	NA	0.108 (0.008)	NA
1990	0.313 (0.023)	-0.113 (0.015)	0.080 (0.004)	-0.027 (0.003)
1991	0.309 (0.021)	-0.155 (0.020)	0.096 (0.005)	-0.053 (0.005)
1992	0.345 (0.028)	-0.135 (0.022)	0.127 (0.010)	-0.042 (0.007)
1993	0.327 (0.031)	-0.154 (0.031)	0.115 (0.009)	-0.048 (0.008)
1994	0.314 (0.034)	-0.118 (0.025)	0.111 (0.012)	-0.039 (0.013)
1995	0.343 (0.041)	-0.177 (0.045)	0.139 (0.016)	-0.062 (0.015)
1996	0.343 (0.045)	NA	0.135 (0.015)	NA
1998	0.355 (0.030)	-0.149 (0.028)	0.114 (0.009)	-0.060 (0.010)
2000	0.414 (0.034)	-0.163 (0.028)	0.154 (0.012)	-0.057 (0.010)
2002	0.445 (0.046)	-0.223 (0.056)	0.165 (0.013)	-0.070 (0.016)
2004	0.540 (0.058)	-0.232 (0.044)	0.216 (0.019)	-0.101 (0.016)
2006	0.472 (0.038)	-0.171 (0.028)	0.206 (0.015)	-0.055 (0.010)
2008	0.549 (0.041)	NA	0.199 (0.015)	NA

Note: This table reports variances (columns (1) and (4)), first-order autocovariances (columns (2) and (5)) and second-order autocovariances (columns (3) and (6)) of the income process where idiosyncratic household and family-aggregate components are defined as in Section 3.3. The sample includes “families” in the 1979-2011 PSID, where we define “families” as a cohabiting couple under 65 and their adult children over 25. Consumption is imputed and is derived from a food demand system estimated on the CEX (see 5.2). Standard errors in parentheses are based on 100 block bootstrap replications, clustered by extended family.

Appendix Table 3: Income Parameter Estimates

	Variance of Permanent Shocks			Variance of Transitory Shocks		
	$\text{var}(u_{i,j,t}^I)$ (idiosyncratic)	$\text{var}(u_{i,j,t}^F)$ (family)	% insurable by family	$\text{var}(e_{i,j,t}^I)$ (idiosyncratic)	$\text{var}(e_{i,j,t}^F)$ (family)	% insurable by family
1979-81	0.014 (0.006)	0.014 (0.003)	0.495 (0.128)	0.018 (0.010)	0.009 (0.005)	0.673 (0.163)
1982	0.015 (0.008)	0.013 (0.003)	0.540 (0.141)	0.026 (0.012)	0.012 (0.005)	0.674 (0.166)
1983	0.012 (0.006)	0.009 (0.003)	0.556 (0.148)	0.020 (0.010)	0.008 (0.004)	0.710 (0.138)
1984	0.018 (0.006)	0.016 (0.003)	0.529 (0.091)	0.019 (0.010)	0.002 (0.004)	0.905 (0.136)
1985	0.031 (0.009)	0.024 (0.003)	0.559 (0.073)	0.024 (0.009)	0.005 (0.004)	0.836 (0.100)
1986	0.013 (0.006)	0.019 (0.003)	0.410 (0.130)	0.024 (0.008)	0.014 (0.005)	0.639 (0.072)
1987	0.016 (0.007)	0.011 (0.003)	0.603 (0.145)	0.013 (0.007)	0.007 (0.004)	0.654 (0.093)
1988	0.023 (0.008)	0.011 (0.003)	0.667 (0.130)	0.018 (0.009)	0.011 (0.005)	0.623 (0.102)
1989	0.018 (0.005)	0.009 (0.003)	0.675 (0.103)	0.015 (0.009)	0.005 (0.004)	0.734 (0.152)
1990	0.019 (0.006)	0.012 (0.003)	0.611 (0.086)	0.022 (0.013)	0.012 (0.006)	0.644 (0.151)
1991	0.022 (0.008)	0.003 (0.003)	0.872 (0.097)	0.022 (0.013)	0.013 (0.005)	0.623 (0.193)
1992	0.000 (0.007)	0.007 (0.003)	0.000 (0.358)	0.051 (0.014)	0.016 (0.005)	0.761 (0.089)
1993	0.004 (0.012)	0.008 (0.003)	0.365 (0.314)	0.053 (0.017)	0.022 (0.008)	0.707 (0.080)
1994-96	0.022 (0.013)	0.013 (0.007)	0.619 (0.196)	0.047 (0.011)	0.028 (0.006)	0.626 (0.057)
1998-00	0.029 (0.010)	0.010 (0.005)	0.747 (0.119)	0.020 (0.010)	0.022 (0.007)	0.472 (0.121)
2002	0.024 (0.010)	0.006 (0.004)	0.790 (0.177)	0.031 (0.015)	0.022 (0.008)	0.579 (0.094)
2004	0.007 (0.006)	0.012 (0.004)	0.375 (0.217)	0.024 (0.010)	0.016 (0.009)	0.598 (0.091)
2006	0.013 (0.003)	0.010 (0.003)	0.566 (0.101)	0.017 (0.008)	0.017 (0.007)	0.492 (0.119)
2008-10	0.013 (0.004)	0.004 (0.002)	0.748 (0.117)	0.028 (0.010)	0.024 (0.007)	0.544 (0.062)
θ (Serial correlation of transitory shock)			0.163	(0.142)		
Income measurement error variance			0.015	(0.011)		

Note: This table reports year-by-year estimates of the income process where idiosyncratic household and family-aggregate components are defined as in Section 3.3. Columns (1)-(3) report estimates for permanent shocks and columns (5)-(7) report estimates for transitory shocks. Columns (3) and (7) report the percent of the total income shock that is potentially insurable by the extended family. The sample includes “families” in the 1979-2011 PSID, where we define “families” as a cohabiting couple under 65 and their adult children over 25. Income is defined as total income net of taxes and family transfers. Standard errors in parentheses are based on 100 block bootstrap replications, clustered by extended family.

Appendix Table 4: Moments associated with mismeasuring the group

Group type	$\text{var}(\Delta \log y_{i,j,t})$	$\text{var}(\overline{\Delta \log y_{i,j,t}})$	$\text{var}(\Delta \log c_{i,j,t})$	$\text{var}(\overline{\Delta \log c_{i,j,t}})$
True group	$\text{var}(u^F) + \text{var}(u^I)$	$\text{var}(u^F)$	$\delta_F^2 \text{var}(u^F) + \delta_I^2 \text{var}(u^I)$	$\delta_F^2 \text{var}(u^F)$
Missing member	$\frac{N}{N-1} \text{var}(u^F) + \frac{N(N-2)}{(N-1)^2} \text{var}(u^I)$	$\frac{N}{N-1} \text{var}(u^F)$	$\delta_F^2 \text{var}(u^F) + \delta_I^2 \text{var}(u^I)$	$\delta_F^2 \text{var}(u^F) + \delta_I^2 \frac{1}{(N-1)^2} \text{var}(u^I)$
Extra member	$\frac{N}{N+1} \text{var}(u^F) + \frac{N^2}{(N+1)(N-1)} \text{var}(u^I)$	$\frac{N}{N+1} \text{var}(u^F)$	$\delta_F^2 \text{var}(u^F) + \delta_I^2 \text{var}(u^I)$	$\delta_F^2 \frac{1+N^2}{(1+N)^2} \text{var}(u^F) + \delta_I^2 \frac{1}{(1+N)^2} \text{var}(u^I)$
Wrong member	$\text{var}(u^F) + \text{var}(u^I)$	$\text{var}(u^F)$	$\delta_F^2 \text{var}(u^F) + \delta_I^2 \text{var}(u^I)$	$\delta_F^2 \frac{(N-1)^2+1}{N^2} \text{var}(u^F) + \delta_I^2 \frac{2}{N^2} \text{var}(u^I)$