

NBER WORKING PAPER SERIES

INSURANCE IN EXTENDED FAMILY NETWORKS

Orazio Attanasio
Costas Meghir
Corina Mommaerts

Working Paper 21059
<http://www.nber.org/papers/w21059>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2015

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 and the ARIA 2014 Annual Meeting. We thank Luigi Pistaferri for generously providing imputed consumption data. Support from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy (CPP) at the IFS, reference RES-544-28-0001, is gratefully acknowledged. Orazio Attanasio thanks the ESRC Professorial Fellowship ES/K010700/1 and the ESRC CPP centre 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Orazio Attanasio, Costas Meghir, and Corina Mommaerts. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Insurance in extended family networks
Orazio Attanasio, Costas Meghir, and Corina Mommaerts
NBER Working Paper No. 21059
April 2015
JEL No. D12,D31,D91,E21,E24,H31

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 is insured, having accounted for public insurance programs. We apply our framework to extended family networks in the United States by exploiting the unique intergenerational structure of the PSID. We find that over 60% of shocks to household income are potentially insurable within family networks. However, we find little evidence that the extended family provides insurance for such idiosyncratic shocks.

Orazio Attanasio
Department of Economics
University College London
Gower Street
London WC1E 6BT
UNITED KINGDOM
and NBER
o.attanasio@ucl.ac.uk

Corina Mommaerts
Department of Economics
Yale University
37 Hillhouse Av
New Haven, CT 06511
corina.mommaerts@yale.edu

Costas Meghir
Department of Economics
Yale University
37 Hillhouse Avenue
New Haven, CT 06511
and IZA
and also NBER
c.meghir@yale.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. The standard life cycle model of individual behaviour, in which households are endowed with a concave utility function, posits a strong incentive for intertemporal smoothing of income changes, both at low and high frequency. The concavity of the utility function implies that households will prefer a smooth consumption to a variable one, given the level of interest rates and discount factor. How much smoothing a household can achieve depends on the instruments at their disposal for such smoothing.

In the standard life cycle model, it is often posited that households can save and possibly borrow using financial assets, which pay an interest rate that is possibly uncertain. More generally, the nature of the assets that individual households can access determines the intertemporal budget constraint that is relevant for the dynamic optimization problem they solve. And the markets individuals have access to in turn determine how much of shocks to individual resources or income that households can smooth and the intertemporal prices they pay for such smoothing.

In a world of complete markets, households can completely diversify idiosyncratic risk and achieve first best intertemporal allocations, given the aggregate shocks that affect their economy. The set of assets necessary to achieve such allocations can be very complex, depending on the nature of income processes, and might include a variety of state contingent arrangements. Such an equilibrium relies critically on the assumption of full information about idiosyncratic shocks and enforceability of contracts. Therefore, it is not surprising that many empirical tests reject the implications of full risk sharing.

At the other extreme, many studies have looked at Bewley models where individual households are endowed with relatively simple assets and can smooth only temporary shocks and, in the presence of borrowing restrictions, not even all of them. In reality, consumption smoothing households participate in a variety of markets and interact with other households both formally and informally. These interactions, even if they might fail to achieve full risk sharing, may afford more consumption smoothing and insurance possibilities relative to those that can be attained trading a set of

exogenously given assets. Therefore, when studying consumption smoothing behaviour, one should include in the intertemporal budget constraints any claim (contingent or not) that the household might be buying and/or holding. In particular, in addition to the standard assets considered in simple versions of the life cycle model, one should include contingent claims that might take many forms, especially where informal transfers and interpersonal ties might be playing an important role in smoothing consumption. Failure to do so, can lead to rejections of the model considered, as pointed out, for instance, by Attanasio and Pavoni (2011), which can take the form of ‘excess smoothness’ of consumption. Such a situation is certainly relevant for developing countries, but the same is true for developed countries where households might have access to a wide network of interpersonal ties that can be used to smooth out certain types of shocks as well as contingent assets, such as insurance or the availability of credit.

From an empirical point of view, the considerations above pose a big challenge, as it may be difficult to have complete information on the intertemporal budget constraints relevant for the individual households and the position on all assets and formal and informal insurance contracts in which they are active. An attractive approach to the study of risk sharing was used in a pathbreaking paper by Townsend (1994), who developed ideas in Wilson (1968) and Altug and Miller (1990), to focus on the properties of first best allocations, independently of the specific decentralization mechanism and assets used to achieve that allocation.

While the Townsend (1994) approach makes clear the implications of full risk sharing, things become more complicated when different imperfections prevent the attainment of first best allocations. However, characterising the deviations of actual allocations from those that would prevail under full risk sharing can be informative about the nature of the imperfections that characterise real economies.

Versions of the Townsend (1994) test have been used in many different contexts, both in developing countries, as in the original application, and in advanced economies (such as Cochrane (1991), Attanasio and Davis (1996)). These tests require the identification of a ‘risk sharing’ group, which could be a village or an entire economy, but can then be used to test the hypothesis that the intertemporal allocation of consumption *within that group* is intertemporally efficient, using only data on consumption and income. In particular, data on asset holdings or transfers are not necessary to perform the exercise.

More recently, Blundell, Pistaferri, and Preston (2008) have proposed an approach that estimates the fraction of temporary and permanent income shocks that are transmitted into consumption. By doing so, they can estimate which fraction of shocks is insured and which is not. Although the approach is different from Townsend's, the spirit is similar, in that it relies only on observations on income and consumption. The approach we take in what follows builds on Blundell, Pistaferri, and Preston (2008) by incorporating a smaller risk sharing group into the framework. In our empirical application, we focus on the extended family as a potential risk sharing group. Such an exercise is interesting because some of the imperfections that prevent full risk sharing, such as information and enforceability of contracts, might be less relevant within the extended family than in the economy at large.

As mentioned above, much of the available evidence thus far rejects the hypothesis of perfect risk sharing. Idiosyncratic shocks to income are reflected, to an extent, in consumption. However, the same evidence indicates that income shocks are not *fully* reflected in consumption. Blundell, Pistaferri, and Preston (2008), for instance, using US data, conclude that most of transitory shocks and about 60% of permanent shocks to income pass on to consumption, with the rest being insured. The fact that transitory shocks are not reflected in consumption is probably a consequence of *self-insurance*, that is, households use simple assets, such as savings, to absorb the temporary fluctuations in income. The fact that some of the permanent shocks seem to be insured is reminiscent of the *excess smoothness* finding discussed in Campbell and Deaton (1989). Using UK data, Attanasio and Pavoni (2011) report similar findings and interpret the fraction of shocks that are *not* insured as a consequence of information frictions and moral hazard.¹

Empirically, a major challenge involves defining the network within which households share risk. Much of the work that empirically tests these partial insurance models focuses on settings in developing countries which offer plausibly exogenous, well-defined networks (e.g. villages in India and Thailand: Townsend (1994), Kinnan (2014); sub-castes in India: Mobarak and Rosenzweig (2012)). In advanced economies, examples are harder to come by, but one network that is both well-defined and ubiquitous is the extended family.

¹An important literature has developed deriving formally the conditions under which partial insurance would occur. In these models the departure from complete markets occurs either because of limited commitment (e.g. Thomas and Worrall (1988), Ligon, Thomas, and Worrall (2002), Kocherlakota (1996)) or because of moral hazard with different amounts of information assumed to be observable or verifiable (e.g. Cole and Kocherlakota (2001), Golosov, Tsyvinski, and Werning (2007), Attanasio and Pavoni (2011)).

In this paper, we investigate, in a partial insurance framework, the extent to which extended families share risk in the United States. The extended family constitutes a natural network within which to look for risk sharing. Comparing the risk sharing that takes place within families to the risk sharing that occurs within the society at large can be informative about the mechanisms that are used to achieve certain allocations. In particular, even in the absence of detailed information on intra-personal transfers one could infer whether intra-family transfers play a big role in achieving the level of risk sharing observed in the data. As in Townsend (1994) and following an approach similar to Blundell, Pistaferri, and Preston (2008), we only use data on consumption and income.

An earlier literature focused specifically on risk sharing within the extended family. 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. However, their method does not allow them to quantify the extent of family-insurance. We follow up and extend this work in a number of ways. First, although perfect risk sharing is nested within our approach in that it implies specific values for the parameters we estimate, we do not consider it as the main hypothesis to test. Instead, we explicitly estimate the extent of partial risk sharing; indeed a motivation of this paper is to estimate the extent to which purely idiosyncratic shocks to household income are insured within the extended family. By considering the difference between “family-aggregate” and idiosyncratic shocks, and using explicitly the information of who is in the family network, we can detect and quantify the amount of extended-family insurance that takes place and distinguish this from self-insurance and intra-household insurance. Moreover, by estimating household and extended family income processes we are able to understand better the relative importance of family versus household level shocks as well as the extent to which intra-family insurance is feasible.² Second, by modeling consumption jointly with the income process we are able to estimate how much of that insurance is actually achieved. Finally, we exploit broader measures of consumption than did the earlier papers, which relied solely on food consumption.

The framework we propose is based on modeling jointly the stochastic process of income and

²Clearly this distinction may be endogenous, as households sort into occupations and sectors and may differentially invest in tasks or family members to take into account the insurance possibilities, minimizing the correlation (or even achieving negative ones), hence the interpretation of our model is post-sorting. In rural India, for example, Rosenzweig and Stark (1989) argue that parents marry their daughters to males in other villages to diversify the risk of weather-related income shocks.

consumption allowing for 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 now extend this work to distinguish between a group-aggregate and a purely idiosyncratic process. Consumption growth is modeled as a function of innovations to the income process as in Hall (1988), Hall and Mishkin (1982), and Blundell and Preston (1998) amongst others. The relationship of consumption growth to income shocks can be rationalized by an approximation of the Euler equation for consumption when preferences are CRRA, as in Blundell, Pistaferri, and Preston (2008).

Our framework also provides several new methodological insights. First, by distinguishing between idiosyncratic and group-aggregate shocks, we are able to identify completely the income and consumption processes, including measurement error in income, which is notoriously difficult to separate from transitory shocks. Second, since we can separately characterize and identify the time-series processes of purely idiosyncratic and group-aggregate shocks, we can additionally identify the amount of income fluctuations (both permanent and transitory) that are insurable by the group. This allows us to evaluate the potential opportunity of a network to share risk. Finally, besides group membership, our framework does not depend on knowledge of the risk sharing arrangements in place within the group.

We use the 1980-2010 waves of the Panel Study of Income Dynamics (PSID) coupled with the 1980-2008 Consumer Expenditure Survey (CEX) to test the model on extended family networks. Our decomposition of income shocks suggests that over 60% of shocks are potentially insurable by family risk sharing networks. However, even though extended families appear to be well-positioned to share risk between member households, we find no evidence of any insurance within the family network.

The paper proceeds in Section 2 with a discussion of alternative approaches using data on direct transfers. Section 3 presents our model and Section 4 discusses identification. Section 5 describes our data and estimation procedure and Section 6 reports our results. Section 7 concludes.

2 Evidence of family insurance using direct transfers

A direct approach to studying whether extended families share risk is to analyze transfers among family members. One strand of the 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 example, find that three-fourths of child care provided to working mothers by relatives is unpaid. In what follows, we ignore co-residence and other specific in-kind transfer decisions and focus instead on non-durable consumption and income co-movements. Our tests are valid under the assumption of separability between the definition of consumption we consider and the consumption of housing services or other in-kind transfers (such as care or baby-sitting). We leave the analysis of these other mechanisms to future work.

Some authors have also look directly at cash transfer data. Recently, McGarry (2012) 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.

Table 1: Parent-Child Transfers (2013 PSID)

	% Any Amt	Cond. Mean	Cond. p25	Cond. Median	Cond. p75
<i>Transfers Given</i>					
Hours	0.368	272 (569)	20	60	240
Money	0.321	3390 (8629)	500	1100	3500
<i>Transfers Received</i>					
Hours	0.204	168 (295)	15	50	208
Money	0.068	382 (361)	100	300	500

Data from parent reports. Standard deviations in parentheses

The test we propose below do not use direct information on transfers and focus, instead, on the

relationship between the distribution of consumption and income. Information on direct transfers, however, is useful to assess the importance that these informal mechanisms have in risk sharing. For this reason, in this section, we present some descriptive evidence on the prevalence of intra-family transfers.

We can perform an analysis similar to that in McGarry (2012) 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 one. 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). However, while there does appear to be a flow of transfers in both directions, the majority of it appears to be from parents to children. Table 2 presents a similar analysis using supplementary data from 1988 with an expanded universe of transfer recipients and finds that most transfers are between parents and children.

Table 2: Family and Friends Transfers (1988 PSID)

	Children		Parents	
	% Any Amt	Conditional Amt	% Any Amt	Conditional Amt
<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)

Total transfers include family and non-family transfers. Standard deviations in parentheses

Additionally, we are able to link the 1988 transfer data to (unexplained) income changes between 1987 and 1988. In Panel A of Table 3 we report the marginal effects of the quartile of a household's unexplained income change on the probability of receiving a monetary transfer, as estimated by a probit regression. From column 1, we see that households in the bottom quartile of income shocks are 7 percent more likely to receive a transfer than households in the top quartile of income shocks. This evidence suggests that these transfers may be playing an insurance role. The next columns presents estimates derived in 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.

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
<i>Panel A: Money Transfers</i>					
Income change quartile					
1 (negative)	0.073** (0.037)	0.133*** (0.049)	0.149*** (0.048)	-0.011 (0.023)	-0.021 (0.031)
2	0.012 (0.037)	0.005 (0.048)	0.021 (0.048)	-0.007 (0.023)	-0.029 (0.035)
3	-0.004 (0.036)	0.017 (0.049)	0.018 (0.048)	0.010 (0.022)	-0.044 (0.034)
4 (positive) - omitted					
N	1033	701	701	701	332
<i>Panel B: Time Transfers</i>					
Income change quartile					
1 (negative)	0.010 (0.044)	0.052 (0.058)	0.033 (0.054)	0.054 (0.038)	-0.049 (0.046)
2	-0.035 (0.043)	-0.066 (0.054)	-0.061 (0.053)	0.005 (0.036)	-0.052 (0.052)
3	-0.027 (0.041)	-0.026 (0.054)	-0.020 (0.054)	0.005 (0.037)	-0.030 (0.046)
4 (positive) - omitted					
N	1033	701	701	701	332

Marginal effects; Standard errors clustered by family in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The simple correlations presented in this section, therefore, suggest that: (1) parents and adult children might 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. In the next section, we discuss a model of partial insurance and extended family risk sharing that might be useful to explain these findings and give to them a structural interpretation. In particular, we want to quantify the role that extended family might have in insuring idiosyncratic income shocks of individual family members. The idea is that the type of imperfections that prevent full risk sharing in the economy may be less relevant for the extended

family.

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. Individual households are seen as a part of a group, such as the extended family, and the income processes will be written, without loss of generality, to reflect this. That is, we decompose the individual 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 that goes on within the group. Obviously, one could also decompose individual 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 groups.

We consider different market environments, ranging from a 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) with a bond paying a fixed interest rate. The consideration of these different cases and some approximations of the consumption function allow us to consider intermediate cases where households are able to insure part of certain idiosyncratic shocks. Throughout this section we assume that the only source of uncertainty is exogenous, post-tax and government transfer household income and preferences over consumption are separable from leisure. While this is a major simplification, we abstract from labor supply decisions and 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).³

³See Blundell, Pistaferri, and Saporta-Ekstein (2012) and Attanasio, Low, and Sanchez-Marcos (2005) for models of consumption insurance that incorporates household labor supply decisions and Lamadon (2014) for a model of firm-worker contracts and insurance.

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, even if they might belong to the same group.

The household is entitled to streams of uncertain income that are seen as exogenous stochastic processes $Y_{i,t}$. Following a well established tradition, we model household income as a permanent-transitory process (MaCurdy (1983), Abowd and Card (1989) and Meghir and Pistaferri (2004)) which is made of three components: (1) a deterministic component which we model as a function of demographics $z_{i,t}$,⁴ (2) a permanent component $P_{i,t}$, (3) a transitory component $\nu_{i,t}$. In addition, measured income is affected by a multiplicative measurement error $r_{i,t}^y$.

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

The permanent component follows a random walk in which the innovations, $u_{i,t}$, are 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 serially un-

⁴Specifically, in estimation we control for year, age, education, race, family size, number of kids, region, and urbanicity, and interactions of year with education, race, region, and urbanicity.

correlated as well:

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

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:

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

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. We will analyse two different risk sharing set ups: on the one hand, we consider the entire economy as a potential risk sharing group; on the other, we consider 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 type of frictions, such as informational asymmetries.

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. Such an aggregate budget constraint may or may not allow for aggregate savings. As we are not making any use of the condition relating to aggregate savings, for the sake of notational simplicity, we write the planner problem at time 0, without aggregate savings; the conditions we use would not be different in the presence of aggregate savings. Our approach focuses on intertemporal allocations *within group* G and it is completely agnostic about what happens *across groups*. We also specify the uncertainty in the economy in a slightly different fashion, following the literature. In particular, we assume that there is a state variable \mathbf{s}_t

which is a multidimensional vector which fully describes the state of the economy which can take discrete values and that evolves according to a Markov chain, so that $Pr\{\mathbf{s}_t = \mathbf{s}' | \mathbf{s}_{t-1} = \mathbf{s}\} = \pi_{s',s}$. The realization of a specific vector for the state variable determines completely the income received by all households. These realizations are fully observable and contractible upon. The social planner therefore maximizes:

$$Max_{\{\mathbf{C}_{i,0}\}} E_t[\sum_{i \in G} \lambda_i V(\mathbf{C}_{i,t})] \quad (2)$$

$$s.t. \quad \sum_{i \in G} Y_{i,\tau}(\mathbf{s}_\tau) = \sum_{i \in G} C_{i,\tau}(\mathbf{s}_\tau) \quad \forall \tau \geq t \quad (3)$$

In equation (2), λ_i is the weight given by the social planner to household i . Different weights correspond to different competitive equilibria and might reflect differences in ownership rights within the risk sharing group considered. Given that all income realization for all consumers are fully contractible, the social planner problem first order conditions⁵ for consumption of household i , state of the world \mathbf{s}_τ at time τ is:

$$\lambda_i \beta U'(C_{i,\tau}) \pi_t^{s(\tau)} = \mu(\mathbf{s}_\tau) \quad (4)$$

where $\mu(\mathbf{s}_\tau)$ is the multiplier associated to constraint (3), $\pi_t^{s(\tau)}$ is the probability of state \mathbf{s}_τ given the current state at time t and $U'(C_{i,\tau})$ the marginal utility of consumption for household i .

Equation (4) is key to characterize the properties of first best allocations of resources where idiosyncratic risk is fully diversified. Notice that, as all states of the world are fully contractible, the equation holds state by state and not in expectation. It is useful to re-write equation (4) as:

$$\lambda_i \beta U'(C_{i,\tau}) = \mu(\mathbf{s}_\tau) / \pi_t^{s(\tau)} \quad (5)$$

so that all the household-specific variables are on the left-hand-side of the equation. If we consider it at two different points in time, τ and τ' , and take the ratios of the two equations, we obtain:

⁵There is a first order condition for each state of the world at τ .

$$\frac{U'(C_{i,\tau})}{U'(C_{i,\tau'})} = \frac{\mu(\mathbf{s}_\tau)/\pi_t^{s(\tau)}}{\mu(\mathbf{s}_{\tau'})/\pi_t^{s(\tau')}} = \nu(\tau, \tau') \quad (6)$$

Notice that in equation (6), 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 (6) considered at two adjacent time periods and obtain:

$$\Delta \log(c_{i,t}) = \psi_t + \varepsilon_{i,t} \quad (7)$$

Townsend (1994) tests 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, individual 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 individual income should not enter significantly in such an equation.

Another interesting way to consider the implications of equation (4) is to take logs of both sides (again under the assumption of CRRA utility), rearrange it and take its cross-sectional variance within the risk sharing group. In this case, we obtain:

$$Var_G(\log(c_{it})) = Var(\lambda_i)/\gamma^2 \quad (8)$$

where γ is the coefficient of relative risk aversion (which is assumed to be constant across households). Notice that under perfect risk sharing, the Pareto weights λ_i are constant, implying that the right-hand side of equation (8) is also a constant. Another implication of perfect risk sharing, therefore, is that the cross sectional variance of log-marginal utility (here approximated by log consumption) is constant over time.

Notice that this characterization of perfect risk sharing only requires data on consumption allocations and idiosyncratic shocks. The test is silent and agnostic about the specific decentralization through which first best allocations can be achieved or about the specific assets and contracts

(formal and informal) that households might be using. Notice also that under perfect risk sharing there is no distinction between (idiosyncratic) permanent and transitory shocks.

At the other extreme of the assumption of perfect risk sharing, one can consider an economy where individual 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, which implies a very simple individual budget constraint, has been referred to as the Bewley model. In such a situation, one is within the realm of a standard life cycle model: transitory shocks are almost fully smoothed out and permanent ones are, instead, almost completely reflected in consumption. The 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.

As is well know, 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, such as 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:

$$\Delta \log c_{i,t} = \delta u_{i,t} + \gamma(1 + \theta)e_{i,t} + \Delta r_{i,t}^c + \xi_{i,t} \quad (9)$$

where δ measures the degree to which permanent shocks pass through to consumption and γ measures the degree to which transitory shocks pas through to consumption. In addition, we allow for classical measurement error in log consumption, $r_{i,t}^c$ and permanent innovations to consumption that are independent of income, $\xi_{i,t}$, possibly reflecting innovations to preferences.

Under this model, the values of δ and γ should solely be dictated by the ability to smooth shocks through self-insurance. Under CRRA preferences, self-insurance is attained through both the potential to borrow from future income streams as well as precautionary savings. In such a set up, 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 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. Intuitively, younger households with low current savings relative to lifetime savings ($\pi_{i,t}$ closer to one) are less able to effectively self-insure through savings, while older households who have realized more of their savings potential ($\pi_{i,t}$ closer to zero) can smooth shocks through savings. All else equal, it follows that older households should have more insurance against permanent shocks (lower δ) than younger households. Meanwhile, in the absence of liquidity constraints, households can borrow against future income to cushion transitory shocks. This helps younger households, who have a longer time horizon over which to borrow (lower $\alpha_{i,t}$), smooth transitory shocks beyond precautionary savings.

Between the two extremes of perfect risk sharing and the Bewley model, there are a variety of intermediate cases where households might be able to smooth parts of permanent shocks. Attanasio and Pavoni (2011), for instance, consider a model with endogenously incomplete markets caused by information frictions in both effort (moral hazard) and assets and show that, in equilibrium, households can insure a part of (but not all) permanent shocks and achieve consumption allocations that exhibit, relative to the Bewley model, ‘excess smoothness’ in the sense of Campbell and Deaton (1989). One can also consider the possibility that, unlike in the standard life cycle model, a fraction of transitory shocks are reflected into consumption, perhaps due to binding liquidity constraints that prevent households from consumption smoothing.

An equation like (9) is particularly useful in this context, as it can be used to identify the fraction of permanent and transitory idiosyncratic income shocks that are transmitted to consumption. The size of the coefficients identified 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 $\pi_{i,t} = 0.8$ and $\delta = 0.64$. 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. This is consistent with the results in Attanasio and Pavoni (2011) for the UK. Attanasio and Pavoni (2011) provide a structural interpretation of the parameter δ as reflecting the extent of informational frictions. Next we turn to a model of family risk sharing that may help provide an explanation for this additional insurance.

3.3 Incorporating Family Risk Sharing

In the previous subsection, we stressed that the risk sharing group G considered there was somewhat arbitrary. Here we consider explicitly risk sharing within the extended family. As we discuss below, with the appropriate data, one can identify the parameters of equation (9), which define the extent to which idiosyncratic shocks are insured within a given group, for the special group defined by the extended family. The family might be particularly interesting as a risk sharing institution because it may be able to deal more effectively with the reasons 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 avoiding moral hazard, 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 (9), where one considers implicitly risk sharing across the whole economy, with those that one obtains considering the extended family as a risk sharing group. To start, we can express the income process in equation (1) in terms of deviation of the individual household idiosyncratic component from the extended family aggregate. In particular, we define $u_{j,t}^F$ as the aggregate permanent shock to family resources for family j , and $u_{i,j,t}^I$ as the idiosyncratic permanent shock to member i in family j , such that $u_{j,t}^F + u_{i,j,t}^I = u_{i,t}$. Analogously, let $e_{j,t}^F + e_{i,j,t}^I = e_{i,t}$ for transitory shocks. Then rewriting equation (1), the growth in log income is:

$$\Delta \log y_{i,j,t} = u_{j,t}^F + u_{i,j,t}^I + \Delta(e_{j,t}^F + \theta e_{j,t-1}^F) + \Delta(e_{i,j,t}^I + \theta e_{i,j,t-1}^I) + \Delta r_{i,j,t}^y. \quad (10)$$

By definition, it must be the case that the sum of the idiosyncratic shocks across family members is zero for both permanent and transitory shocks: $\sum_{i=1}^{n_j} u_{i,j,t}^I = 0$ and $\sum_{i=1}^{n_j} e_{i,j,t}^I = 0$. There is no loss of generality and no particular restriction implied by the way we have written equation (10). We allow the variance of the individual and family component of both the transitory and permanent shocks of the income process to be different. Notice, however, that we assume that the persistence parameter of the temporary shocks θ is assumed to be the same for the family and individual components.

The decomposition of income shocks into idiosyncratic and family-aggregate components allows us to quantify what percentage of shocks *could* be insured by the family, which effectively defines

the risk sharing opportunity that the family has. Idiosyncratic shocks are by definition household-level deviations from the family-average shock, and hence the family network can redistribute funds between households to smooth these shocks. Family-aggregate shocks, on the other hand, cannot be smoothed by family networks. Therefore, the pass-through of idiosyncratic income shocks to consumption may differ from the pass-through of family-aggregate shocks.

To study these differences in pass-through rates, we rewrite equation (9), the growth in log consumption, as:

$$\Delta \log c_{i,j,t} = \delta_I u_{i,j,t}^I + \delta_F u_{j,t}^F + \gamma_I(1 + \theta)e_{i,j,t}^I + \gamma_F(1 + \theta)e_{j,t}^F + \Delta r_{i,j,t}^c + \xi_{i,j,t} \quad (11)$$

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.

As we discussed before, equation (11) nests a wide variety of models, ranging from the Bewley model to perfect risk sharing. Moreover, if we can identify all the parameters of this equation, we can consider simultaneously risk sharing within and across 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 (11). We focus the discussion on the insurance parameters for permanent 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 Family Risk Sharing. Under perfect family risk sharing, the distribution of income between family members has no effect on the distribution of consumption between family members (Hayashi, Altonji, and Kotlikoff (1996)). Thus, controlling for shocks to the family 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 member 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 family j . Overall, perfect family risk sharing predicts that $0 = \delta_I \leq \delta_F$.

In sum, our framework allows us to distinguish between two extreme cases of family risk sharing behavior: zero risk sharing (self-insurance) and perfect family risk sharing. In addition, it allows us to quantify the amount of family insurance by using the null of self-insurance, $\delta_I = \delta_F \leq 1$ and estimating the degree to which $\delta_I < \delta_F$.

4 Identification

The model we have presented can be seen as a stochastic factor model. The econometric structure consists of two equations, one for income growth and one for consumption growth. Each of these equations depends on some unobserved common factors, namely the permanent and transitory shocks for the family and the individual as well as mutually independent shocks affecting each of the two processes, including measurement error or taste shocks. The covariance structure of the factors is informative because it defines the extent of uncertainty facing the households as well as the extent to which this is insurable within the family. In addition, we are interested in the coefficients associated to the factors in the consumption growth equation because they reflect the amount of insurance that occurs within and between extended families. The parameters of the factor model need to be estimated from panel data on consumption and income. We now discuss identification of such a model.

The 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 $\text{var}(u^F)$ and $\text{var}(u^I)$ and transitory income variances $\text{var}(e^F)$ and $\text{var}(e^I)$, (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.⁶ To identify parameters, we use covariances that exploit both time and within-family dimensions.

Define the vector $Y_{ijt} = [\Delta y_{i,j,t}, \overline{\Delta y_{j,t}}, \Delta c_{i,j,t}, \overline{\Delta c_{j,t}}]$, where y_{it} represents residual log income and c_{ijt} residual log consumption for household i in extended family j at time period t respectively, after removing the effects of education, age, demographic composition, and aggregate trends. The expressions with the overbar relate to family wide aggregates, as defined above. Finally Δ represents the first difference operator ($\Delta x_t = x_t - x_{t-1}$).

We have already defined the factor structure of individual income and consumption growth. We also need to define the model for family level income and consumption growth. This is obtained by aggregating equation (10) and equation (11) within a family:

$$\begin{aligned}\overline{\Delta \log y_{j,t}} &\equiv \frac{1}{n_j} \sum_{i=1}^{n_j} \Delta \log y_{i,j,t} = u_{j,t}^F + \Delta(e_{j,t}^F + \theta e_{j,t-1}^F) + \frac{1}{n_j} \sum_{i=1}^{n_j} \Delta r_{i,j,t}^y \\ \overline{\Delta \log c_{j,t}} &\equiv \frac{1}{n_j} \sum_{i=1}^{n_j} \Delta \log c_{i,j,t} = \delta_F u_{j,t}^F + \gamma_F (1 + \theta) e_{j,t}^F + \frac{1}{n_j} \sum_{i=1}^{n_j} \Delta r_{i,j,t}^c + \frac{1}{n_j} \sum_{i=1}^{n_j} \xi_{i,j,t}\end{aligned}$$

Deviations from these measures (e.g. $\Delta \log y_{i,j,t} - \overline{\Delta \log y_{j,t}}$) are by construction idiosyncratic to household i in family j and perfectly insurable within the extended family network.

Our observable moments relate to the cross sectional and time series covariance matrix of Y_{ijt} . For example we use the covariance of consumption growth with income growth. We also use the autocovariances of consumption growth and of income growth both within one time period and across time periods (e.g. $\text{cov}(\Delta c_{i,j,t}, \Delta y_{i,j,t+1})$ and $\text{cov}(\Delta c_{i,j,t}, \Delta y_{i,j,t+2})$).

The structure we use implies a number of restrictions on these covariances. In addition, we make some additional substantive assumptions that allow the identification of our parameters of interest. A first set of restrictions follow from the definition of family level shocks: by construction the overall family shocks and the idiosyncratic shocks for each household within the family are uncorrelated. Moreover, the idiosyncratic shocks of individual households within a family are correlated in a specific way implied by the fact that they have to add up to zero. These restrictions are completely

⁶We 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.

innocuous as they follow from the definition.

We assume that the shocks to income, whether transitory or permanent, are independent of any taste shocks. Moreover, we assume that measurement error is independent and identically distributed as well as independent between income and consumption. Our final restriction follows from earlier work by Meghir and Pistaferri (2004) that show that we can represent log income (conditional on aggregate shocks, age, education and demographic characteristics) as following a random walk and an MA(1) error. Here we assume this structure, which provides a number of restrictions. It should be stressed that this is not an identifying assumption and that more general structures can be allowed, but they would have no empirical content here, as we show in Section 6. The variance of measurement error in income is identified separately from the variance of the transitory shock if the latter follows a moving average process of order one or higher, by using cross family information. This is in contrast to the result in Meghir and Pistaferri (2004) which did not use information from related extended family units.

Finally, note that our 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.

We provide a detailed explanation of the moments used and how these lead to the estimated parameters in Appendix A.

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).

5.1 PSID Sample

We use the 1979-2011 PSID, a longitudinal study of US households that started in 1968 with a sample of about 5,000 households. Of these, around 3,000 were representative of the US population

(the core sample), and around 2,000 were low-income families (the Census Bureau’s SEO sample).⁷ After 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 family risk sharing network definition.

For each survey wave, the PSID collected information on annual household income and weekly food expenditures.⁸ Starting in 1999, the PSID also began collecting information on specific consumption expenditures that cover around 70% of total consumer expenditures. We supplement our analysis with this additional measure of consumption for these more recent years. In addition, the PSID collects a variety of demographic information that we use in our analysis, including age, race, education, employment status, household size, number of children, and state and urban status of residence, among others. Finally, the PSID collected more detailed information on time and monetary transfers between family and friends in the 1988 and 2013 waves which we also exploit.

We define a “family” in our sample as a cohabiting couple and their⁹ adult children. We only include adult children who are at least 25 years old (to avoid large education changes) and once they have split from the original household unit.¹⁰ 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 yearly income and consumption changes). This results in a final sample of 1,410 unique 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 consist of three individuals on average, and child

⁷In addition, in 1997 around 500 immigrant families were added.

⁸Households are asked about income from the previous calendar year, while the reference period for food expenditure is not clear. We assume that the reference period is the same as that for income, as in previous work.

⁹From the main data, it is not obvious whether the spouse of the head is also a parent, thus the PSID provides a supplemental dataset that links children in the sample to their parents. However, while this parent file contains most child-parent links, it does not cover the entire universe of children. Thus we only keep family units in which we are sure that the spouse of the head is also a parent (meaning that they are in the parent file). This reduces the sample by about 4%.

¹⁰We 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.

households have at least one young child on average. Income and consumption are slightly lower for child households, but not drastically so. Average family size is three households.

Table 4: Summary Statistics

	Parents	Children
Age	56.03 (4.94)	30.17 (3.95)
Household size	2.79 (1.33)	3.09 (1.45)
Number of kids 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	0.35 (0.48)	0.38 (0.48)
Lives in big city	0.33 (0.47)	0.35 (0.48)
Annual income	34195 (19009)	28215 (15112)
Annual imputed consumption	17719 (11259)	14874 (9754)
Annual food consumption	4125 (1817)	3749 (1740)
Annual reported consumption	8862 (3650)	8418 (3909)
Family size		3.28 (1.25)
Number of unique families		1410
Number of family-year observations		8783
Years of data		25 (1980-2010)

Standard deviations in parentheses. See text for income and consumption definitions.

The PSID collects a measure of household total income, which is the sum of taxable (wage and salary income, asset income, and business profits) and transfer (alimony, annuity, child support, help from family and friends, retirement, SSI, TANF, UI, VA pension, welfare, workers compensation, and food stamps) income of members of the household. In our analysis we use a modified income definition which nets out taxes and does not include family transfers. In this way we control for

the insurance that is offered (implicitly or explicitly) by the tax and benefit system. 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. Finally, we deflate income by the CPI-Urban deflator (with a base year of 1983-1984).

Our main consumption measure is imputed from food consumption using estimates from a food demand system estimated on the CEX (see Section 5.2). Food consumption is defined as expenditures on food eaten inside and outside the home as well as food stamps. In addition to our food and imputed consumption measures, we also provide results using “reported consumption” for 1999-2009, defined as the sum of expenditures on food, gasoline, other transportation expenses, utilities, and home insurance.¹¹ Reported and imputed consumption are deflated by the CPI-Urban deflator, and food consumption is deflated by the food CPI deflator.

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 (2004) 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.

¹¹The PSID also collected data on health, childcare, education, and rent, but we exclude these categories in our “reported consumption” definition to be more consistent with our imputed consumption definition.

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 consumption for households and calculate data covariances. Second, we estimate the parameters of the income process, and 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 $y_{i,j,t}$ and $c_{i,j,t}$ of a regression of each variable on dummies of family 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 family, we take the average log unexplained growth of all households in the family: $\overline{\Delta \log y_{i,t}} = \frac{1}{N_j} \sum_j \Delta \log y_{i,j,t}$ and $\overline{\Delta \log c_{i,j,t}} = \frac{1}{N_j} \sum_j \Delta \log c_{i,j,t}$, 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.¹² 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 family level (Hall and Horowitz (1996), Horowitz (2001)). This method allows us to account for arbitrary serial correlation between family members and for the fact that the second and third steps took the estimates from the previous steps as given.

¹²Estimating 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.

Income and consumption data is only available every other year starting in 1999. To account for this in the estimation procedure, we work with changes in income and consumption over a two year period (“long-difference”: $y_{i,j,t} - y_{i,j,t-2}$ instead of $y_{i,j,t} - y_{i,j,t-1}$). This slightly modifies the income and consumption growth equations, and hence the covariances and identification. See Appendix C for more details.

6 Results

Our goal is to quantify the extent of insurance provided by a household’s family risk sharing network. To arrive at this, we must first characterize the structure of the income and consumption processes, and then estimate the relevant parameters.

6.1 Characterization of the Income and Consumption Processes

In Figure 1, we plot autocovariances of (residual) log income growth by year at both the individual household level and the family-average level (see Table 8 in the appendix for point estimates as well as second order autocovariances). Variances after 1996 are estimated to be considerably larger after 1996 than before. This difference, however, is explained by changes in the survey structure. After 1996, the PSID started collecting data bi-annually: we expect the “long-difference” variances (those in more recent years) to be larger, since, for example, they include the effect of two permanent shocks as opposed to one for variances of single-year differences.

The family-average variances are much smaller than the household variances. This result is also to be 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.

In Figure 1, we also plot the first order autocovariance at the household and extended family level. We observe that, whilst they are both negative (consistently with previous evidence), the absolute value of the household level autocovariance is considerably larger than the one for the extended family shocks. However, the ratio of the first order autocovariates to the variances for the individual and family level shocks is roughly of the same order of magnitude, justifying our

assumption that the parameter θ is the same for the family and individual components of the temporary shocks.

The other set of moments that is used to identify the parameters of the income process is the second order autocorrelation of both family level and individual income. We do not plot those, but the point estimates of these estimates (and their standard errors) are reported in Table 8 in the Appendix. These autocorrelation 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.

In Figure 2, we plot the variance and first order auto-covariance of consumption (see Table 9 in the appendix for point estimates). Again, we note the large difference between the figures prior to and after 1996. We also observe, as with income, large differences between household and family-average variances. Measurement error in consumption can be inferred from the first-order autocovariance (for single-year differences, at least; see equations (19) and (24) in the appendix). This includes error from the imputation procedure. Higher-order autocovariances are insignificant (not shown), which is consistent with the implications of the model.

Figure 1: Income variances and autocovariances





Figure 2: Consumption variances and autocovariances, imputed from CEX



The vertical line at 1996 indicates a change in PSID survey frequency. Gaps in late 1980's are due to missing food consumption questions in 1988 and 1989.

6.2 Parameter Estimation

To the best of our knowledge, ours is 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.¹³ As mentioned above, we assume that the persistence parameter of transitory shocks θ is the same for the individual and family components of income, while the variances are allowed to differ.

The first thing to note is that both permanent (columns 1-2) and transitory (columns 4-5) 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 shocks are perfectly positively correlated between family members, the family is an ineffective group within which to share risk, because each member is faced with the same shock. At the other extreme, if shocks are perfectly negatively correlated between family members, the

¹³Table 10 in the Appendix reports estimates and bootstrapped standard errors of the parameters of the income process year by year.

family can completely smooth fluctuations between family members.

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 over 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 family risk sharing networks. The question, of course, is the extent to which it does.

Table 5: Average Income Parameter Estimates

Variance of Permanent Shocks			Variance of Transitory Shocks		
Idiosyncratic	Fam-Agg	% Fam-Insurable	Idiosyncratic	Fam-Agg	% Fam-Insurable
0.018 (0.002)	0.011 (0.001)	0.608 (0.038)	0.026 (0.006)	0.015 (0.003)	0.635 (0.024)
θ (Serial correlation of transitory shock)			0.152 (0.085)		
Income measurement error variance			0.014 (0.008)		

Standard errors based on 100 block bootstrap replications in parentheses.

Another important aspect to note is that the estimates of the transitory shock variances do not include the effects of measurement error. Instead, as we show in Appendix A, we are able to separately identify the variance of measurement error from transitory shocks. Our estimate (0.014), 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 ($\theta = 0.152$ with a standard error of 0.085). If the true θ is in fact zero, then the variance of the measurement error of income is unidentified 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.

Table 6: Partial Insurance Estimates

	δ (Permanent)	γ (Transitory)
Imputed Consumption, 1980-2008	0.519 (0.085)	-0.026 (0.045)
Food Consumption, 1980-2008	0.314 (0.048)	-0.029 (0.030)
Reported Consumption, 1998-2008	0.273 (0.074)	0.064 (0.058)

Standard errors based on 100 block bootstrap replications in parentheses.
Slope heterogeneity and time-varying measurement error variances also
estimated but not reported.

Before discussing the family insurance estimates, in Table 6, we present partial insurance estimates based on a model where the household may be a member of some unspecified insurance network, or indeed just self insured based on own assets. In other words, we ignore the division of income into a family and household component and ignore the possibility that income shocks might be smoothed out through family networks. 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, which Blundell, Pistaferri, and Preston (2008) do not consider, allowing in effect only for shocks to wages.¹⁴

To perform this exercise, we sum the idiosyncratic and family-aggregate components of income estimated above and estimate the consumption parameters as in equation (9). Consistently with prior work, we cannot reject the hypothesis that transitory shocks are completely smoothed and not reflected in consumption. Of course this could happen either through intrahousehold and intrafamily transfers or through self-insurance (savings). In contrast, a 10% decrease in permanent income is associated with a 5.19% decrease in non-durable consumption, a 3.14% decrease in food consumption, and a 2.73% decrease in reported non-durable consumption.¹⁵ This evidence 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 48% of the decrease in income does not translate into non-durable consumption indicates that there is a substantial amount of insurance

¹⁴There 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).

¹⁵Note that reported consumption is only available from 1996-2008. When we estimate δ using our measure of imputed consumption over the same time period, we find a value of δ around two-thirds the size of the estimate over the full period. This may explain part of the difference in the estimates using the two measures of consumption, but because the estimate using reported consumption is so low, we do not focus our analysis on it.

against shocks to disposable income above and beyond the government tax and transfer system, which we explicitly account for. 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 on the partial within family insurance coefficients. The second column, δ_F , reports the transmission of permanent shocks that are aggregate to the family and hence not insurable through family networks, while the first column, δ_I , reports the transmission of the idiosyncratic portions of permanent shocks which are therefore insurable through family networks. The difference between these columns, $\delta_F - \delta_I$, reported in column 3, quantifies the level of family insurance. We find no evidence of any extra insurance within the family against permanent shocks or transitory shocks for imputed or food consumption. Specifically the idiosyncratic and family level shocks transmit to the same extent. The largest difference we find with some within family insurance suggested is for reported consumption, but the effect is insignificant. This suggests that at least overall the extended family in the US is not a source of even partial insurance.

One reason risk sharing might break down is moral hazard or incomplete and asymmetric information: in an environment where households have an incentive to shirk or hide assets insurance may break down. One interpretation of these results 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 or other reasons that inhibit full insurance within the family, we can provide suggestive evidence by performing our within-family insurance test on specific subsamples. Rows 4 and 5 of Table 7 show estimates for families that all live in the same state and for families that do not.

Table 7: Family Insurance Estimates

	Permanent			Transitory		
	δ_I (Idio)	δ_F (Fam-Agg)	$\delta_F - \delta_I$	γ_I (Idio)	γ_F (Fam-Agg)	$\gamma_F - \gamma_I$
Imputed	0.509 (0.132)	0.537 (0.092)	0.028 (0.153)	-0.050 (0.065)	0.023 (0.061)	0.074 (0.085)
Food	0.326 (0.080)	0.298 (0.061)	-0.028 (0.097)	-0.045 (0.044)	-0.001 (0.038)	0.044 (0.053)
Reported	0.253 (0.198)	0.353 (0.129)	0.100 (0.255)	0.068 (0.122)	0.052 (0.078)	-0.016 (0.151)
<i>Subsamples:</i>						
Same State	0.315 (0.167)	0.689 (0.150)	0.374 (0.219)	0.071 (0.098)	-0.128 (0.110)	-0.199 (0.147)
Different State	0.470 (0.181)	0.289 (0.109)	-0.181 (0.215)	-0.065 (5.796)	0.291 (3.091)	0.355 (3.480)
High Wealth	0.264 (0.171)	0.278 (0.095)	0.014 (0.207)	-0.039 (1.379)	0.135 (1.741)	0.174 (1.285)
Low Wealth	0.669 (0.197)	0.475 (0.132)	-0.194 (0.235)	-0.127 (0.148)	0.048 (0.128)	0.175 (0.210)
High Education	0.433 (0.166)	0.312 (0.114)	-0.121 (0.207)	0.094 (0.407)	0.082 (0.567)	-0.012 (0.388)
Low Education	0.675 (0.302)	0.582 (0.185)	-0.093 (0.358)	-0.084 (0.176)	0.025 (0.170)	0.109 (0.251)

Standard errors based on 100 block bootstrap replications in parentheses. Slope heterogeneity and time-varying measurement error variances also estimated but not reported. Subsample analyses (run separately) use imputed consumption and set $\text{var}(m_y) = 0.014$.

Family members who all live near each other are probably better able to monitor the behavior and events of other family members; we use living in the same state as a proxy for this: for families that live close by we cannot reject they hypothesis that there is complete within family insurance for permanent (as well as for transitory) shocks. For those who live far apart we reject complete family insurance. In both cases the hypothesis that family level shocks are fully insured is again rejected. While highly suggestive, this evidence cannot be taken as conclusive because we cannot reject the hypothesis that the coefficients differ from each other: a larger data set would be very important here as well as more detailed information on geographic distance. Of course there is 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.

In the remaining rows of Table 7 we explore how insurance estimates vary by assets and education. In rows 6 and 7 we split the sample into families whose parents have assets, respectively, in the top and bottom 40% of the asset distribution of parents. Families with larger asset holdings are better able to self-insure, which may translate into a higher willingness to insure other family mem-

bers. We find that generally the transmission of permanent shocks is lower for high-asset families than low-asset families, which suggests that families with higher assets are better able to self-insure. Indeed the transmission coefficients are not jointly significant in families with high wealth parents, while they are highly significant for low wealth households (however, we are not able to reject that the difference between the two sets of estimates is zero, as shown in column 3). This points in the direction of self-insurance as the main mechanism, rather than intra-family transfers. Finally, rows 8 and 9 splits the sample into families in which the father went to college or the father did not graduate high school. Although the point estimates are suggestive that the networks with the higher education father are better at insuring shocks the estimates are too imprecisely estimated to draw conclusions.

7 Discussion and Conclusions

Income shocks to households are substantial, even accounting for the mitigating effects of taxes and welfare. It is well understood that such fluctuations have very large welfare effects (see for example Low, Meghir, and Pistaferri (2010)) and yet on average only about half of the fluctuations are insured by some channel - probably the holding of assets.

There is a growing theoretical and empirical literature on partial insurance in which commitment issues or moral hazard prevent full insurance. However, in many of these models some degree of partial insurance can arise at least under some conditions that allow members of a network to commit to some level of transfers. An example is the work of Ligon, Thomas, and Worrall (2002) where partial insurance can arise with sufficiently patient and infinitely lived agents. Attanasio and Pavoni (2011) derive conditions for partial insurance when the imperfection is moral hazard. Since the immediate family should be able to resolve in part the commitment and moral hazard issues one may expect the family network to provide some degree of insurance. Indeed we show that such a system of transfers would go a long way towards mitigating the effects of risk since over half of the fluctuations are insurable within the network.

However, Hayashi, Altonji, and Kotlikoff (1996) have already shown that complete within family insurance does not occur. Kaplan (2012) has already demonstrated that the ability to move back in with parents provides some insurance for young adults against adverse shocks; hence there is some

evidence that family networks may provide some insurance. Here we now show that the extent of insurance for insurable shocks within the extended family network consisting of all siblings and the parents is no better than the mitigation we see for overall or family level shocks. Of course, this is an average effect and our results are suggestive of better insurance for families that live close to each other and where presumably the ties and the monitoring are tighter.

Beyond the implication that these estimates have for the welfare costs of risk they also have policy implications. Despite the fact that public welfare programs are unable to offer sufficient insurance, private networks do not replace them. Thus it is not the case that extending public insurance programs necessarily would crowd out private networks. While public programs may have other moral hazard consequences, crowding out does not appear one of them, at least from the baseline of where the US programs are now.

In the first part of the paper, we present a framework of partial insurance and group risk sharing that decomposes income shocks into group-aggregate and idiosyncratic components. 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 completely identify - including the identification of 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 amount (over 60%) of income shocks are idiosyncratic within extended family networks. We argue that this implies a large potential for the family risk sharing network to have a non-trivial impact on the transmission of income shocks into consumption. However, no such insurance occurs on average. Suggestive evidence shows that some insurance is provided by family networks in narrower geographic areas.

In conclusion, while we apply this framework to study extended families in the US, it is more generally applicable to any group of households whom we suspect may collectively share risk, such as rural villagers in developing countries (Meghir, Mobarak, Mommaerts *et al.* (2014) apply this framework to villages in rural Bangladesh). By considering the difference between 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 this from 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 pp. 411–445. ISSN 00129682.
- Altonji, J. G., Hayashi, F., and Kotlikoff, L. J. (1992). ‘Is the Extended Family Altruistically Linked? Direct Tests Using Micro Data.’ *The American Economic Review*, pages 1177–1198.
- 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.
- Altug, S. and Miller, R. A. (1990). ‘Household Choices in Equilibrium.’ *Econometrica*, vol. 58(2).
- 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. (2004). ‘Imputing consumption in the PSID using food demand estimates from the CEX.’
- Blundell, R., Pistaferri, L., and Preston, I. (2008). ‘Consumption Inequality and Partial Insurance.’ *The American Economic Review*, vol. 98(5):pages pp. 1887–1921. ISSN 00028282.
- Blundell, R., Pistaferri, L., and Saporta-Ekstein, I. (2012). ‘Consumption Inequality and Family Labor Supply.’ *Working Paper*.
- Blundell, R. and Preston, I. (1998). ‘Consumption inequality and income uncertainty.’ *The Quarterly Journal of Economics*, vol. 113(2):pages 603–640.
- 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*, pages 1–24.
- Campbell, J. J. and Deaton, A. (1989). ‘Why Is Consumption So Smooth?’ *The Review of Economic Studies*, pages 357–374.
- 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.
- 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: Journal of the Econometric Society*, pages 891–916.
- Hall, R. E. (1988). ‘Intertemporal Substitution in Consumption.’ *Journal of Political Economy*, vol. 96(2).
- 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. ISSN 00129682.
- Hayashi, F., Altonji, J., and Kotlikoff, L. (1996). ‘Risk-Sharing between and within Families.’ *Econometrica*, vol. 64(2):pages pp. 261–294. ISSN 00129682.
- 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(3):pages 446–512.
- Kinnan, C. (2014). ‘Distinguishing barriers to insurance in Thai villages.’
- Kocherlakota, N. R. (1996). ‘Implications of efficient risk sharing without commitment.’ *The Review of Economic Studies*, vol. 63(4):pages 595–609.
- Lamadon, T. (2014). ‘Productivity shocks, dynamic contracts and income uncertainty.’ Tech. rep., Working paper.
- Ligon, E., Thomas, J., and Worrall, T. S. (2002). ‘Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies.’ *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. (2012). ‘Dynamic aspects of family transfers.’ Tech. rep., National Bureau of Economic Research.
- Meghir, C., Mobarak, A., Mommaerts, C., and Morten, M. (2014). ‘Migration, Seasonality and Consumption Insurance.’ *Working Paper*.
- Meghir, C. and Pistaferri, L. (2004). ‘Income Variance Dynamics and Heterogeneity.’ *Econometrica*, vol. 72(1):pages pp. 1–32. ISSN 00129682.
- Meyer, B. D. and Mok, W. K. C. (2006). ‘meyer.’ Working Papers 0610, Harris School of Public Policy Studies, University of Chicago.

- Mobarak, A. M. and Rosenzweig, M. R. (2012). ‘Selling formal insurance to the informally insured.’
- Rosenzweig, M. R. and Stark, O. (1989). ‘Consumption smoothing, migration, and marriage: Evidence from rural India.’ *The Journal of Political Economy*, 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*, pages 84–112.
- Thomas, J. and Worrall, T. S. (1988). ‘Self-enforcing Wage Contracts.’ *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.
- Wilson, R. (1968). ‘The Theory of Syndicates.’ *Econometrica*, vol. 36(1):pages 119–132.

Appendix A: Moments used in identification and estimation

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

$$\text{cov}(\Delta y_{i,j,t}, \Delta y_{i,j,t}) = \text{var}(u_t^I) + \text{var}(e_t^I) + (\theta - 1)^2 \text{var}(e_{t-1}^I) + \theta^2 \text{var}(e_{t-2}^I) \quad (12)$$

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

$$\text{cov}(\Delta y_{i,j,t}, \Delta y_{i,j,t+1}) = (\theta - 1)\text{var}(e_t^I) - \theta(\theta - 1)\text{var}(e_{t-1}^I) \quad (13)$$

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

$$\text{cov}(\Delta y_{i,j,t}, \Delta y_{i,j,t+2}) = -\theta\text{var}(e_t^I) - \theta\text{var}(e_t^F) \quad (14)$$

Family-average autocovariances

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t}}) = \text{var}(u_t^F) + \text{var}(e_t^F) + (\theta - 1)^2 \text{var}(e_{t-1}^F) + \theta^2 \text{var}(e_{t-2}^F) + \frac{2}{n_j} \text{var}(m_y) \quad (15)$$

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t+1}}) = (\theta - 1)\text{var}(e_t^F) - \theta(\theta - 1)\text{var}(e_{t-1}^F) - \frac{1}{n_j} \text{var}(m_y) \quad (16)$$

$$\text{cov}(\overline{\Delta y_{j,t}}, \overline{\Delta y_{j,t+2}}) = -\theta\text{var}(e_t^F) \quad (17)$$

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

1. From (17) 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 (14) 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 (16) from (13) 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, (17) identifies $\text{var}(e_t^F)$ and consequently (14) identifies $\text{var}(e_t^I)$

2. (16) or (13) then identifies $\text{var}(m_y)$.
3. (15) identifies $\text{var}(u_t^F)$.
4. (12) 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 (12) through (14)), measurement error co-moves one-for-one with transitory shocks (similarly to most income models). In family-average autocovariances (equations (15) through (17)), 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.¹⁶

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.

¹⁶Note 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

$$\begin{aligned} \text{cov}(\Delta c_{i,j,t}, \Delta c_{i,j,t}) &= \delta_I^2 \text{var}(u_t^I) + \gamma_I^2 (1 + \theta)^2 \text{var}(e_t^I) + \delta_F^2 \text{var}(u_t^F) + \gamma_F^2 (1 + \theta)^2 \text{var}(e_t^F) \\ &\quad + \text{var}(m_{c,t}) + \text{var}(m_{c,t-1}) + \text{var}(\xi) \end{aligned} \quad (18)$$

$$\text{cov}(\Delta c_{i,j,t}, \Delta c_{i,j,t+1}) = -\text{var}(m_{c,t}) \quad (19)$$

$$\text{cov}(\Delta c_{i,j,t}, \Delta y_{i,j,t}) = \delta_I \text{var}(u_t^I) + \gamma_I (1 + \theta) \text{var}(e_t^I) + \delta_F \text{var}(u_t^F) + \gamma_F (1 + \theta) \text{var}(e_t^F) \quad (20)$$

$$\text{cov}(\Delta c_{i,j,t}, \Delta y_{i,j,t+1}) = \gamma_I (1 + \theta) (\theta - 1) \text{var}(e_t^I) + \gamma_F (1 + \theta) (\theta - 1) \text{var}(e_t^F) \quad (21)$$

$$\text{cov}(\Delta c_{i,j,t}, \Delta y_{i,j,t+2}) = -\gamma_I \theta (1 + \theta) \text{var}(e_t^I) - \gamma_F (1 + \theta) \theta \text{var}(e_t^F) \quad (22)$$

Family-average autocovariances

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta c_{j,t}}) = \delta_F^2 \text{var}(u_t^F) + \gamma_F^2 (1 + \theta)^2 \text{var}(e_t^F) + \frac{1}{n_i} (\text{var}(m_{c,t}) + \text{var}(m_{c,t-1}) + \text{var}(\xi)) \quad (23)$$

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta c_{j,t+1}}) = -\frac{1}{n_i} \text{var}(m_{c,t}) \quad (24)$$

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta y_{j,t}}) = \delta_F \text{var}(u_t^F) + \gamma_F (1 + \theta) \text{var}(e_t^F) \quad (25)$$

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

$$\text{cov}(\overline{\Delta c_{j,t}}, \overline{\Delta y_{j,t+2}}) = -\gamma_F \theta (1 + \theta) \text{var}(e_t^F) \quad (27)$$

Identification of consumption parameters is as follows:¹⁷

1. (19) or (24) identifies $\text{var}(m_{c,t})$.
2. (26) or (27) identifies γ_F .
3. (21) or (22) identifies γ_I .
4. (25) identifies δ_F .
5. (20) identifies δ_I .
6. Either of (18) or (23) 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

¹⁷Note 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 family size. This means that the distribution of income is a mixture distribution in which the components are defined by the size of the 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

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

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 $\text{var}(e_{t-1}^S)$.

$$\begin{aligned} \text{cov}(\Delta_2 y_{i,j,t}, \Delta_2 y_{i,j,t}) &= 2\text{var}(u_t^I) + (1 + \theta^2)\text{var}(e_t^I) + (1 + \theta^2)\text{var}(e_{t-2}^I) \\ &\quad + 2\text{var}(u_t^F) + (1 + \theta^2)\text{var}(e_t^F) + (1 + \theta^2)\text{var}(e_{t-2}^F) + 2\text{var}(m_y) \end{aligned} \quad (28)$$

$$\text{cov}(\Delta_2 y_{i,j,t}, \Delta_2 y_{i,j,t+2}) = -(1 + \theta^2)\text{var}(e_t^I) - (1 + \theta^2)\text{var}(e_t^F) - \text{var}(m_y) \quad (29)$$

$$\text{cov}(\overline{\Delta_2 y_{j,t}}, \overline{\Delta_2 y_{j,t}}) = 2\text{var}(u_t^F) + (1 + \theta^2)\text{var}(e_t^F) + (1 + \theta^2)\text{var}(e_{t-2}^F) + \frac{2}{n_j}\text{var}(m_y) \quad (30)$$

$$\text{cov}(\overline{\Delta_2 y_{j,t}}, \overline{\Delta_2 y_{j,t+2}}) = -(1 + \theta^2)\text{var}(e_t^F) - \frac{1}{n_j}\text{var}(m_y) \quad (31)$$

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: (31) identifies $\text{var}(e_t^F)$, (30) identifies $\text{var}(u_t^F)$, (29) identifies $\text{var}(e_t^I)$, and (28) identifies $\text{var}(u_t^I)$.

Consumption long-difference

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

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

$$\begin{aligned} \text{cov}(\Delta_2 c_{i,j,t}, \Delta_2 c_{i,j,t}) &= 2\delta_I^2 \text{var}(u_t^I) + 2\delta_F^2 \text{var}(u_t^F) + 2\gamma_I^2 (1+\theta)^2 \text{var}(e_t^I) + 2\gamma_F^2 (1+\theta)^2 \text{var}(e_t^F) \\ &\quad + \text{var}(m_{c,t}) + \text{var}(m_{c,t-2}) + 2\text{var}(\xi) \end{aligned} \quad (32)$$

$$\text{cov}(\Delta_2 c_{i,j,t}, \Delta_2 c_{i,j,t+2}) = -\text{var}(m_{c,t}) \quad (33)$$

$$\text{cov}(\Delta_2 y_{i,t}, \Delta_2 c_{i,t}) = 2\delta_I \text{var}(u_t^I) + 2\delta_F \text{var}(u_t^F) + \gamma_I (1+\theta)^2 \text{var}(e_t^I) + \gamma_F (1+\theta)^2 \text{var}(e_t^F) \quad (34)$$

$$\text{cov}(\Delta_2 y_{i,t+1}, \Delta_2 c_{i,t}) = -\gamma_I (1+\theta)^2 \text{var}(e_t^I) - \gamma_F (1+\theta)^2 \text{var}(e_t^F) \quad (35)$$

$$\text{cov}(\overline{\Delta_2 c_{i,t}}, \overline{\Delta_2 c_{i,t}}) = 2\delta_F^2 \text{var}(u_t^F) + 2\gamma_F^2 (1+\theta)^2 \text{var}(e_t^F) + \frac{1}{n_i} (\text{var}(m_{c,t}) + \text{var}(m_{c,t-2}) + 2\text{var}(\xi)) \quad (36)$$

$$\text{cov}(\overline{\Delta_2 c_{i,t}}, \overline{\Delta_2 c_{i,t+2}}) = -\frac{1}{n_i} \text{var}(m_{c,t}) \quad (37)$$

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

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

Identification

Identification of consumption parameters using long-difference covariances is analogous to the short difference covariances: (33) or (37) identifies $\text{var}(m_{c,t})$, (39) identifies γ_F , (38) identifies δ_F , (35) identifies γ_I , (34) identifies δ_I , and finally (32) or (36) identifies $\text{var}(\xi)$.

Appendix Tables

Table 8: 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	0.1031 (0.0074)	-0.0290 (0.0057)	-0.0055 (0.0061)	0.0430 (0.0027)	-0.0116 (0.0020)	0.0015 (0.0021)
1980	0.1019 (0.0101)	-0.0400 (0.0080)	-0.0017 (0.0086)	0.0386 (0.0030)	-0.0134 (0.0025)	-0.0051 (0.0021)
1981	0.1113 (0.0095)	-0.0433 (0.0081)	0.0073 (0.0071)	0.0432 (0.0032)	-0.0139 (0.0022)	0.0024 (0.0018)
1982	0.1164 (0.0099)	-0.0442 (0.0084)	-0.0009 (0.0048)	0.0434 (0.0029)	-0.0143 (0.0021)	-0.0021 (0.0018)
1983	0.1047 (0.0098)	-0.0315 (0.0059)	-0.0108 (0.0076)	0.0371 (0.0029)	-0.0105 (0.0017)	-0.0072 (0.0024)
1984	0.1007 (0.0077)	-0.0269 (0.0055)	0.0018 (0.0056)	0.0337 (0.0021)	-0.0052 (0.0015)	0.0018 (0.0015)
1985	0.1256 (0.0113)	-0.0329 (0.0066)	-0.0097 (0.0061)	0.0397 (0.0027)	-0.0082 (0.0019)	-0.0017 (0.0018)
1986	0.1192 (0.0096)	-0.0407 (0.0062)	-0.0142 (0.0068)	0.0454 (0.0032)	-0.0149 (0.0020)	-0.0045 (0.0020)
1987	0.1018 (0.0079)	-0.0235 (0.0058)	-0.0082 (0.0054)	0.0362 (0.0024)	-0.0079 (0.0017)	-0.0025 (0.0018)
1988	0.1033 (0.0083)	-0.0286 (0.0063)	-0.0101 (0.0051)	0.0352 (0.0025)	-0.0114 (0.0020)	-0.0017 (0.0017)
1989	0.0984 (0.0092)	-0.0293 (0.0059)	-0.0071 (0.0054)	0.0301 (0.0024)	-0.0077 (0.0021)	-0.0016 (0.0015)
1990	0.1131 (0.0101)	-0.0434 (0.0075)	0.0100 (0.0065)	0.0380 (0.0030)	-0.0148 (0.0023)	0.0051 (0.0025)
1991	0.1110 (0.0099)	-0.0405 (0.0084)	0.0032 (0.0074)	0.0341 (0.0025)	-0.0133 (0.0024)	-0.0038 (0.0023)
1992	0.1281 (0.0106)	-0.0660 (0.0092)	-0.0149 (0.0092)	0.0419 (0.0028)	-0.0172 (0.0024)	-0.0073 (0.0031)
1993	0.1612 (0.0132)	-0.0675 (0.0123)	0.0028 (0.0075)	0.0494 (0.0036)	-0.0195 (0.0033)	0.0040 (0.0035)
1994	0.2022 (0.0224)	-0.0666 (0.0124)	-0.0191 (0.0126)	0.0745 (0.0073)	-0.0252 (0.0037)	-0.0122 (0.0038)
1995	0.1840 (0.0229)	-0.0696 (0.0186)	NA	0.0613 (0.0055)	-0.0260 (0.0066)	NA
1996	0.1831 (0.0203)	NA	NA	0.0722 (0.0061)	NA	NA
1998	0.2080 (0.0223)	-0.0813 (0.0159)	NA	0.0768 (0.0075)	-0.0304 (0.0047)	NA
2000	0.1791 (0.0187)	-0.0396 (0.0129)	NA	0.0739 (0.0071)	-0.0257 (0.0058)	NA
2002	0.1834 (0.0277)	-0.0686 (0.0181)	NA	0.0685 (0.0069)	-0.0271 (0.0054)	NA
2004	0.1640 (0.0179)	-0.0613 (0.0111)	NA	0.0770 (0.0066)	-0.0274 (0.0042)	NA
2006	0.1566 (0.0144)	-0.0442 (0.0087)	NA	0.0701 (0.0053)	-0.0217 (0.0038)	NA
2008	0.1615 (0.0139)	-0.0738 (0.0100)	NA	0.0637 (0.0038)	-0.0319 (0.0039)	NA
2010	0.1552 (0.0125)	NA	NA	0.0633 (0.0054)	NA	NA

Table 9: 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	0.3453 (0.0282)	-0.1319 (0.0189)	0.1188 (0.0086)	-0.0419 (0.0049)
1982	0.3063 (0.0270)	-0.1283 (0.0208)	0.1098 (0.0082)	-0.0342 (0.0060)
1983	0.3124 (0.0256)	-0.1142 (0.0147)	0.1044 (0.0066)	-0.0338 (0.0041)
1984	0.2929 (0.0224)	-0.1316 (0.0156)	0.1060 (0.0069)	-0.0471 (0.0063)
1985	0.3380 (0.0277)	-0.1416 (0.0265)	0.1140 (0.0085)	-0.0390 (0.0072)
1986	0.3268 (0.0270)	NA	0.1076 (0.0078)	NA
1990	0.3106 (0.0230)	-0.1114 (0.0151)	0.0803 (0.0044)	-0.0267 (0.0031)
1991	0.3075 (0.0213)	-0.1553 (0.0203)	0.0957 (0.0048)	-0.0532 (0.0051)
1992	0.3481 (0.0283)	-0.1337 (0.0221)	0.1260 (0.0099)	-0.0412 (0.0069)
1993	0.3290 (0.0309)	-0.1559 (0.0314)	0.1158 (0.0093)	-0.0486 (0.0077)
1994	0.3153 (0.0343)	-0.1177 (0.0250)	0.1115 (0.0116)	-0.0393 (0.0133)
1995	0.3440 (0.0414)	-0.1776 (0.0448)	0.1396 (0.0160)	-0.0623 (0.0154)
1996	0.3444 (0.0451)	NA	0.1358 (0.0151)	NA
1998	0.3573 (0.0303)	-0.1492 (0.0281)	0.1131 (0.0088)	-0.0605 (0.0100)
2000	0.4142 (0.0336)	-0.1629 (0.0284)	0.1537 (0.0115)	-0.0566 (0.0097)
2002	0.4450 (0.0463)	-0.2228 (0.0564)	0.1651 (0.0128)	-0.0696 (0.0158)
2004	0.5401 (0.0580)	-0.2420 (0.0448)	0.2163 (0.0193)	-0.1014 (0.0154)
2006	0.4792 (0.0390)	-0.1708 (0.0281)	0.2081 (0.0150)	-0.0551 (0.0104)
2008	0.5499 (0.0414)	NA	0.1992 (0.0150)	NA

Table 10: Income Parameter Estimates

Year	Variance of Permanent Shocks			Variance of Transitory Shocks		
	Idiosyncratic	Fam-Agg	% Fam-Insurable	Idiosyncratic	Fam-Agg	% Fam-Insurable
1979-81	0.014 (0.005)	0.014 (0.003)	0.507 (0.111)	0.018 (0.007)	0.010 (0.004)	0.657 (0.090)
1982	0.017 (0.007)	0.013 (0.003)	0.560 (0.135)	0.024 (0.009)	0.012 (0.004)	0.663 (0.093)
1983	0.013 (0.006)	0.009 (0.003)	0.596 (0.123)	0.018 (0.008)	0.009 (0.004)	0.671 (0.112)
1984	0.018 (0.005)	0.016 (0.003)	0.530 (0.081)	0.017 (0.007)	0.001 (0.004)	0.935 (0.098)
1985	0.032 (0.008)	0.025 (0.003)	0.568 (0.074)	0.022 (0.007)	0.005 (0.004)	0.829 (0.077)
1986	0.015 (0.007)	0.019 (0.003)	0.434 (0.123)	0.024 (0.007)	0.013 (0.005)	0.644 (0.057)
1987	0.017 (0.006)	0.011 (0.003)	0.599 (0.109)	0.012 (0.007)	0.006 (0.004)	0.654 (0.127)
1988	0.027 (0.006)	0.012 (0.003)	0.687 (0.070)	0.013 (0.007)	0.009 (0.004)	0.579 (0.091)
1989	0.023 (0.007)	0.009 (0.003)	0.720 (0.085)	0.016 (0.009)	0.005 (0.004)	0.753 (0.086)
1990	0.022 (0.005)	0.013 (0.003)	0.628 (0.066)	0.021 (0.010)	0.012 (0.005)	0.643 (0.074)
1991	0.021 (0.006)	0.004 (0.003)	0.843 (0.115)	0.021 (0.009)	0.013 (0.005)	0.622 (0.117)
1992	0.002 (0.007)	0.006 (0.003)	0.241 (0.312)	0.049 (0.012)	0.018 (0.005)	0.738 (0.067)
1993	0.008 (0.009)	0.006 (0.003)	0.554 (0.286)	0.049 (0.013)	0.021 (0.007)	0.701 (0.056)
1994-96	0.024 (0.010)	0.013 (0.006)	0.656 (0.131)	0.044 (0.010)	0.027 (0.005)	0.622 (0.047)
1998-00	0.030 (0.009)	0.009 (0.005)	0.769 (0.105)	0.018 (0.010)	0.023 (0.006)	0.446 (0.134)
2002	0.023 (0.010)	0.006 (0.004)	0.784 (0.136)	0.032 (0.012)	0.021 (0.007)	0.604 (0.069)
2004	0.006 (0.005)	0.011 (0.004)	0.341 (0.196)	0.025 (0.008)	0.021 (0.008)	0.539 (0.075)
2006	0.013 (0.003)	0.009 (0.003)	0.587 (0.117)	0.017 (0.008)	0.018 (0.007)	0.485 (0.101)
2008-10	0.012 (0.004)	0.004 (0.003)	0.745 (0.153)	0.028 (0.007)	0.024 (0.007)	0.542 (0.063)
θ (Serial correlation of transitory shock)				0.152 (0.085)		
Income measurement error variance				0.014 (0.008)		

Standard errors based on 100 block bootstrap replications in parentheses.