Heterogeneity and Tests of Risk Sharing

Sam Schulhofer-Wohl*

June 17, 2009

Abstract

How well do people share risk? Do non-market institutions – charity, progressive taxes, transfer payments – make up for the lack of complete insurance markets? Or is risk sharing far worse than what complete markets could achieve? Standard risk-sharing regressions assume that any variation in households' risk preferences is uncorrelated with variation in income. I combine administrative and survey data to show that this assumption fails; risk-tolerant workers hold jobs where earnings carry more aggregate risk. The correlation makes risk-sharing regressions in the previous literature too pessimistic. I derive techniques that eliminate the bias, apply them to U.S. data, and find that the effect of idiosyncratic income shocks on consumption is practically small and statistically difficult to distinguish from zero.

Keywords: risk sharing, risk preferences, heterogeneity, imperfect insurance.
JEL classification: E21, E24.

*Department of Economics and Woodrow Wilson School of Public and International Affairs, Princeton University, 363 Wallace Hall, Princeton, NJ 08544, and NBER. E-mail: sschulho@princeton.edu. Phone: (609) 258-7392. I thank Pierre-André Chiappori, James Heckman, Robert Shimer and especially Robert Townsend for many invaluable discussions. John Cochrane, Angus Deaton, Andra Ghent, Paul Heaton, Steven Levitt, Derek Neal, Stephen O'Connell, Genevieve Pham-Kanter, Jesse Shapiro and numerous seminar participants also made helpful suggestions. The University of Chicago, the National Institute on Aging and the Chicago Center of Excellence in Health Promotion Economics provided generous financial support during my work on initial drafts of this paper.

1 Introduction

Risk pervades economic life. Workers lose their jobs or win promotions. Investments fail or succeed. Economies enter downturns or booms. The literature on risk sharing has focused on how well people are insured against idiosyncratic risks, such as a decrease in income. This paper investigates whether some people are also insured against aggregate risks, such as a recession that causes many families' incomes to fall, and the consequences of this aggregate risk sharing for measurements of insurance against idiosyncratic risk.

Although the idea may seem counterintuitive, aggregate risk is insured all the time. Someone who invests in Treasury bills rather than stocks, forgoing high average returns to reduce the variance of returns, effectively pays an insurance premium for protection against aggregate risk. Insurance against aggregate risk transfers it to those who are least risk averse.

Most empirical analyses of risk sharing, however, have concentrated on insurance against idiosyncratic risk. If insurance is Pareto efficient, households' consumption depends only on aggregate shocks and not at all on idiosyncratic ones. A long literature therefore tests for efficiency by regressing consumption on idiosyncratic shocks such as income, using time indicator variables to control for aggregate shocks. Insurance is found to be imperfect: The coefficient on income is almost always positive.[1]

In this paper, I show that ignoring insurance against aggregate risk can lead to mistaken conclusions about how well people are insured against idiosyncratic risk. Using time indicator variables to represent aggregate shocks, as in the existing literature, assumes that aggregate shocks affect all households equally. But when risk preferences vary across households, less-risk-averse households bear more aggregate risk, and their consumption moves more strongly with aggregate shocks. If less-risk-averse households' income also moves more strongly with

[1] Among many examples, full insurance has been rejected in data from the United States (Attanasio and Davis, 1996; Cochrane, 1991; Dynarski and Gruber, 1997; Hayashi et al., 1996), Côte d'Ivoire (Deaton, 1997), India (Munshi and Rosenzweig, 2009; Townsend, 1994), Nigeria (Udry, 1994) and Thailand (Townsend, 1995a,b). Mace (1991) does not reject efficiency in U.S. data, but Nelson (1994) overturns this result.

aggregate shocks, the usual regression will find a spuriously large correlation between income and consumption and can spuriously reject the null hypothesis of full insurance when insurance is actually perfect.

After showing the theoretical possibility of this bias in section 2, I demonstrate in section 3 that income *does* vary more with aggregate shocks for less-risk-averse households. I classify respondents to the Health and Retirement Study by how they answer hypothetical questions about taking risky jobs. Restricted-access Social Security records show that earnings are more volatile and more correlated with GDP for men who claim to be less risk averse.

The rest of the paper measures the relationship between idiosyncratic income shocks and consumption while accounting for the relationship between risk preferences and income processes. Section 4 derives econometric methods for testing the null hypothesis of full insurance. I focus on data with only a few observations per household and therefore treat households' preferences as nuisance parameters that must be eliminated from the equation. Section 5 extends the tests to allow estimates of the extent of partial insurance if full insurance is rejected. My approach is to interpret standard risk-sharing regressions as structural estimates of a simple model of imperfect insurance, in which households trade in complete markets but transferring resources between households is costly. The coefficient on income in a risk-sharing regression then measure the relative costs and benefits of risk sharing; the effect is large if transferring resources is very costly or households are not very risk averse. Section 6 applies the methods to income and consumption data from the Panel Study of Income Dynamics. Holding aggregate shocks constant, household consumption rises with income, but accounting for heterogeneous preferences reduces the effect by one-fourth to one-half. The effect is small and, in many specifications, statistically indistinguishable from zero. I reject the hypothesis that common-preferences regressions are correctly specified.

This paper is related to small but growing literatures on the relationship between preferences and income processes, and on risk sharing with heterogeneous preferences. Bonin et al.

(2007) show that more-risk-averse people hold jobs with less idiosyncratic risk, as measured by cross-sectional dispersion in self-reported earnings. They do not analyze macroeconomic risk or the time series properties of individual earnings, and because they use survey data on earnings instead of administrative data, they cannot distinguish measurement error from true variability in earnings. Fuchs-Schündeln and Schündeln (2005) analyze the endogeneity of income risk in the context of measuring the rate of precautionary savings: Ignoring the possibility that more prudent people choose less risky jobs will lead to downward-biased estimates of the importance of precautionary savings. Dynarski and Gruber (1997) test for full insurance while allowing heterogeneity in time preference but not in risk preference. Townsend (1994) allows heterogeneous risk preferences in part of his seminal paper on testing for full insurance, but he employs a short panel and runs a separate regression for each household, implying that hypotheses can be tested only by assuming a parametric distribution for the regression error term. Kurosaki (2001) makes a similar analysis. Dubois (2001) tests for full insurance with heterogeneous preferences while allowing the coefficient of relative risk aversion to be a function of observed household characteristics, but he rules out unobserved heterogeneity. Finally, Mazzocco and Saini (2009) test for full insurance with heterogeneous preferences in data from rural India. Their dataset has more than 100 observations per household, which allows them to avoid functional form assumptions about utility functions; to estimate each household's preferences, rather than treat preferences as nuisance parameters; and to test for heterogeneity in preferences. However, they do not estimate a model of partial insurance and so cannot say how economically important a rejection of full insurance is. Further, their methods do not apply to short panels and thus cannot be used to study risk sharing using available household-level data from the United States.[2]

[2]Mazzocco and Saini (2009) suggest that their methods could be applied to long panels of synthetic cohorts. However, constructing synthetic cohorts would average out all within-cohort heterogeneity, reducing the need for tests that allow heterogeneity in the first place.

2 Heterogeneity Bias in Risk-Sharing Tests

Diamond (1967) and Wilson (1968) show that a Pareto-efficient consumption allocation depends only on aggregate shocks and not at all on idiosyncratic shocks, with more risk-tolerant households bearing a larger share of the aggregate risk. This section shows that when preferences vary, standard tests of full insurance do not capture the true Pareto-efficient allocation and thus may generate spurious rejections of full insurance.

Let households' preferences depend on consumption, c, and leisure, ℓ. At each date t, denote the state of the economy by s_t. Each household i maximizes discounted expected utility:

$$\mathrm{E}_0 \sum_{t=0}^{T} \beta_i^t u_i[c_{it}(s_t), \ell_{it}(s_t)],$$

where u_i is concave in c. Let $X_{it}(s_t)$ denote household i's income in state s_t. For simplicity, assume income and leisure are given exogenously; alternatively, one can view the analysis as studying the optimal allocation of consumption given an already optimal choice of income and leisure. Likewise, one can either assume there is no storage or interpret the analysis as describing the optimal allocation conditional on an already Pareto-optimal aggregate amount of storage.

The optimal allocation can be decentralized, but it is convenient to study a social planner's problem. Given Pareto weights α_i, the planner assigns consumption to households to maximize the weighted sum of their discounted expected utilities,

$$\max_{\{c_{it}(s_t)\}} \sum_i \alpha_i \mathrm{E}_0 \sum_{t=0}^{T} \beta_i^t u_i[c_{it}(s_t), \ell_{it}(s_t)],$$

subject to the constraint that, for each date and state, aggregate income is at least as large as aggregate consumption:

$$\sum_i X_{it}(s_t) \geq \sum_i c_{it}(s_t) \quad \forall\, t, s_t.$$

Assume an interior solution and let $\Pr(s_t)\lambda_t(s_t)$ be the Lagrange multiplier on the aggregate resource constraint. The first-order condition for consumption is

$$\alpha_i \beta_i^t \frac{\partial}{\partial c} u_i[c_{it}^*(s_t), \ell_{it}(s_t)] = \lambda_t(s_t). \tag{1}$$

The crucial implication of (1) is the familiar fact that, under full insurance, a household's income does not affection its consumption at all, conditional on the aggregate shock $\lambda_t(s_t)$.

To bring the model to data, take constant relative risk aversion (CRRA) preferences,

$$u_i(c, \ell) = \frac{c^{1-\gamma_i}}{1-\gamma_i} \ell^{\xi_i}, \quad \gamma_i > 0,$$

as an approximation to the household's utility function.[3] This approximation can model differences in utility functions or differences in relative risk aversion among households that have identical non-CRRA preferences but varying wealth. Assume that consumption is measured with multiplicative error: Observed consumption is $c_{it} = e^{\epsilon_{it}} c_{it}^*$. Then the optimal allocation (1) satisfies

$$\log c_{it} = \frac{\log \alpha_i}{\gamma_i} - \frac{\log \lambda_t}{\gamma_i} + \frac{t \log \beta_i}{\gamma_i} + \frac{\xi_i}{\gamma_i} \log \ell_{it} + \epsilon_{it}. \tag{2}$$

Equation (2) says aggregate shocks λ_t have a larger effect on households that have smaller

[3] More formally, when preferences do not depend on leisure, define $h_i(c) = \log[u_i'(c)]$ and $f_i(w) = \log[h_i^{-1}(w)]$. Fix $\bar{c}$ and let $k_i - w/\gamma_i$ be the first-order Taylor series approximation to $f_i(w)$ around $\bar{w} = h_i(\bar{c})$. Notice that $\gamma_i = -1/f_i'(\bar{w}) = -\bar{c}u_i''(\bar{c})/u_i'(\bar{c})$, which is the coefficient of relative risk aversion at $\bar{c}$. Substitute the Taylor series into the definitions of $f_i(w)$ and $h_i(c)$, solve for $u_i'(c)$ and integrate to obtain $c^{1-\gamma_i}/(1-\gamma_i)$ as an approximation to $u_i(c)$ in the neighborhood of $\bar{c}$. This is the approximation implicit in Theorem 5 of Wilson (1968).

coefficients of relative risk aversion γ_i. Consumption rises faster for households with larger rates of time preference β_i or larger elasticities of intertemporal substitution $1/\gamma_i$; these differences in patience can also be interpreted as differences in consumption trends for households at different points in the life cycle. Income does not enter the equation at all, so one could test for full insurance by adding income to the equation:

$$\log c_{it} = \frac{\log \alpha_i}{\gamma_i} - \frac{\log \lambda_t}{\gamma_i} + \frac{t \log \beta_i}{\gamma_i} + \frac{\xi_i}{\gamma_i} \log \ell_{it} + g \log X_{it} + \epsilon_{it}, \tag{3}$$

then estimating the coefficient on income g and testing the hypothesis that $g = 0$.

Almost all previous analyses of risk sharing have not estimated equation (3), however.[4] The studies note that with identical risk and time preferences – $\gamma_i = \gamma$ and $\beta_i = \beta$ – one can define $\tilde{\lambda}_t = \lambda_t \beta^{-t}$ and estimate

$$\log c_{it} = \frac{\log \alpha_i}{\gamma} + \frac{1}{\gamma}(-\log \tilde{\lambda}_t) + \frac{\xi_i}{\gamma} \log \ell_{it} + g \log X_{it} + \epsilon_{it}^{equal}. \tag{4}$$

The central point of this paper is that the estimated coefficient on income in (4) is too large.[5]

The flaw in (4) is omitted variable bias. If the true model is (3) but a researcher mistak-

[4] Dynarski and Gruber (1997) allow heterogeneity in time preference but not in risk preference. Townsend (1994) and Kurosaki (2001) study approximately 100 households observed over 10 years in three Indian villages. They replace $\log \lambda_t$ with the logarithm of the households' total measured consumption C_t, then estimate (3) one household at a time, computing a coefficient on income for each household. This approach misspecifies the equation because $\log \lambda_t$ is not a linear function of $\log C_t$ when preferences vary. Also, the income coefficient is biased when the regression includes measured consumption aggregates (Ravallion and Chaudhuri, 1997), and with 10 observations per household, hypotheses about each household's income coefficient can be tested only by assuming a parametric form for the distribution of the errors.

[5] Most studies also impose a common coefficient on leisure in (4). Deaton and Paxson (1994) propose an alternative test of full insurance that also relies crucially on common preferences: According to (4), if consumption is separable from leisure, the cross-sectional variance of log consumption is constant over time under full insurance. If the cross-sectional variance rises over time, insurance is deemed imperfect. Equation (2) shows, however, that when households have different preferences, the cross-sectional variance of log consumption changes over time even under full insurance.

enly estimates (4), the error term in (4) is

$$\epsilon_{it}^{equal} = \left(\frac{1}{\gamma_i} - \frac{1}{\gamma}\right)(-\log\lambda_t) + \frac{\log\beta_i}{\gamma_i}t + \epsilon_{it}.$$

The least squares estimator of the coefficient on income in equation (4) is unbiased if $\text{Cov}(\log X_{it}, \epsilon_{it}^{equal}) = 0$, biased upward if $\text{Cov}(\log X_{it}, \epsilon_{it}^{equal}) > 0$ and biased downward if $\text{Cov}(\log X_{it}, \epsilon_{it}^{equal}) < 0$.

For simplicity, consider the case where all households have the same rate of time preference, so the time trend can be absorbed into the aggregate shock $\log\lambda_t$. Suppose that income depends on common and idiosyncratic shocks: $\log X_{it} = q_i m_t + u_{it}$. Assuming u_{it} and ϵ_{it} are i.i.d.,

$$\text{Cov}\big(\log X_{it}, \epsilon_{it}^{equal}\big) = \text{Cov}\big(q_i m_t, (1/\gamma_i - 1/\gamma)(-\log\lambda_t)\big). \tag{5}$$

Aggregate income and aggregate consumption are likely to be positively correlated. Thus the expression in (5) is positive – and the income coefficient in (4) is biased upward – if $\text{Cov}(q_i, 1/\gamma_i) > 0$. The parameter q_i is the effect of aggregate shocks on income; γ_i is the coefficient of relative risk aversion. Thus the income coefficient in (4) is *biased upward if income responds more strongly to aggregate shocks for less-risk-averse households.* A similar argument applies if rates of time preference are heterogeneous: The coefficient in (4) is *biased upward if income rises more quickly for more patient households.*

Other than the Dubois (2001) and Mazzocco and Saini (2009) papers on heterogeneity and risk sharing described in the introduction, previous studies have addressed the possible bias by finding assumptions under which it does not arise. For example, Cochrane (1991) tests how households smooth shocks other than income that he argues are uncorrelated with preferences, while Ogaki and Zhang (2001) allow decreasing relative risk aversion, $u_i(c) = (c-\theta)^{1-\gamma}/(1-\gamma)$, so relative risk aversion depends on the level of consumption but nothing else. My approach is different. In section 3, I show empirically that preferences *are* correlated with income processes even after controlling for observed characteristics. Then, in section 4, I derive estimators that eliminate the resulting bias in risk-sharing tests.

3 Occupation Choice and Risk Preferences

Simple models of both optimal portfolio choice and imperfect insurance suggest that, at least in a world without full insurance, less-risk-averse people are likely to hold riskier jobs. For example, suppose that each worker consumes his income and that all workers are equally productive in all jobs, but the mean and variance of income are higher in some jobs than in others. Then less-risk-averse workers will optimally choose the riskier jobs because these workers are more willing to accept a higher variance in return for a higher mean. This sorting will arise whether the income risk is idiosyncratic or aggregate risk, just as long as it is uninsured. Alternatively, suppose that transferring resources between workers after they receive their incomes is feasible but costly. Then an optimal assignment of workers to jobs will minimize the *ex post* transfer costs by placing each worker in a job where the income process resembles her optimal consumption allocation. Since less-risk-averse workers optimally bear a larger share of aggregate risk, they will hold jobs where income is more highly correlated with aggregate shocks. Of course, if workers' skills are also heterogeneous and if skills are correlated with risk preferences, the equilibrium sorting would not necessarily put less-risk-averse workers in riskier jobs. In addition, risk preferences and income risk could be correlated even in a world with full insurance, either because skills could be correlated with risk preferences or because preferences for smoothing labor supply over time (which affect the Pareto-optimal assignment of workers to jobs where productivity is more or less cyclical) could be correlated with preferences for smoothing consumption over time. The relationship between preferences and income processes is thus an empirical question, and I now turn to empirical evidence.

3.1 Data

The Health and Retirement Study (HRS), a panel survey of more than 22,000 Americans born in 1923 to 1947, contains both lifetime earnings histories and experimental questions that give evidence on respondents' preferences. These data permit a direct test of the hypothesis that incomes are more strongly correlated with aggregate shocks for less-risk-averse workers. On entering the study, each HRS respondent is asked:

> Suppose that you are the only income earner in the family, and you have a good job guaranteed to give you your current (family) income every year for life. You are given the opportunity to take a new and equally good job, with a 50-50 chance it will double your (family) income and a 50-50 chance that it will cut your (family) income by a third. Would you take the new job?[6]

Depending on how they answer, respondents are then asked about jobs that give a fifty-fifty chance of doubling income or of cutting it by 20 percent or 50 percent.[7] The majority reject even the job that might cut income by 20 percent, and I classify them as having low risk tolerance. I classify those who accept any risky job as having high risk tolerance. Respondents answered the question in 1992 or 1998, when they were 51 to 75 years old.

Barsky et al. (1997) establish that HRS respondents who say they would take risky jobs take more risks in real life. They are more likely to be self-employed, to be immigrants, to smoke and to drink heavily. They also invest more money in stocks and less in savings accounts. I focus here on whether more risk-tolerant workers face greater macroeconomic risks through their jobs.

[6] To avoid status-quo bias, the question in some years specifies that the respondent must leave the current job due to allergies and choose between two new jobs, one providing the current income for certain and the other offering a fifty-fifty chance of higher or lower income.

[7] Some respondents are also asked about jobs that might cut income by 10 percent or 75 percent.

The HRS includes restricted-access Social Security records that show each respondent's annual earnings in jobs and self-employment from 1951 to either 1991 or 1997, depending on when the respondent entered the HRS sample. The main weakness is that through 1979, earnings are top-coded at the Social Security taxable earnings maximum; 30 to 60 percent of observations on prime-age male workers are censored in each year. The pre-1980 data also omit jobs not covered by Social Security, including, in some years, most government jobs. Starting in 1980, the data come from W-2 tax forms, which include jobs not covered by Social Security and are top-coded at higher levels, but which omit earnings from self-employment.[8] I deflate earnings by the Consumer Price Index.

I restrict my analysis to men ages 23 to 61, since older and younger men are less likely to work full time, as are women in the HRS cohorts. I drop respondents who did not answer the risk tolerance questions or did not release their Social Security records, as well as observations with zero earnings. In the empirical analysis, I assume that observations with missing or zero earnings are missing at random.[9] I also drop individuals who have fewer than five uncensored earnings observations, because I want to analyze the time-series properties of each individual's earnings and, with fewer than five uncensored observations, some of the specifications I estimate are not identified.[10] After applying these restrictions, I have 115,424 annual observations for 4,090 workers. Table 1 gives summary statistics.

[8]In the Internet appendix, I report results using Social Security covered earnings for all years; they are largely similar.

[9]Treating zeros as censored low values would be inappropriate because – given that prime-age male workers are unlikely to remain unemployed for an entire calendar year – many zeros are likely to be for people who held jobs not covered by Social Security.

[10]Dropping workers with relatively few uncensored earnings observations introduces the risk of selection bias since the expected number of censored observations for a given worker depends in part on the variance of the worker's income. I defer discussion of this and other, related problems until after I present the results.

3.2 Analysis

The simplest way to study the relationship between risk preferences, earnings and aggregate shocks would be to compute the correlation of earnings with an aggregate variable such as gross domestic product for workers in each risk tolerance group. Because earnings are top-coded, I cannot directly compute this correlation. Instead, I measure the relationship between earnings and aggregate shocks in a regression framework, where I can account for censoring with a tobit model. The specifications I estimate take the form

$$\log\left(earnings_{it}\right) = \pi_{0i} + \pi_{1j}\log\left(GDP_t\right) + \mathbf{x}_{it}'\mathbf{\Pi}_j + \upsilon_{ijt}, \tag{6}$$

where j = low, high indexes risk-tolerance groups, π_{0i} is an individual random effect, and $\mathbf{x}_{it}$ is a vector of time-invariant and time-varying controls: education, experience, experience squared, and indicator variables for workers who are white, immigrants and veterans. I put GDP in real, per-capita terms, and I include a quadratic time trend in the controls to account for common trends in GDP and earnings.

I interpret (6) as an equation that a worker could use to predict his lifetime earnings profile. The coefficient on GDP shows how income varies over the business cycle. The experience and trend variables included as controls reflect how income would vary over time if there were no aggregate shocks. The variance of the error term measures idiosyncratic fluctuations in the worker's earnings. Equation (6) is an admittedly simplistic specification for earnings dynamics – for example, one might prefer to examine income growth rather than the level of income – but with top-coded data, I cannot take first differences to compute innovations to income. In addition, my goal here is only to study the correlation between earnings and aggregate shocks, not to estimate a fully specified model of the time series properties of individual earnings. In the Internet appendix, I use data from the Panel Study of Income Dynamics, where top-coding is minimal but risk-preference data are lacking, to show that regressions in differences would likely produce similar results.[11]

The tobit model for censoring assumes that the error term is normally distributed and has constant variance over time for each worker. These assumptions allow one to write down the likelihood for the censored data. Actually estimating equation (6) by maximum likelihood would be computationally challenging, however, because it would require calculating high-dimensional integrals of normal densities. Instead, I employ the Bayesian Markov Chain Monte Carlo method to find the posterior distribution of the parameters.[12] I allow the error term to be correlated across workers within a risk-tolerance group in each year: $\upsilon_{ijt} = \pi_{jt} + \nu_{it}$, where π_{jt} is a normally distributed random effect. I employ an uninformative prior for the regression coefficients and a standard weakly informative conjugate prior for the variances of the errors. The Internet appendix gives details.

The top panel of table 2 shows the results. A 1 percent change in GDP raises earnings by 0.4 percentage point more for workers in the high risk tolerance group than for workers in the low risk tolerance group. (For both groups, however, the elasticity is surprisingly high, given that the elasticity of aggregate earnings to GDP has to be near one unless labor income is highly procyclical. Below, I explore whether misspecification causes this problem.) The difference between the risk tolerance groups persists when I use aggregate personal consumption expenditures or aggregate wages and salaries as the macroeconomic variable in place of GDP. For GDP and aggregate wages and salaries, the difference is statistically significant at the 5 percent level (if we interpret Bayesian posterior density intervals as frequentist confidence intervals). The variance of idiosyncratic errors is also higher for workers in the high-risk-tolerance group. In other words, high-risk-tolerance workers bear both more aggregate risk and more idiosyncratic risk.

The more risk-tolerant group includes more immigrants and has, on average, about seven

[11] Although the PSID asked risk tolerance questions in 1996, most people who answered the questions have too few years of earnings data to analyze the relationship between their earnings and aggregate shocks.

[12] See algorithms 10 and 16 in Chib, 2001.

more months of education. To rule out the possibility that these differences in characteristics mean men in the more risk-tolerant group never had the opportunity to take less risky jobs, in the bottom panel of table 2 I repeat the regressions on a sample restricted to white, native-born men with exactly 12 years of education. The data contain 1,075 such workers, the largest race-by-education cell in the sample. Restricting the sample reduces the precision of the results but does not change the pattern of point estimates: Earnings vary more with GDP for more risk-tolerant men, who also face larger idiosyncratic fluctuations.

Equation (6) is misspecified if workers within a risk-tolerance group have different income processes. To see whether this biases my estimates, I estimate regressions in which the coefficients differ for each worker:

$$\log\left(earnings_{it}\right) = \pi_{0i} + \pi_{1i}\log\left(GDP_t\right) + \pi_{2i}experience_{it} + \pi_{3i}experience_{it}^2 + \nu_{it}. \quad (7)$$

In this specification, the individual-specific intercept π_{0i} absorbs all time-invariant individual characteristics such as education. Since a quadratic in experience is equivalent to a quadratic in time for any given worker, the experience variables now also account for common trends in GDP and earnings, or for differences in rates of time preference that affect job choice. I estimate these regressions by maximum likelihood one worker at a time, again using a tobit model to account for censored earnings, and examine the relationship between the estimated parameters and risk preferences in the cross section of workers.

Table 3 shows the means of the regression coefficients and of the variance of idiosyncratic errors in (7) for workers with high and low risk tolerance. Allowing the coefficients to vary across workers greatly reduces the precision of the results, but the point estimates tell the same story as before: High-risk-tolerance workers face more aggregate and idiosyncratic risk. For example, a 1 percent rise in per capita real GDP, relative to trend, increases earnings by an average of 2.0 percent for high-risk-tolerance workers but by just 1.65 percent for low-risk-tolerance workers. The differences in point estimates persist when I use different aggregate variables and when I restrict the sample to white, native-born men with exactly 12 years of education.

Allowing the coefficients to differ for each worker reduces the estimated average elasticity of earnings to aggregate variables, but the elasticity remains substantially larger than one. Aggregation effects are a possible explanation: If earnings are more procyclical for low-income workers – see, e.g., Solon et al. (1994) – the average elasticity of income to GDP, weighting workers equally, can exceed one even though the average elasticity, weighting by income, must be near one. Top-coding prevents me from computing HRS workers' permanent income and reweighting to test this hypothesis. However, reweighting is possible in my experiment with PSID data in the Internet appendix, and shows that the coefficient is statistically indistinguishable from one when workers are weighted by permanent income.

3.3 Caveats

The results support the idea that people share aggregate risk by sorting into jobs according to risk preferences, with several caveats. The largest concern is that because HRS respondents are interviewed late in life, people who express more risk tolerance may do so precisely because they experienced volatile incomes and learned to tolerate fluctuations. In other words, preferences may be a function of income processes rather than the other way around. If so, the results do not demonstrate that people deliberately sort into occupations on the basis of risk preferences. However, the results still show that risk tolerance is correlated with the aggregate income risk a person faces – and this correlation, not any particular direction of causality, is all that is needed to bias the results of risk-sharing regressions that assume common preferences. More broadly, a different model of efficient risk sharing is needed if preferences adapt to the shocks a person happens to experience.

Another concern is that differences in reported risk tolerance may reflect something other than differences in actual preferences. For example, holding the utility function constant, people who have low assets may be less willing to take risky jobs because they cannot self-insure. If people who have low assets also have less cyclical labor income, I would then find that people who report lower risk tolerance hold less risky jobs, without any true heterogeneity in preferences. However, Solon et al. (1994) find that low-income workers have *more* procyclical incomes, so accounting for the effect of assets on reported preferences would likely increase, not decrease, the measured differences in earnings risk.

Measurement problems are also a concern. I do not correct for noise in responses to the risk tolerance questions, but if any workers are misclassified, the true difference between the groups is larger than the estimated difference. Earnings as measured by Social Security records will spuriously move with GDP if workers shift during recessions from jobs covered by Social Security to jobs not covered by Social Security. Also, Social Security earnings may also omit some earnings from self-employment if people under-report their earnings to avoid taxes, so – even though Social Security covers self-employment – measured earnings could move spuriously with GDP if people shift into self-employment during recessions.

Finally, two kinds of non-random sample selection could affect the results. First, because HRS respondents are surveyed late in life, people who died young are selected out of the sample. If more risk-tolerant people take more potentially fatal risks, then the workers who appear in the HRS sample are likely to be less risk-tolerant than the population as a whole. The results thus reflect the relationship between earnings and risk tolerance among people who survive to late middle age, rather than among all people. Second, and more problematic, I exclude anyone with fewer than five uncensored earnings observations because the individual tobit regression is not necessarily identified for these workers. Among low-income workers, this selection criterion has little effect, because their earnings rarely exceed the Social Security taxable maximum. But among higher-income workers, selection on the number of uncensored observations matters. Consider two workers who have the same mean earnings but different variances, and assume the mean exceeds the Social Security cap. The worker with the higher variance will have, on average, more years of earnings below the cap. Thus, workers whose earnings have a high mean and low variance are less likely to appear in my sample. If the high-risk-tolerance group had substantially higher mean earnings, this non-random selection could cause me to find more variable earnings in the high-risk-tolerance group without any true sorting into different occupations. However, the summary statistics in table 1 show that the mean earnings of the two groups are virtually identical, reducing the likelihood that selection on the number of uncensored observations drives the results.

4 Econometrics: Robust Risk-Sharing Tests

This section develops econometric methods for accomplishing two goals. The first goal is to test the null hypothesis of full insurance – that is, the null that idiosyncratic income shocks do not affect consumption – while allowing risk preferences to be heterogeneous and correlated with income processes. We can do so by estimating equation (3) and testing the null that the income coefficient $g = 0$. The second goal is to determine whether risk-sharing tests that ignore heterogeneity are correctly specified. We can do so by testing whether estimates of the heterogeneous-preferences equation (3) and the common-preferences equation (4) yield the same results.

Two methods are known for estimating equations like (3) in samples with many households i and a few dates t. The equation can be treated as a factor model and estimated by minimizing a sum of squared residuals. Alternatively, Ahn et al. (2001) propose Generalized Method of Moments estimators for a class of equations that includes (3).

The two approaches are complementary. The factor approach imposes strong conditions on the distribution of the errors and assumes consumption and leisure are separable. However, if these assumptions hold, the factor test is valid whether preferences are correlated with income processes or not. Thus comparing results from the factor test with results from a common-preferences regression can show whether the common-preferences regression is misspecified. If we statistically reject the hypothesis that the factor model and common preferences estimates are the same, then we conclude the common-preferences regression is misspecified. By contrast, the GMM approach makes fewer assumptions but is valid only when preferences are correlated with income processes. It is meaningless to compare a GMM estimate of (3) with an estimate of the common preferences model (4): The estimates are guaranteed to differ since each is valid only under assumptions that make the other invalid.

The two methods also perform different experiments. The factor test analyzes how consumption responds to deviations of income from its mean over time. The GMM test analyzes how consumption responds to changes in income from one year to the next. Hayashi et al. (1996) note that if insurance is imperfect but households know about income shocks in advance and adjust consumption accordingly, consumption growth is uncorrelated with contemporaneous income growth. The factor model measures the relationship between consumption and income over longer periods and thus may detect more failures of insurance.

In the empirical analysis, I will apply the tests to household survey data. Because households may report income inaccurately, the tests must account for measurement error. In fact, doing so is crucial to demonstrating that ignoring heterogeneity leads to spurious rejections of full insurance. Suppose there is *not* full insurance and the true coefficient on income is positive. Classical measurement error in income will bias the estimated coefficient toward zero if the model is estimated by ordinary least squares (OLS). The lower the signal-to-noise ratio, the smaller the estimated coefficient. Adding controls to a regression reduces the signal-to-noise ratio in the regressors, so adding controls for heterogeneous preferences to an OLS risk-sharing regression will make the attenuation bias worse. In other words, if I use OLS and reject full insurance when assuming homogeneity but not when allowing heterogeneity, it could be either because there is full insurance but the homogeneous-preferences coefficient is biased away from zero, or because there is imperfect insurance but the heterogeneous-preferences coefficient is biased toward zero. To show that assuming homogeneity makes the test spuriously reject full insurance, I must show that accounting for heterogeneity changes the results even after accounting for measurement error. Dealing with measurement error is straightforward for GMM, which already uses instrumental variables. For the factor test, existing estimators assume regressors are measured without error; I derive an instrumental variables method that is valid when income is measured with error.

The rest of this section gives technical details of the tests and can be skipped by readers interested only in the results.

4.1 Notation

I consider an unbalanced panel of data on N households in T years. Let y_{it} stand for log consumption, x_{it} for log income and z_{it} for log leisure. Let ι_{it} equal one if household i is observed at t and equal zero otherwise. Let x_{it}, y_{it} and z_{it} equal zero if household i is not observed at t.

Following Wooldridge's (2005) analysis of panel random coefficients models, decompose leisure into a household-specific projection on the aggregate shock and time trend and a residual from this projection:

$$z_{it} = \tilde{a}_i^z + \tilde{b}_i^z d_t + \tilde{r}_i^z t + \ddot{z}_{it}^{het}, \quad \mathrm{E}[\ddot{z}_{it}^{het}(\tilde{a}_i^z + \tilde{b}_i^z d_t + \tilde{r}_i^z t)] = 0.$$

The residual $\ddot{z}_{it}^{het}$ is the idiosyncratic component of leisure. Substituting this decomposition into (3) yields

$$y_{it} = a_i + b_i d_t + r_i t + f z_{it} + g x_{it} + e_{it}^{het}, \quad e_{it}^{het} = f_i \ddot{z}_{it}^{het} + \epsilon_{it}. \tag{3'}$$

where $a_i = \log \alpha_i/\gamma_i + f_i \tilde{a}_i^z$, $1/\gamma_i + f_i \tilde{b}_i^z$, $r_i = \log \beta_i/\gamma_i + f_i \tilde{r}_i^z$ and $d_t = -\log \lambda_t$. The first goal is to estimate (3′) and test the null hypothesis that $g = 0$.

One can also decompose leisure into a projection on a household fixed effect and a common aggregate shock,

$$z_{it} = \tilde{a}_i^z + \tilde{b}^z d_t + \ddot{z}_{it}^{equal}, \quad \mathrm{E}[\ddot{z}_{it}^{equal}(\tilde{a}_i^z + \tilde{b}^z d_t)] = 0.$$

The common-preferences regression (4) can then be written

$$y_{it} = a_i + d_t + f z_{it} + g x_{it} + e_{it}^{equal}, \quad e_{it}^{equal} = f_i \ddot{z}_{it}^{equal} + \epsilon_{it}. \tag{4′}$$

The second goal is to test whether estimates of (3′) yield different conclusions about the null that $g = 0$ than estimates of (4′). (More precisely, the goal is to test the null hypothesis that the two equations give the same results.)

4.2 Factor model test

The factor model test begins by assuming that preferences are separable in consumption and leisure:

Assumption F.1. *$\xi_i = 0$ for all households i.*

Under assumption F.1, leisure z does not appear on the right-hand side of (3′), so $e_{it}^{het} = \epsilon_{it}$. Subtracting the within-household mean of each variable from (3′) removes the fixed effect a_i and yields:

$$u_{it}^y = b_i u_{it}^d + r_i u_{it}^t + g u_{it}^x + u_{it}^\epsilon, \tag{3″}$$

where

$$u_{it}^y = y_{it} - \frac{\sum_{t=1}^T \iota_{it} y_{it}}{\sum_{t=1}^T \iota_{it}} \text{ if } \iota_{it} = 1 \text{ and } u_{it}^y = 0 \text{ if } \iota_{it} = 0,$$

and where u^x_{it}, u^d_{it}, u^t_{it} and u^ϵ_{it} are defined similarly. Likewise, (4′) becomes

$$u^y_{it} = d_t + g u^x_{it} + u^\epsilon_{it}. \tag{4''}$$

Suppose that de-meaned leisure u^z_{is} at all dates s is uncorrelated with ϵ_{it} at each date t. If all households have the same preferences, (4″) can be estimated by instrumental variables, with u^z_{it} as an instrument for de-meaned income u^x_{it} and with time indicator variables included as regressors to account for the aggregate shocks d_t. The factor model test I derive next produces valid estimates of (3″) whether households have identical preferences or not.

I adopt several assumptions that are standard in instrumental variables applications:

Assumption F.2. $\mathrm{E}[x_{is}\epsilon_{it}|\boldsymbol{\iota}_i] = \mathrm{E}[b_i\epsilon_{it}|\boldsymbol{\iota}_i] = \mathrm{E}[r_i\epsilon_{it}|\boldsymbol{\iota}_i] = \mathrm{E}[\epsilon_{it}|\boldsymbol{\iota}_i] = 0$ *for all* i, s, t.

Assumption F.3. *Instead of observing actual log income,* x_{it}*, we observe a mismeasured variable* $x^*_{it} = x_{it} + v_{it}$.

Assumption F.4. $\mathrm{E}[\mathbf{z}_i\boldsymbol{\epsilon}_i{}'] = \mathrm{E}[\mathbf{z}_i\mathbf{v}_i'] = \mathbf{0}$.

Assumption F.2 says measurement error in consumption is random in the sense that it is uncorrelated with income, preferences and the pattern of missing data. Assumption F.3 says log income is measured with additive error or, equivalently, that income is measured with multiplicative error. Assumption F.4 says leisure is uncorrelated with measurement error in consumption and income.

I also adopt a homoscedasticity assumption that is standard in factor analysis:

Assumption F.5. $\mathrm{E}[\boldsymbol{\epsilon}_i\boldsymbol{\epsilon}_i{}'] = \sigma_i^2\mathbf{I}$ *and* $\mathrm{E}[\mathbf{v}_i\mathbf{v}_i'] = \zeta_i^2\mathbf{I}$ *for all* i.

Assumption F.5 says measurement error in both consumption and income is serially uncorrelated and homoscedastic within each household.[13]

[13]Testing this assumption may seem desirable and, in a balanced panel, would be straightforward. For

With these assumptions in hand, it is helpful to start by considering the least-squares factor estimator in the literature. Equation (3″) requires a normalization because, for any nonzero constants k_1, k_2 and k_3, we can replace d_t with $k_1 d_t + k_2 t + k_3$, b_i with b_i/k_1, r_i with $r_i - k_2 b_i/k_1$ and a_i with $a_i - k_3 b_i/k_1$ without changing the model. That is, the level, scale, sign and trend of the aggregate shock d_t are not identified. Define the estimators $\{\hat{b}_i, \hat{r}_i\}_{i=1}^N$, $\{\hat{d}_t\}_{t=1}^T$ and $\hat{g}$ as the solutions to

$$\min_{\substack{\{\tilde{b}_i,\tilde{r}_i\}_{i=1}^N,\\ \{\tilde{d}_t\}_{t=1}^T,\tilde{g}}} \frac{1}{N}\sum_{i=1}^N\sum_{t=1}^T \iota_{it}(u_{it}^y - \tilde{b}_i u_{it}^{\tilde{d}} - \tilde{r}_i u_{it}^t - \tilde{g} u_{it}^x)^2$$

$$\text{such that} \quad \sum_{t=1}^T \tilde{d}_t = \sum_{t=1}^T t\tilde{d}_t = 0,\ \sum_{t=1}^T \tilde{d}_t^2 = 1,\ \tilde{d}_1 > 0. \quad (8)$$

The constraints are normalizations. The first normalization fixes the level and trend of d_t, the second fixes the scale, and the third fixes the sign. To guarantee a unique solution, I also assume that there are at least four time periods, that each household is observed at least three times and that some households are observed four times.[14]

The estimator in (8) appears to suffer from an incidental parameters problem because the number of parameters increases as either N or T increases. We can eliminate the parameters indexed by i to solve this problem. Let $\mathbf{u}_i^x = [u_{i1}^x \cdots u_{iT}^x]'$ and define $\mathbf{u}_i^y$, $\mathbf{u}_i^d$, $\mathbf{u}_i^t$, $\mathbf{u}_i^\epsilon$ and $\mathbf{d}$ analogously. Let $\mathbf{\Delta}_i^d = [\mathbf{u}_i^d\ \mathbf{u}_i^t]$, and let $\tilde{\mathbf{\Delta}}_i^d$ denote the matrix in which $\tilde{d}_t$ replaces d_t. Also let $\mathbf{I}_i$ be the $T \times T$ matrix with ι_{it} on the diagonal and zeros elsewhere. The objective function in (8) is strictly convex, and first-order conditions are necessary and sufficient for a minimum.

example, to test homoscedasticity, one could estimate the parameters and then test whether the cross-sectional variance of residuals was the same at each date t. However, in an unbalanced panel, matters become more difficult. Continuing with the example of homoscedasticity, the cross-sectional variance of residuals does not have to be the same at each date even if $\mathrm{Var}[\epsilon_{it}]$ is the same at each date for any given household i, because different households i may be in the sample at different dates. A test of homoscedasticity would actually test the joint hypothesis that the errors are homoscedastic and that there is no selection on the variance of the error term. Since selection on the variance of the error term would not bias the test, rejecting this joint hypothesis would not prove that the test is invalid, and I do not pursue such tests here.

[14] Given only three time periods, we could fit the data perfectly with any g, d_0, d_1 and d_2 by choosing a_i, b_i and r_i to solve $y_{i0} = a_i + b_i d_0 + g x_{i0}$, $y_{i1} = a_i + b_i d_1 + r_i + g x_{i1}$ and $y_{i2} = a_i + b_i d_2 + 2r_i + g x_{i2}$.

Thus we can replace $\tilde{b}_i$ and $\tilde{r}_i$ in (8) with the corresponding first-order conditions to produce the following equivalent estimator of $\mathbf{d}$ and g:

$$(\hat{\mathbf{d}}, \hat{g}) = \min_{\tilde{\mathbf{d}},\tilde{g}} \frac{1}{N} \sum_{i=1}^{N} (\mathbf{u}_i^y - \tilde{g}\mathbf{u}_i^x)' \left[\mathbf{I}_i - \tilde{\mathbf{\Delta}}_i^d \left(\tilde{\mathbf{\Delta}}_i^{d}{}' \hat{\mathbf{\Delta}}_i^d \right)^{-1} \tilde{\mathbf{\Delta}}_i^{d}{}' \right] (\mathbf{u}_i^y - \tilde{g}\mathbf{u}_i^x), \tag{9}$$

again subject to normalizations on $\mathbf{d}$. The estimator in (9) employs a generalized fixed effects transformation. To see this, let $V(\mathbf{d}) = \left[\mathbf{I}_i - \mathbf{\Delta}_i^d (\mathbf{\Delta}_i^{d\prime} \mathbf{\Delta}_i^d)^{-1} \mathbf{\Delta}_i^{d\prime}\right]$. Since $V(\mathbf{d})\mathbf{\Delta}_i^d = 0$, we could – if $\mathbf{d}$ were known – multiply ($3''$) by $V(\mathbf{d})$ to obtain

$$V(\mathbf{d})\mathbf{u}_i^y = gV(\mathbf{d})\mathbf{u}_i^x + V(\mathbf{d})\mathbf{u}_i^\epsilon, \tag{10}$$

which does not contain the household fixed effects b_i and r_i. Because $\mathbf{d}$ is not known, (9) uses the estimated $V(\hat{\mathbf{d}})$ to transform the model and eliminate the fixed effects.

Kiefer (1980) shows that if ϵ_{it} is homoscedastic and serially uncorrelated – as guaranteed by assumption F.5 – and if there is no individual-specific intercept or trend, the estimator in (9) is consistent in a balanced panel as long as income is measured without error. The extension to individual-specific intercepts and trends and an unbalanced panel is straightforward.[15] However, if income is measured with error and there is not full insurance, the estimated income coefficient in (9) will be biased toward zero and the test allowing heterogeneity will be underpowered. I derive an instrumental variables version of the estimator to eliminate this attenuation bias. The idea of the estimator is that if $\mathbf{d}$ were known, the income coefficient g could be estimated from (10) using an instrument for x, while if g were known, $\mathbf{d}$ could be estimated by factor analysis regardless of measurement error in x. Because neither g nor $\mathbf{d}$ is known, the procedure iterates between estimating the two parameters.

The procedure requires a version of the usual rank condition for instrumental variables:

[15] See the Internet appendix.

Assumption F.6. $\mathrm{E}[\mathbf{z}_i' V(\tilde{\mathbf{d}})\mathbf{x}_i] \neq 0$ *for all* $\tilde{\mathbf{d}}$ *in a neighborhood of* $\mathbf{d}$.

Assumption F.6 says that the instrument is correlated with income after controlling for aggregate shocks, not just for the true aggregate shocks but for incorrect aggregate shocks in the neighborhood of the true shocks. The full rank condition must hold for incorrect aggregate shocks because, in any finite sample, the estimated aggregate shocks will not equal the true aggregate shocks.

Define an estimator of $\mathbf{d}$ and g by the following system of equations:

$$\tilde{\mathbf{d}}(\tilde{g}) = \arg\min_{\hat{\mathbf{d}}} \frac{1}{N} \sum_{i=1}^{N} (\mathbf{u}_i^y - \tilde{g}\mathbf{u}_i^{x^*})' \left[\mathbf{I}_i - \hat{\mathbf{\Delta}}_i^d \left(\hat{\mathbf{\Delta}}_i^{d\prime} \hat{\mathbf{\Delta}}_i^d \right)^{-1} \hat{\mathbf{\Delta}}_i^{d\prime} \right] (\mathbf{u}_i^y - \tilde{g}\mathbf{u}_i^{x^*}) \tag{11a}$$

$$\tilde{g}(\tilde{\mathbf{d}}) = \frac{\sum_{i=1}^{N} \mathbf{u}_i^{y\prime} \left[\mathbf{I}_i - \tilde{\mathbf{\Delta}}_i^d \left(\tilde{\mathbf{\Delta}}_i^{d\prime} \tilde{\mathbf{\Delta}}_i^d \right)^{-1} \tilde{\mathbf{\Delta}}_i^{d\prime} \right] \mathbf{u}_i^z}{\sum_{i=1}^{N} \mathbf{u}_i^{x^*\prime} \left[\mathbf{I}_i - \tilde{\mathbf{\Delta}}_i^d \left(\tilde{\mathbf{\Delta}}_i^{d\prime} \tilde{\mathbf{\Delta}}_i^d \right)^{-1} \tilde{\mathbf{\Delta}}_i^{d\prime} \right] \mathbf{u}_i^z} \tag{11b}$$

Equation (11a) uses the least squares procedure (9) to estimate $\mathbf{d}$, taking as given some value $\tilde{g}$ for the coefficient on income. If $\tilde{g} = g$, the measurement error in x_{it}^* simply increases the variance of the residuals and the least squares procedure will continue to provide a consistent estimator of $\mathbf{d}$. Hence $\tilde{\mathbf{d}}(\tilde{g})$ will be consistent for $\mathbf{d}$ if $\tilde{g}$ is a consistent estimator of g. Equation (11b) uses z_{it} as an instrument for x_{it}^* to estimate the income coefficient g. As in (10), knowledge of $\mathbf{d}$ is required to eliminate the household-specific preference parameters. Assumptions F.3, F.4 and F.6 guarantee that z_{it} is a valid instrument. Thus $\tilde{g}(\tilde{\mathbf{d}})$ will be consistent for g if $\tilde{\mathbf{d}}$ is a consistent estimator of $\mathbf{d}$. It follows that, if (11a) and (11b) together have a unique solution, this solution is a consistent estimator of $(g, \mathbf{d})$.

A pair $(\tilde{\mathbf{d}}, \tilde{g})$ that solves both (11a) and (11b) can be found iteratively: Start with a guess for $\tilde{g}$; use (11a) to compute $\tilde{\mathbf{d}}$; use this new $\tilde{\mathbf{d}}$ to produce a new $\tilde{g}$ using (11b); and repeat until convergence. In the data I analyze, the iteration always converges in 10 or fewer steps. I investigate whether the fixed point is unique by starting the iteration from several places and examining whether all give the same result. The estimator can be adapted to allow heterogeneity only in risk preference or only in time preference by omitting **d** or the time trend from the $\mathbf{\Delta}$ matrix. Interpreting the estimator in (11a) as just-identified GMM applied to moment conditions defined by (11b) and the first-order conditions of (11a) shows that the estimator is root-N consistent and asymptotically normal.[16]

The estimator of g in (11) allows me to perform hypothesis tests that meet both of the goals set out at the beginning of this section. First, to test the null of full insurance while allowing heterogeneity, I test the hypothesis that $g = 0$ using (11). Such a test has correct size as long as assumptions F.1 to F.6 hold. Second, to test the null that allowing heterogeneity does not change the test results – in other words, that the common-preferences regression (4) is correctly specified – I compare estimates of g from the heterogeneous-preferences model (11), $\hat{g}^{het}$, with estimates from the IV estimator based on the common-preferences equation ($4''$), $\hat{g}^{equal}$. Consider constructing a confidence interval for $\hat{g}^{het} - \hat{g}^{equal}$. If the interval excludes zero, one rejects the hypothesis that common-preferences and heterogeneous-preferences methods give identical results. Now, under the joint hypothesis that there is full insurance and that preferences are uncorrelated with income processes, the heterogeneous-preferences estimator and the common-preferences estimator are both consistent; if there is full insurance but preferences are correlated with income processes, only the heterogeneous-preferences estimator is consistent. Thus if the confidence interval for $\hat{g}^{het} - \hat{g}^{equal}$ excludes zero, one must either reject full insurance or reject the hypothesis that the common-preferences test gives correct results. In cases where the confidence interval for $\hat{g}^{het} - \hat{g}^{equal}$ excludes zero but where one does not reject full insurance using the heterogeneous-preferences test, one must conclude that the common-preferences regression gives incorrect results and is misspecified. The test has correct size as long as assumptions

[16]The model cannot actually be estimated by GMM applied to these moment conditions because they may not have a unique solution. For example, in a balanced panel, all eigenvectors of a particular matrix solve the first-order conditions of (11a).

F.1 to F.6 hold, under the maintained hypothesis of full insurance.

4.3 GMM tests

The GMM test does not require leisure to be separable from consumption but does require leisure to be strictly exogenous with respect to the error term:

Assumption G.1. *For all s and t,*

$$E[z_{is}e_{it}] = 0, \tag{12}$$

where e_{it} represents the error term in $(3')$ *or* $(4')$ *depending on which equation is estimated.*

Moment condition (12) forms the basis of GMM tests for whether $g = 0$ in $(3')$ and $(4')$ under a variety of assumptions about preferences. It differs in several ways from the assumptions of the factor model. Assumption G.1 potentially allows consumption and leisure to be nonseparable, and it says nothing about the variance or serial correlation of the errors. However, if leisure appears on the right-hand side of the equation, assumption G.1 also prohibits classical measurement error in leisure – in contrast to the factor model, which assumed measurement error in leisure.[17] Further, if leisure appears on the right-hand side, then a heterogeneous coefficient on leisure appears in the error term in (12): $e_{it} = \epsilon_{it} + f_i\ddot{z}_{it}$. Assumption G.1 thus also requires that $\mathrm{E}[z_{is}\ddot{z}_{it}f_i] = 0$. That is, the covariance of leisure with the idiosyncratic component of leisure must not depend on the random coefficient.[18] My instruments are leads and lags of leisure. It is difficult to know whether these instruments satisfy the conditional homoscedasticity requirement; I assume that they do or that

[17] It would be difficult to allow classical measurement error in leisure by employing lags of leisure as instruments: As shown below, the first, second and third lags of leisure will appear in the residuals $\hat{e}$ in the moment conditions, making it necessary to employ the fourth or higher lags as instruments – but such high lags are likely to be quite weak instruments. In results not reported here, I find no statistically significant relationship between income and the fourth lag of leisure.

[18] See Wooldridge (1997, 2003, 2005) and Heckman and Vytlacil (1998) for other, related assumptions under which least squares and instrumental variables estimators also recover the means of random coefficients.

the coefficients on leisure are identical across households, and I use a test of overidentifying restrictions to check my assumptions. Note that since $\ddot{z}_{it}^{equal} = \ddot{z}_{it}^{het} + (\tilde{b}_i^x - \tilde{b}^x)d_t + \tilde{r}_i^x t$, the conditional homoscedasticity requirement is less likely to hold for the common-preferences equation (4′) than for the heterogeneous-preferences equation (3′). Thus, as observed by Wooldridge (2005) for random-coefficients panel data models in general, heterogeneous coefficients on leisure will cause fewer problems in (3′) than in (4′) even if all households have the same preferences.

All my tests examine the relationship between consumption and income over a four-year time span. I use both long and short time differences of variables. To write these differences concisely, for any variable ζ, define $\Delta_s \zeta_{it} = \zeta_{it} - \zeta_{i,t-s}$ and $\Delta_s^2 \zeta_{it} = (\zeta_{it} - \zeta_{i,t-1}) - (\zeta_{i,t-s} - \zeta_{i,t-s-1})$.

Throughout, I consider only a subset of the T^2 moment conditions that could be generated from (12). Long leads and lags of leisure are likely to be weak instruments. For example, leisure at some date is unlikely to be a good predictor of income 20 years in the future. GMM estimators that use many weak instruments can have poor finite sample properties (Stock et al., 2002). In addition, not all leads and lags are available for any given observation in an unbalanced panel. I therefore limit myself to a relatively small set of relatively strong instruments: leisure and its first lag.[19] Let $\mathbf{h}_{it}$ denote the vector of instruments $[1\ z_{it}\ z_{i,t-1}]'$.

Consider taking differences of (4′) to eliminate the household-specific intercept:

$$\Delta_3 y_{it} = \Delta_3 d_t + f\Delta_3 z_{it} + g\Delta_3 x_{it} + \Delta_3 e_{it}^{equal}.$$

(If preferences are separable in consumption and leisure, omit leisure from the right-hand side.) If leisure and its first lag – $z_{it}, z_{i,t-1}$ – are uncorrelated with the error e_{it}^{equal}, the

[19] Alternatively, one could use the full set of moment conditions but employ an estimator that has better finite-sample properties, such as generalized empirical likelihood (Newey and Smith, 2004). GEL methods involve optimizing highly nonlinear objective functions and proved to be computationally infeasible given the large number of nonlinear moment conditions available here.

following moment conditions hold:

$$\mathrm{E}\left[\mathbf{h}_{it}(\Delta_3 y_{it} - \Delta_3 d_t - f\Delta_3 z_{it} - g\Delta_3 x_{it})\right] = \mathbf{0}, \quad t = 4, \ldots, T. \tag{13}$$

In an unbalanced panel, the moment conditions for each date t can be constructed using all observations with data at $t-3$, $t-2$, $t-1$ and t. Not using income as an instrument prevents attenuation bias due to measurement error in income. Equation (13) lists $3(T-3)$ moment conditions in $T-1$ parameters $(\Delta_3 d_4, \ldots, \Delta_3 d_T, f, g)$. The parameters can be estimated by GMM and the $2T-8$ overidentifying restrictions tested using a chi-squared statistic. To test for full insurance while assuming all households have the same risk and time preferences, I test whether the estimate of g based on the moment conditions (13) is zero.

Next suppose that households have identical risk preferences but different rates of time preference. Then the aggregate shocks in (3′) can still be represented by time dummy variables, but the right-hand side includes a household-specific time trend, which can be eliminated by taking second differences:

$$\Delta_2^2 y_{it} = \Delta_2^2 d_t + f\Delta_2^2 z_{it} + g\Delta_2^2 x_{it} + \Delta_2^2 e_{it}^{het}.$$

If leisure and its lag are uncorrelated with the error term, these moment conditions hold:

$$\mathrm{E}\left[\mathbf{h}_{it}(\Delta_2^2 y_{it} - \Delta_2^2 d_t - f\Delta_2^2 z_{it} - g\Delta_2^2 x_{it})\right] = \mathbf{0}, \quad t = 4, \ldots, T. \tag{14}$$

The parameters can again be estimated by GMM. A test for whether $g = 0$ using these moment conditions is a valid test for full insurance under assumption G.1 and the maintained hypothesis that all households have the same risk preferences, whether time preferences vary or not.

Consider eliminating the risk preference parameters b_i from (3′). If households have the same rate of time preference, but risk preferences vary, the second difference of (3′) is

$$\Delta_2 y_{it} = b_i \Delta_2 d_t + f \Delta_2 z_{it} + g \Delta_2 x_{it} + \Delta_2 e_{it}. \tag{15}$$

Equation (15) is equivalent to the model studied by Ahn et al. (2001):[20] $\tilde{y}_{it} = b_i \tilde{d}_t + \tilde{\mathbf{x}}_{it}' \boldsymbol{\theta} + \tilde{e}_{it}$. Ahn et al. propose a quasi-differencing estimator. Equation (15) implies that

$$\begin{aligned}\Delta_2 y_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 y_{i,t-1} = f\left(\Delta_2 z_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 z_{i,t-1}\right) \\ + g\left(\Delta_2 x_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 x_{i,t-1}\right) + \Delta_2 e_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 e_{i,t-1},\end{aligned}$$

which does not contain the household-specific risk preference parameter b_i. Then the following moment conditions hold for $t = 4, \ldots, T$:

$$\mathrm{E}\left[\mathbf{h}_{it} \left\{ \begin{array}{r} \Delta_2 y_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 y_{i,t-1} - f\left(\Delta_2 z_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 z_{i,t-1}\right) \\ -g\left(\Delta_2 x_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 x_{i,t-1}\right) \end{array} \right\}\right] = \mathbf{0}. \tag{16}$$

A test for whether $g = 0$ using moment conditions (16) is a valid test for full insurance under assumption G.1 and the maintained hypothesis that all households have the same time preferences, whether risk preferences vary or not.

Second differences of (3′) can be quasi-differenced to produce moment conditions valid when both risk and time preferences vary: For $t = 4, \ldots, T$,

$$\mathrm{E}\left[\mathbf{h}_{it} \left\{ \begin{array}{r} \Delta_1^2 y_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 y_{i,t-1} - f\left(\Delta_1^2 z_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 z_{i,t-1}\right) \\ -g\left(\Delta_1^2 x_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 x_{i,t-1}\right) \end{array} \right\}\right] = \mathbf{0}. \tag{17}$$

Under assumption G.1, a test for whether $g = 0$ using moment conditions (17) is a valid test for full insurance regardless of variation in risk or time preferences.

[20] Indeed, Ahn et al. (2001) mention risk-sharing tests as a possible application of their model.

Identification of the parameters in all four sets of moment conditions – (13), (14), (16) and (17) – requires that the instruments (differences of leisure) be correlated with the right-hand-side variables (differences of income and leisure). Ahn et al. (2001) show that, in addition, identification in (16) and (17) requires that at least one instrument be correlated with b_i. If this condition fails, the moment conditions (16) and (17) have infinitely many solutions and the GMM estimators are not defined. To see why an instrument must be correlated with b_i, consider the case where $f = g = 0$, so income and leisure do not appear on the right-hand side. Then the moment conditions in (16) reduce to the cross-sectional regression of $\Delta_2 y_{it}$ on $\Delta_2 y_{i,t-1}$ using z_{it} and its lags as an instrument for $\Delta_2 y_{i,t-1}$. Identification requires that $\Delta_2 y_{i,t-1}$ be correlated with z_{it} and its lags. But $\Delta_2 y_{i,t-1} = b_i \Delta_2 d_{t-1} + \Delta_2 e_{i,t-1}$; $e_{i,t-1}$ is assumed uncorrelated with z_{it}; and $\Delta_2 d_{t-1}$ does not vary in a cross-section of households at a given date. Hence $\Delta_2 y_{i,t-1}$ is correlated with z_{it} only if b_i is correlated with z_{it}.

The only way to avoid assuming that b_i is correlated with z_{it} is to somehow restrict the distribution of the error term, since then one can identify $\mathbf{d}$ from moment conditions of the form $\mathrm{E}[(\mathbf{u}_i^y - g\mathbf{u}_i^x - r_i\mathbf{u}^t)(\mathbf{u}_i^y - g\mathbf{u}_i^x - r_i\mathbf{u}^t)'] = \mathrm{E}[b_i^2]\mathbf{u}^d\mathbf{u}^{d\prime} + \boldsymbol{\Sigma}$, where $\boldsymbol{\Sigma} = \mathrm{E}[\mathbf{u}_i^e\mathbf{u}_i^{e\prime}]$ is known up to a small number of parameters. In the case of an unbalanced panel, there would be different moment conditions for each possible pattern of missing data: $\mathrm{E}[(\mathbf{u}_i^y - g\mathbf{u}_i^x - r_i\mathbf{u}_i^t)(\mathbf{u}_i^y - g\mathbf{u}_i^x - r_i\mathbf{u}_i^t)'|\boldsymbol{\iota}_i] = \mathrm{E}[b_i^2|\boldsymbol{\iota}_i]\mathbf{u}_i^d\mathbf{u}_i^{d\prime} + \mathrm{E}[\mathbf{u}_i^e\mathbf{u}_i^{e\prime}|\boldsymbol{\iota}_i]$. Unless one is willing to assume that the error covariance matrix $\mathrm{E}[\mathbf{u}_i^e\mathbf{u}_i^{e\prime}|\boldsymbol{\iota}_i]$ does not depend on the pattern of missing data $\boldsymbol{\iota}_i$, this implies a vast number of moment conditions each involving only a small number of observations, and GMM estimation would be infeasible.[21] The factor estimator, which in effect assumes $\mathrm{E}[\mathbf{u}_i^e\mathbf{u}_i^{e\prime}|\boldsymbol{\iota}_i] = \mathrm{E}[\sigma_i^2|\boldsymbol{\iota}_i](\mathbf{I} - \boldsymbol{\iota}_i\boldsymbol{\iota}_i'/\boldsymbol{\iota}_i'\boldsymbol{\iota}_i)$, is a special case where it is feasible to use moment conditions based on the distribution of the error term in an unbalanced panel.

Leisure is correlated with the risk preference parameter b_i, and the parameters in (16)

[21] In a panel of length T, there are 2^T possible patterns of missing data, so the number of possible moment conditions would be $O(2^T)$. The dataset in this paper has $T = 22$, implying more than 4 million moment conditions.

and (17) are identified, if and only if the common-risk-preferences models (13) and (14) are misspecified. Moment conditions (16) and (17) therefore may produce different results than (13) and (14) whether risk preferences vary or not; unlike with the factor model, comparing results between GMM estimators will not show whether risk preferences vary. However, if we maintain that heterogeneity in risk preferences is important, the moment conditions in (16) and (17) estimate the effect of income on consumption without using the factor model's assumptions about homoscedasticity, no serial correlation and separability.

All of the moment conditions examine the association between income and consumption over a four-year period. However, the particular income variations that are examined differ. Conditions (13) look at innovations to income, conditions (14) look at the difference between two consecutive innovations, and conditions (16) and (17) consider even more complicated differences. The need for quasidifferencing in (16) and (17) makes it impossible to examine exactly the same innovations in all four cases, but I try to make the results as comparable as possible by using the same instruments and a four-year period in each case.

5 Interpreting Risk-Sharing Regressions Under Partial Insurance

Anecdotal evidence tells us that at least some idiosyncratic shocks – for example, winning a large lottery jackpot – are not fully insured. It is therefore interesting to go beyond simple tests of the null hypothesis of full insurance and investigate the magnitude of departures from full insurance: Are failures of insurance economically important or not? In this section, I develop a simple model that shows how coefficients in risk-sharing reflections measure the extent of partial insurance. I then derive econometric assumptions under which the tests of full insurance discussed above can be interpreted as providing structural estimates of this partial-insurance parameter.

5.1 Model

The framework is identical to that of the full-insurance model in section 2, except that transferring resources between households to provide insurance is costly. Specifically, if household i has income X and consumes $c \neq X$, an additional quantity $\phi_i h(X, c)$ of the consumption good is destroyed, where $\phi_i \geq 0$. The parameter ϕ_i measures how difficult it is to insure household i. I assume that $h(X, c) = 0$ if $X = c$ and $h(X, c) > 0$ otherwise, and that h is convex and twice differentiable, with $\partial^2 h / \partial c \partial X < 0$. There is no household-level storage. If aggregate storage (implemented by a social planner) is feasible, then, as with full insurance, the model should be interpreted as describing the optimal allocation conditional on a prior aggregate decision about how much to store.

This model of transactions costs does not correspond to any real-world institution. However, it would be difficult to formally model all of the myriad ways that households share risk, from insurance contracts to financial transactions to informal gifts between friends and relatives. I interpret the transactions costs in my model as a reduced form for all of the institutions that households use to share risk and all of the information and incentive problems that make these institutions less than ideal. Formally modeling only one risk-sharing institution, such as the optimal contract under limited commitment (e.g., Ligon et al., 2002), would also be a reduced form, since it would set aside all other institutions that are present in reality. The recent literature has also proposed reduced-form imperfect-insurance models based on the permanent income hypothesis: Some shocks are fully insured while households use bond-holdings to self-insure against other shocks, and the relative variances of the two kinds of shocks characterize the overall degree of insurance (Blundell et al., 2008; Heathcote et al., 2009). Although such models may feel more realistic, they are not consistent with the excess sensitivity of consumption to anticipated income shocks (e.g., Flavin, 1981). My model is compatible with excess sensitivity since transactions costs keep households from smoothing predictable as well as unpredictable fluctuations.

The assumption of no household-level storage is admittedly a weakness. It implies that the cost of smoothing over time is the same as the cost of smoothing across states. Although a model with smaller costs of smoothing over time than across states would be more realistic, such a model would not lead to an intertemporally separable regression such as (3) or (4). Thus, for purposes of interpreting the coefficient on income in a standard risk-sharing regression, the model constructed here may be more helpful.

As under full insurance, it is convenient to find the optimal allocation in the model with costly transfers via a social planner's problem. The planner's objective function is the same as before, but the planner now faces the constraint that, for each date and state, aggregate income must be at least as large as aggregate consumption plus the cost of transfers between households:

$$\sum_i X_{it}(s_t) \geq \sum_i c_{it}(s_t) + \sum_i \phi_i h[X_{it}(s_t), c_{it}(s_t)] \quad \forall\, t, s_t.$$

Assuming an interior solution, the first-order condition is now

$$\alpha_i \beta_i^t \frac{\partial}{\partial c} u_i[c_{it}^*(s_t), \ell_{it}(s_t)] = \lambda_t(s_t) \left(1 + \phi_i \frac{\partial}{\partial c} h[X_{it}(s_t), c_{it}^*(s_t)]\right). \tag{18}$$

Equation (18) shows that transfer costs distort the allocation away from the first best. In particular, holding leisure constant, consumption rises with income: Differentiating (18) with respect to income and applying the implicit function theorem yields

$$\frac{\partial c_{it}^*}{\partial X_{it}(s_t)} = \frac{\lambda_t \phi_i \frac{\partial^2 h}{\partial c \partial X}}{\alpha_i \beta_i^t \frac{\partial^2 u_i}{\partial c^2} - \lambda_t \phi_i \frac{\partial^2 h}{\partial c^2}} = \frac{\phi_i \frac{\partial^2 h}{\partial c \partial X}}{-\left(1 + \phi_i \frac{\partial h}{\partial c}\right) \left| \frac{\partial^2 u_i}{\partial c^2} \Big/ \frac{\partial u_i}{\partial c} \right| - \phi_i \frac{\partial^2 h}{\partial c^2}} > 0.$$

The magnitude of the relationship between consumption and income depends on the relative costs and benefits of risk sharing. Income has a large effect on consumption if ϕ_i is large, in which case risk sharing is expensive, or if the coefficient of absolute risk aversion $|(\partial^2 u_i/\partial c^2)/(\partial u_i/\partial c)|$ is small, in which case the household does not greatly mind fluctuations in consumption.

To take the model to data, we need a functional form for h. Let the cost of transferring resources between households be

$$h(X,c) = \frac{c}{2}\left(\log\frac{c}{X}\right)^2,$$

which satisfies the stated assumptions on h as long as $\log(X/c) < 1$. The optimal allocation (18) then satisfies

$$\log c_{it} = \frac{\log\alpha_i}{\gamma_i} - \frac{\log\lambda_t}{\gamma_i} + \frac{t\log\beta_i}{\gamma_i} + \frac{\xi_i}{\gamma_i}\log\ell_{it} - \frac{1}{\gamma_i}\log\left[1 + \phi_i\log\frac{c_{it}}{X_{it}} + \frac{\phi_i}{2}\left(\log\frac{c_{it}}{X_{it}}\right)^2\right] + \epsilon_{it}.$$

For $\phi_i \log(c_{it}/X_{it})$ close to zero – that is, when the resources lost to imperfect risk sharing are small – this equation is approximately equivalent to

$$\log c_{it} = \frac{\log\alpha_i}{\phi_i+\gamma_i} + \frac{1}{\phi_i+\gamma_i}(-\log\lambda_t) + \frac{\log\beta_i}{\phi_i+\gamma_i}t + \frac{\xi_i}{\phi_i+\gamma_i}\log\ell_{it} + \frac{\phi_i}{\phi_i+\gamma_i}\log X_{it} + \epsilon_{it}. \quad (19)$$

Equation (19) says income has a larger effect on households that are less risk averse (small γ_i) or more difficult to insure (large ϕ_i).[22] As insurance costs go to zero, income does not affect consumption at all, while as insurance costs go to infinity, consumption moves one-for-one with income.

Let g denote the mean effect of income on consumption: $g \equiv \mathrm{E}[\phi_i/(\phi_i+\gamma_i)]$. Then we

[22]One can derive the same equation in levels without any approximation by assuming constant absolute risk aversion preferences, $h(X,c) = (c-X)^2/2$ and additive measurement error, but the non-negativity constraint on measured consumption makes additive errors statistically unattractive.

can rewrite (19) as

$$\log c_{it} = \frac{\log \alpha_i}{\phi_i + \gamma_i} + \frac{1}{\phi_i + \gamma_i}(-\log \lambda_t) + \frac{\log \beta_i}{\phi_i + \gamma_i} t + \frac{\xi_i}{\gamma_i} \log \ell_{it} + g \log X_{it} + \epsilon_{it}^{het},$$

$$\epsilon_{it}^{het} = \epsilon_{it} + \left(\frac{\phi_i}{\phi_i + \gamma_i} - g \right) \log X_{it}, \quad (3''')$$

which is the same risk-sharing regression proposed in (3). The estimated income coefficient in risk-sharing regressions such as (3) and (4) thus measures the average relative costs and benefits of risk sharing, $g \equiv \mathrm{E}[\phi_i/(\phi_i + \gamma_i)]$. However, just as with full insurance, omitted-variable bias will distort the estimate of g if preferences are heterogeneous but we estimate the common-preferences equation (4) instead of the heterogeneous-preferences equation (3). Since the bias in the estimate of g is likely upward, estimates of (4) may lead us to conclude that insurance is worse than it actually is.

5.2 Econometrics

Section 4 derived methods for testing the null hypothesis of full insurance based on testing whether $g = 0$ in (3). Because the methods generate estimates of g, I can also use them to measure the extent of partial insurance if the null of full insurance is rejected. However, I must then ensure the estimates of g are consistent even under partial insurance. This requires one assumption on top of those made to test the null of full insurance, described below.

Given the additional assumption, one can interpret the estimates of g as measuring the extent of partial insurance as long as one accepts the approximations inherent in the reduced-form partial insurance model of this paper. The most serious approximation is that the cost of smoothing across time is the same as the cost of smoothing across states. Also, the linearization of the first-order condition (18) to obtain the risk-sharing regression (3$'''$) is an approximation. The differencing operators used to remove household fixed effects could amplify any biases from this linear approximation. Because different numbers and kinds of differencing operators are used to estimate the models with and without heterogeneity, differences in estimates that do and do not allow heterogeneity may reflect differences in the approximation bias rather than the omitted variable bias from heterogeneity. Estimates of the partial insurance model therefore need to be interpreted with more caution than tests of the null hypothesis of full insurance.

The rest of this subsection gives details of the additional assumption required to estimate the partial-insurance model and may be skipped by readers interested only in the results.

The error term in ($3'''$) contains a heterogeneous coefficient on income, while the error term in (3) does not contain a heterogeneous coefficient on income because, under the null of full insurance, the coefficient on income is zero for all households. To estimate the partial insurance model, I therefore assume:

Assumption P.1. $E[z_{is}\ddot{x}_{it}g_i] = 0$, *where* $g_i = \phi_i/(\phi_i + \gamma_i) - g$ *and where* $\ddot{x}_{it}$ *is the residual from a household-specific projection of log income on a time trend and the aggregate shock, as in* (4.1).

Assumption P.1 parallels the exogeneity assumption for the GMM test in the case where consumption and leisure are nonseparable. It holds trivially if $\phi_i = k\gamma_i$ for some constant k and all households i, so that the coefficient on income is the same for all households. Otherwise, the assumption implies that the covariance between income and the instruments does not depend on the heterogeneous coefficient. It is difficult to know whether this assumption holds; in the GMM estimates, I check it by testing overidentifying restrictions.

Under assumption P.1 as well as assumptions F.1 to F.6, the factor method will provide a consistent estimator of g in ($3'''$). Under assumptions P.1 and G.1, the GMM method will provide a consistent estimator of g in ($3'''$).

6 Robust Risk-Sharing Regressions: Results

This section estimates the effects of idiosyncratic income fluctuations on consumption in equations (3) and (4).

6.1 Data

I analyze data on consumption, income and leisure from the Panel Study of Income Dynamics. The PSID, which has followed thousands of American families since 1968, is among the only panels long enough to permit estimation of a model with multiple household-specific parameters. However, the PSID measures only food consumption.[23] While food is not an ideal proxy for total consumption, it may be more likely to be time separable, as the expected utility formulation assumes.

I use data from the 1974 to 1997 waves of the PSID – a period over which the definitions of food and income variables remained roughly constant – but drop 1988 and 1989, when no food consumption data were collected. I define income as the household's total money income except for Aid to Families with Dependent Children, Supplemental Security Income, other welfare payments, unemployment insurance, worker's compensation, and help from relatives, all of which represent insurance rather than shocks that should be insured. I convert income data to real terms using the Consumer Price Index and food data using the food and beverages component of the CPI. I compute the estimates using total consumption and using consumption adjusted by the number of adult equivalent household members. I measure leisure by 8,760 (the number of hours in a year) minus hours worked by the head of household.[24] I restrict the sample to households with consumption, income and leisure

[23]The Consumer Expenditure Survey, which measures more consumption than the PSID, contains only two observations on income per household – too short a panel for my purposes.

[24]In the web appendix, I report results from a robustness check that defines the leisure variable to include the spouse's leisure for couple-headed households; the parameter estimates are very similar to those obtained using only the head's leisure.

data in at least four consecutive years, the minimum number of observations required for the GMM estimators. The data are well known; Appendix A describes them in detail.

Theory describes the optimal allocation of consumption to well-defined households with fixed utility functions. The reality is that households change constantly as people are born, die, marry, divorce and so on. Because these events could change a household's preferences, I entertain two possible definitions of a household. Under both definitions, I define anyone who moves out of a household as a new household. Under the first definition, I also create a new household when the head of household or the head's spouse changes, on the theory that the head and spouse make decisions and determine preferences. Under the second definition, I create a new household when any household members change.

Because new households can form from old ones, observations on different households may not be independent. In addition, the PSID uses a clustered sampling design in which many households originally lived in the same geographic areas. I adjust all bootstrap confidence intervals and two-step efficient GMM procedures to account for arbitrary correlation in the error terms across households and over time within each of the 119 original PSID primary sampling units.[25]

6.2 Threats to identification

Many previous researchers have identified ways in which either (1) income is correlated with consumption under full insurance or (2) income is uncorrelated with consumption even though insurance is incomplete or nonexistent. The first problem would render risk-sharing

[25] For bootstrap procedures, I draw primary sampling units from the original sample with replacement. I construct equal-tailed confidence intervals, which account for possible bias and asymmetry in the estimator's finite-sample distribution (Horowitz, 2001). For two-step efficient GMM, let $\hat{\boldsymbol{\theta}}$ be a GMM estimator of $\boldsymbol{\theta}$ based on the moment conditions $\mathrm{E}[\mathbf{g}(\mathbf{x}_i, \boldsymbol{\theta})] = \mathbf{0}$. Let j index groups of households. An estimated variance-covariance matrix for $\hat{\boldsymbol{\theta}}$ that accounts for correlation over time and across households within each group is $\hat{\mathbf{V}} = N(\mathbf{G}'\mathbf{W}\mathbf{G})^{-1}\mathbf{G}'\mathbf{W}\mathbf{S}\mathbf{W}\mathbf{G}(\mathbf{G}'\mathbf{W}\mathbf{G})^{-1}$, where $\mathbf{W}$ is the GMM weighting matrix, $\mathbf{G} = \sum_i \partial\mathbf{g}(\mathbf{x}_i, \hat{\boldsymbol{\theta}})/\partial\boldsymbol{\theta}$, $\bar{\mathbf{g}} = \sum_i \mathbf{g}(\mathbf{x}_i, \hat{\boldsymbol{\theta}})/N$ and $\mathbf{S} = \sum_j \left(\sum_{i\in j} \mathbf{g}(\mathbf{x}_i, \hat{\boldsymbol{\theta}}) - \bar{\mathbf{g}}\right)\left(\sum_{i\in j} \mathbf{g}(\mathbf{x}_i, \hat{\boldsymbol{\theta}}) - \bar{\mathbf{g}}\right)'/N$. $\mathbf{W}^* = \mathbf{S}^{-1}$ is an efficient weighting matrix that accounts for clustering.

regressions biased; the second would render them useless. My purpose is not to revisit the debate about whether regressions of consumption on income can ever constitute a good test of insurance, nor can I solve most of the well-known problems here. However, I review these problems as a reminder of important caveats in interpreting my results.

Observed income will be correlated with observed consumption if productivity at work is correlated with preference shocks, if consumption and leisure are nonseparable but leisure is omitted, if measurement errors in consumption and income are correlated, or if income includes insurance payments. None of these correlations necessarily reflects a failure of full insurance. My GMM estimates include leisure on the right-hand side of the consumption equation to account for nonseparable preferences. I try to remove insurance payments from income by subtracting welfare and other transfer payments.

Because I create a new household when the head or spouse changes, my regressions say nothing about insurance against divorce or death. In addition, a test of risk sharing using data on food consumption lacks power against the alternative that consumption of other goods is not well smoothed. More generally, whatever consumption variable a researcher studies may not be the consumption that households care about. Even given data on consumption of many goods besides food, one would face difficult questions of how to treat durable goods, household production and the like.

6.3 Factor estimates

Table 4 shows factor estimates of the coefficient on income for the common-preferences model (4) and for versions of the heterogeneous-preferences model (3) that allow heterogeneity in time preference, risk preference, or both. Leisure is assumed to be separable from consumption and is used as an instrument for income.

When I study households' total food consumption and assume households have identical preferences, the elasticity of food consumption with respect to income is 0.199 to 0.305, depending on the definition of a household. Thus a 1 percent increase in income raises food consumption by 0.199 to 0.305 percent, holding aggregate shocks constant but assuming identical preferences. Allowing variation in preferences reduces the elasticity by one-third to one-half, to between 0.100 and 0.150. Since food is a necessity and has an income elasticity below unity, the coefficients should be adjusted upward to find an estimated elasticity of total consumption to idiosyncratic income shocks. Blundell et al. (2005) estimate that the elasticity of total consumption to food consumption is $1/0.88 = 1.14$ in U.S. data.

In the total-consumption data, I reject at the 5 percent level the hypothesis that the coefficients are equal in models with and without heterogeneity. Because the heterogeneous-preferences estimator is consistent under weaker assumptions than the common-preferences estimator, I therefore reject the hypothesis that the common-preferences estimator is correctly specified.

The effect of accounting for heterogeneity is less clear after I adjust consumption to per-adult-equivalent units. If I define a new household whenever family composition changes in any way, allowing heterogeneity reduces the effect of income on consumption by a factor of between 20 percent and 40 percent; the difference is statistically significant. However, if I define a new household only when the head or spouse changes, the income coefficient falls only slightly when I control for heterogeneity, and the difference is not statistically significant.

6.4 GMM estimates

Tables 5 and 6 show two-step efficient Generalized Method of Moments estimates of the effect of income on consumption. Leisure is used as an instrument and, in some specifications, included on the right-hand side of the regression to account for nonseparability between consumption and leisure.

The overidentifying restrictions are frequently rejected in the models that assume common preferences, suggesting that these models are misspecified. The overidentifying restrictions are never rejected when heterogeneity in both risk and time preferences is allowed. Allowing heterogeneity in time preferences affects the coefficient on income more than allowing heterogeneity in risk preferences. Depending on the definition of a household, the adjustment for household size and whether I allow nonseparability between consumption and leisure, the elasticity of consumption with respect to income ranges from 0.195 to 0.545 when I assume common preferences. The elasticity ranges from -0.013 to 0.141 when I allow heterogeneity in time preferences, from 0.072 to 0.524 when I allow heterogeneity in risk preferences, and from -0.086 to 0.163 when I allow heterogeneity in both time and risk preferences. There is no clear pattern in how allowing nonseparability between consumption and leisure affects the results. Defining a new household if any household member changes tends to reduce the coefficient on income.

6.5 Discussion

The factor model estimates and GMM estimates both show that allowing heterogeneity tends to reduce the estimated effect of income on consumption, but the magnitudes of the estimated coefficients differ substantially between the factor and GMM specifications. The difference is not surprising, since the factor model and the GMM moment conditions rely on different assumptions. The factor model may detect more failures of insurance because it looks at variation in consumption and income over a longer time span.The factor estimates also require consumption and leisure to be separable, while the GMM specifications allow nonseparability if leisure is included as a regressor. However, the GMM estimates that allow nonseparability also require that leisure be measured without error.

The GMM estimates also potentially suffer from a weak instruments problem. The problem is especially severe in the case of estimates that allow heterogeneous risk preferences, because in these specifications the aggregate shocks are identified only from the correlation between preferences and income processes, a correlation that need not be strong. Finite sample bias due to weak instruments potentially explains the very large coefficients on income in some of the GMM specifications that allow heterogeneity only in risk preferences.

The effect of income on consumption differs statistically from zero in the factor model estimates and many GMM specifications, suggesting that insurance is imperfect. But the imperfection is small. In the factor model estimates, a 1 percent increase in a household's income, holding aggregate resources constant, raises the household's food consumption by at most around 0.15 percent, and perhaps much less. By comparison, Blundell et al. (2005) estimate that the cross-sectional elasticity of food spending to total nondurable spending is 0.88 in U.S. data. In other words, the response of consumption to income changes is less than one-fifth of what it would be in the absence of all consumption smoothing.

The partial-insurance model of section 5 provides another useful way to interpret the estimated income coefficient. Recall that, in the model, the elasticity of income to consumption is $\mathrm{E}[\phi_i/(\phi_i + \gamma_i)]$, the ratio of insurance costs to insurance costs plus risk aversion. Suppose for simplicity that $\phi_i/\gamma_i = k$ for all households i and that the average household has log utility. After adjusting the estimated 0.15 elasticity of food consumption to income upward based on the Blundell et al. (2005) results, my estimates imply that it would cost a household with log utility at most 0.1 percent of consumption to have a hypothetical 10 percent difference between consumption and income. By comparison, the cost would be up to 0.27 percent of consumption if one used the coefficients estimated assuming common preferences.

One can also compare my results to two papers that test risk sharing with PSID data and assume income processes are uncorrelated with risk preferences. Cochrane (1991) regresses consumption growth on income growth and finds an elasticity of 0.1 to 0.2. He does not use instrumental variables, so the coefficient is biased toward zero if income is measured with error; IV estimates would likely be larger. Dynarski and Gruber (1997) also regress consumption growth on income growth, but they use instrumental variables and allow variation in time preferences (although not risk preferences); they find an elasticity of 0.205. Both papers omit leisure from the right-hand side, thus assuming either that consumption and leisure are separable or that the social planner can freely transfer leisure across households.

7 Conclusion

Under full insurance, consumption does not depend at all on income after holding aggregate shocks constant. Under imperfect insurance, the effect of income on consumption shows the relative costs and benefits of risk sharing. This paper highlights the importance of accounting carefully for preferences when measuring how income affects consumption.

I show that if households have different risk and time preferences and if income processes are correlated with preferences, then a risk-sharing regression that assumes identical preferences will find too large an effect of idiosyncratic income shocks on consumption. Empirical results confirm the bias. First, income processes *are* correlated with preferences: In the HRS, earnings are more correlated with aggregate shocks among workers with greater risk tolerance. Second, in the PSID, allowing heterogeneity in preferences substantially reduces the estimated effect of income on consumption.

The results suggest several directions for further research. First, the HRS data suggest that people may choose jobs in part on the basis of risk preferences. Sorting on the basis of risk preference means not sorting purely on the basis of productivity; if risk were eliminated or insurance were better, people might sort differently and output might rise. (See Heathcote et al., 2008, for a related analysis of the effects of insurance on aggregate output when workers are identical but face idiosyncratic productivity shocks.) The welfare gain from changes in sorting could be calculated by using panel data on occupation choice, income and preferences to estimate the relationship between preferences, productivity in various occupations and the time series properties of individual income. Second, although I have demonstrated that –

in broad terms – more-risk-averse people bear less aggregate risk, I have not tested whether aggregate risk is allocated in precisely the shares required for Pareto efficiency. Detailed data on risk preferences and consumption processes might allow one to measure how well people share aggregate risk. Finally, I study a static model. The relationship between preferences, income processes and risk sharing could also fruitfully be investigated in a dynamic context.

Table 1: Summary statistics by risk tolerance for Health and Retirement Study earnings samples.

Variable	Low risk tol. mean	Low risk tol. s.d.	High risk tol. mean	High risk tol. s.d.
Variables that are constant for each worker				
Years of education	12.1	3.2	12.8	3.4
Age in 1992	56.1	5.1	55.3	5.2
White	0.78		0.78	
Immigrant	0.07		0.09	
Veteran	0.57		0.55	
Observations on worker	28.6	6.8	27.7	7.0
Uncensored observations on worker	20.4	7.4	19.5	7.1
Variables that change over time				
Experience (age-education-6)	21.0	10.3	19.9	10.2
Annual earnings[a]	19275	12136	19778	13119
Log(annual earnings)[a]	9.61	0.91	9.61	0.96
Number of men	2,528		1,562	
Number of observations	72,201		43,223	

Men in Health and Retirement Study with at least five uncensored earnings observations from age 23 to 61. [a]Social Security earnings through 1979; from 1980, W-2 earnings censored at $120,000 in nominal terms. Deflated by CPI; 1982-1984 dollars.

Table 2: Pooled random-intercept tobit regressions of log(real annual earnings) on aggregate variables, men ages 23-61 in Health and Retirement Study.

Aggregate variable:	GDP[a]		Personal consumption[a]		Wages/salaries[b]	
	Risk tolerance		Risk tolerance		Risk tolerance	
Parameter	low	high	low	high	low	high
			A. All men			
coefficient on	2.624	2.863	2.485	2.592	2.266	2.578
log(aggregate shock)	(0.054)	(0.077)	(0.066)	(0.087)	(0.049)	(0.070)
95% CI for difference	[0.047, 0.425]		[−0.110, 0.328]		[0.143, 0.485]	
standard deviation of	0.838	0.897	0.837	0.897	0.838	0.898
idiosyncratic error	(0.003)	(0.004)	(0.003)	(0.004)	(0.003)	(0.004)
95% CI for difference	[0.051, 0.069]		[0.050, 0.070]		[0.050, 0.069]	
observations	72,201	43,223	72,201	43,223	72,201	43,223
			B. White, U.S.-born men with exactly 12 years of education			
coefficient on	2.474	2.861	2.356	2.541	2.125	2.445
log(aggregate shock)	(0.101)	(0.149)	(0.134)	(0.195)	(0.091)	(0.139)
95% CI for difference	[0.039, 0.753]		[−0.265, 0.663]		[−0.019, 0.643]	
standard deviation of	0.780	0.801	0.780	0.801	0.780	0.802
idiosyncratic error	(0.005)	(0.007)	(0.005)	(0.007)	(0.005)	(0.007)
95% CI for difference	[0.004, 0.038]		[0.004, 0.038]		[0.005, 0.039]	
observations	21,837	9,911	21,837	9,911	21,837	9,911

Regressions also include individual and year random effects and control for experience, experience squared, time trend, time trend squared and an indicator for veteran status. In addition, regressions in panel A include years of education and indicators for immigrants and whites. See Internet appendix for estimated coefficients on controls. Table shows posterior means of parameters, with posterior standard deviations in parentheses. Figures in square brackets are 2.5th and 97.5th percentiles of the posterior distribution for the difference between low- and high-risk-tolerance groups. [a]Per capita, chained 2000 dollars. [b]Per capita, deflated by GDP deflator for personal consumption.

Table 3: Mean parameter estimates from individual tobit regressions of log(real annual earnings) on aggregate variables, men ages 23-61 in Health and Retirement Study.

Aggregate variable:	GDP[a]		Personal consumption[a]		Wages/salaries[b]	
	Risk tolerance		Risk tolerance		Risk tolerance	
Parameter	low	high	low	high	low	high
			A. All men			
coefficient on	1.651	2.034	1.607	1.867	1.320	1.707
log(aggregate shock)	(0.155)	(0.224)	(0.187)	(0.255)	(0.136)	(0.181)
test for difference	$t = 1.41$, $p = 0.08$		$t = 0.82$, $p = 0.21$		$t = 1.71$, $p = 0.04$	
standard deviation of	0.488	0.534	0.485	0.529	0.484	0.529
idiosyncratic error	(0.007)	(0.009)	(0.007)	(0.009)	(0.007)	(0.009)
test for difference	$t = 3.89$, $p = 0.00$		$t = 3.79$, $p = 0.00$		$t = 3.82$, $p = 0.00$	
workers	2,439	1,499	2,422	1,493	2,433	1,492
			B. White, U.S.-born men with exactly 12 years of education			
coefficient on	1.644	2.095	1.438	1.935	1.324	1.501
log(aggregate shock)	(0.277)	(0.414)	(0.323)	(0.479)	(0.230)	(0.345)
test for difference	$t = 0.91$, $p = 0.18$		$t = 0.86$, $p = 0.20$		$t = 0.43$, $p = 0.33$	
standard deviation of	0.434	0.470	0.429	0.468	0.429	0.470
idiosyncratic error	(0.013)	(0.019)	(0.012)	(0.019)	(0.013)	(0.019)
test for difference	$t = 1.61$, $p = 0.05$		$t = 1.71$, $p = 0.04$		$t = 1.80$, $p = 0.04$	
workers	715	330	714	328	714	328

Regressions also control for experience and experience squared. Workers were dropped if tobit estimation did not converge. Means weighted by number of uncensored observations on worker's earnings. Standard error of mean in parentheses. P-values are for test of null hypothesis that mean does not depend on risk tolerance, against alternative that mean is higher in high-risk-tolerance group. [a]Per capita, chained 2000 dollars. [b]Per capita, deflated by GDP deflator for personal consumption.

Table 4: Factor model estimates of the effect of income on consumption after controlling for aggregate shocks.

	log(consumption)							
	New household definition							
	New head/spouse				Any change			
	A. Total consumption data							
log(income)	0.305	0.150	0.147	0.148	0.199	0.132	0.106	0.100
95% CI	(0.261, 0.345)	(0.116, 0.187)	(0.115, 0.184)	(0.098, 0.193)	(0.160, 0.225)	(0.077, 0.187)	0.057, 0.144)	(0.021, 0.128)
95% CI for diff. from common prefs.	-	(-0.203, -0.129)	(-0.194, -0.129)	(-0.198, -0.120)	-	(-0.103, -0.019)	(-0.122, -0.063)	(-0.150, -0.056)
	B. Per adult equivalent consumption data							
log(income)	0.151	0.141	0.141	0.144	0.161	0.129	0.105	0.092
95% CI	(0.121, 0.190)	(0.108, 0.177)	(0.112, 0.179)	(0.107, 0.178)	(0.119, 0.193)	(0.073, 0.184)	(0.056, 0.144)	(0.005, 0.115)
95% CI for diff. from common prefs.	-	(-0.042, 0.022)	(-0.038, 0.026)	(-0.048, 0.031)	-	(-0.067, 0.010)	(-0.082, -0.028)	(-0.129, -0.033)
Heterogeneity:								
risk aversion	no	yes	no	yes	no	yes	no	yes
time preference	no	no	yes	yes	no	no	yes	yes

Equal-tailed 95% confidence intervals are computed using 79 bootstrap samples. To allow for correlation across households and over time within each of the 119 PSID primary sampling units, the bootstrap samples are constructed by drawing PSUs with replacement from the original sample.

Table 5: GMM estimates of the effect of income on consumption after controlling for aggregate shocks, if a new household is defined when the head or spouse changes.

	log(consumption)							
	A. Total consumption data							
log(income)	0.339	0.545	0.046	0.141	0.203	0.524	0.105	0.066
	(0.018)	(0.036)	(0.021)	(0.044)	(0.036)	(0.056)	(0.027)	(0.045)
log(leisure)	-	0.324	-	0.0427	-	0.393	-	-0.057
		(0.004)		(0.003)		(0.006)		(0.003)
Test of overidentifying restrictions:								
χ^2	56.9	37.2	42.0	53.3	37.5	50.8	39.6	35.8
d.f.	31	30	31	30	31	30	31	30
p	0.003	0.170	0.090	0.005	0.197	0.010	0.139	0.215
	B. Per adult equivalent consumption data							
log(income)	0.275	0.195	0.025	-0.001	0.397	0.470	0.097	0.163
	(0.019)	(0.037)	(0.023)	(0.046)	(0.053)	(0.063)	(0.027)	(0.050)
log(leisure)	-	-0.112	-	0.036	-	0.352	-	0.123
		(0.003)		(0.003)		(0.004)		(0.003)
Test of overidentifying restrictions:								
χ^2	41.0	43.8	31.9	29.3	46.4	46.5	38.0	33.4
d.f.	31	30	31	30	31	30	31	30
p	0.108	0.050	0.420	0.504	0.038	0.028	0.182	0.307
Heterogeneous risk preference	no	no	no	no	yes	yes	yes	yes
Heterogeneous time preference	no	no	yes	yes	no	no	yes	yes

Standard errors (in parentheses) and test statistics are robust to heteroscedasticity and to correlation across households and over time within each of the 119 PSID primary sampling units.

Table 6: GMM estimates of the effect of income on consumption after controlling for aggregate shocks, if a new household is defined when any family member changes.

	log(consumption)							
	A. Total consumption data							
log(income)	0.330	0.349	0.035	0.054	0.072	0.498	0.047	-0.070
	(0.023)	(0.037)	(0.025)	(0.053)	(0.032)	(0.083)	(0.031)	(0.066)
log(leisure)	-	0.053	-	0.125	-	0.648	-	-0.148
		(0.004)		(0.006)		(0.016)		(0.007)
Test of overidentifying restrictions:								
χ^2	59.7	48.8	23.4	42.3	50.3	40.2	31.8	32.1
d.f.	31	30	31	30	31	30	31	30
p	0.001	0.017	0.833	0.068	0.016	0.101	0.425	0.364
	B. Per adult equivalent consumption data							
log(income)	0.283	0.234	-0.001	-0.013	0.345	0.123	0.053	-0.086
	(0.023)	(0.041)	(0.026)	(0.055)	(0.047)	(0.076)	(0.029)	(0.068)
log(leisure)	-	-0.104	-	-0.079	-	0.083	-	-0.165
		(0.005)		(0.006)		(0.014)		(0.007)
Test of overidentifying restrictions:								
χ^2	26.1	37.5	27.4	35.4	35.6	29.0	37.6	30.3
d.f.	31	30	31	30	31	30	31	30
p	0.716	0.162	0.652	0.228	0.259	0.516	0.192	0.449
Heterogeneous risk preference	no	no	no	no	yes	yes	yes	yes
Heterogeneous time preference	no	no	yes	yes	no	no	yes	yes

Standard errors (in parentheses) and test statistics are robust to heteroscedasticity and to correlation across households and over time within each of the 119 PSID primary sampling units.

A PSID consumption, income and leisure data

I use the PSID core sample, which began with 3,000 households chosen randomly from the U.S. population in 1968. I examine these variables, for which summary statistics appear in table A.1:

Income: family money income – the sum of labor earnings, capital income and transfer payments received by all household members – minus Aid to Families with Dependent Children, Supplemental Security Income, other welfare payments, unemployment insurance, worker's compensation and help from relatives.

Leisure: 8,760 hours (the number of hours in a year) minus the number of hours that the head reported working.

Food consumption: the sum of annual expenditure on food eaten at home; annual expenditure on food eaten away from home, except at work and school; and annual value of food stamps received.

Adult equivalent household members: The PSID defines a food standard for each household that accounts for economies of scale as well as differences in food needs by age and sex. I divide income and consumption by the food standard to obtain data per adult equivalent household member. Table A.2 shows the equivalence scale.

Price indexes: I deflate the income data using the Consumer Price Index and the food consumption data using the food and beverages component of the CPI.

Dates: PSID questions on income and leisure refer to the previous calendar year. For example, questions in the 1968 survey asked about income and leisure during 1967. The time period covered by the food questions is not specified in the survey, and it is unclear what time period respondents have in mind when they respond. Following the literature, I assume that the food data also refer to the previous year.

Household structure: When the head or spouse changes, it is unclear when during the year the change took place. I therefore use only observations for which the household had the same head and spouse in the previous year as in the current year. I determine household membership by matching individual ID numbers in the PSID individual data file to households in the PSID family data file.

Sample selection: I drop an observation if the PSID flags any of the three food variables as a major or minor assignment (i.e., imputed value). This eliminates less than 1 percent of the observations.

Table A.1: Summary statistics for PSID consumption and income sample.

	New household definition			
	New head/spouse		Any change	
Variable	mean	s.d.	mean	s.d.
annual food consumption[a]	4026	2300	3840	2271
log(annual food consumption)[a]	8.14	0.59	8.09	0.61
adult equivalent food consumption[a,b]	2139	1183	2200	1206
log(adult equivalent food consumption)[a,b]	7.55	0.49	7.58	0.49
annual income net of transfers[c]	33542	32910	32053	33103
log(annual income net of transfers)[c]	10.09	0.98	10.03	0.98
adult equivalent annual income net of transfers[b,c]	17952	18700	18357	19695
log(adult equivalent annual income net of transfers)[b,c]	9.50	0.92	9.52	0.90
head's annual hours not at work	7036	1050	7150	1092
log(head's annual hours not at work)	8.85	0.15	8.86	0.15
Observations	60,820		42,740	
Households	5,677		5,489	
Years of data per household:				
mean	10.7		7.8	
minimum	4		4	
25th percentile	6		5	
median	8		7	
75th percentile	14		9	
maximum	22		22	

PSID core sample households with data on head's work hours, family money income and food consumption in at least four consecutive years. [a]Deflated by food and beverages component of CPI; 1982-1984 dollars. [b]See equivalence scale in table A.2; scaled so adjustment factor is 1 for a man age 21 to 35 living alone. [c]Deflated by CPI.

Table A.2: PSID adult equivalence scale.

Age	Male	Female	Family size	Adjustment
≤ 3	3.9	3.9	1	+20%
4-6	4.6	4.6	2	+10%
7-9	5.5	5.5	3	+5%
10-12	6.4	6.3	4	none
13-15	7.4	6.9	5	-5%
16-20	8.7	7.2	≥ 6	-10%
21-35	7.5	6.5		
35-55	6.9	6.3		
≥ 56	6.3	5.4		

Source: PSID; U.S. Department of Agriculture formula.

References

Ahn, Seung Chan, Young Hoon Lee and Peter Schmidt, 2001, "GMM estimation of linear panel data models with time-varying individual effects," *Journal of Econometrics* 101(2), 219–255.

Attanasio, Orazio and Steven J. Davis, 1996, "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy* 104(6), 1227–1262.

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball and Matthew D. Shapiro, 1997, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics* 112(2), 537–579.

Blundell, Richard, Luigi Pistaferri and Ian Preston, 2005, "Imputing Consumption in the PSID Using Food Demand Estimates from the CEX," Working paper, Institute for Fiscal Studies.

Blundell, Richard, Luigi Pistaferri and Ian Preston, 2008, "Consumption Inequality and Partial Insurance," *American Economic Review* 98(5), 1887–1921.

Bonin, Holger, Thomas Dohmen, Armin Falk, David Huffman and Uwe Sunde, 2007, "Cross-sectional earnings risk and occupational sorting: The role of risk attitudes," *Labour Economics* 14, 926–937.

Chamberlain, Gary and Michael Rothschild, 1983, "Arbitrage, Factor Struture, and Mean-Variance Analysis on Large Asset Markets," *Econometrica* 51(5), 1281–1304.

Chib, Siddhartha, 2001, "Markov Chain Monte Carlo Methods: Computation and Inference," in Heckman, James J. and Edward Leamer, editors, *Handbook of Econometrics*, volume V, pp. 3569–3649, Amsterdam: Elsevier.

Cochrane, John H., 1991, "A Simple Test of Consumption Insurance," *Journal of Political Economy* 99(5), 957–976.

Deaton, Angus, 1997, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, Baltimore: Johns Hopkins University Press for the World Bank.

Deaton, Angus and Christina Paxson, 1994, "Intertemporal Choice and Inequality," *Journal of Political Economy* 102(3), 437–467.

Diamond, Peter A., 1967, "The Role of a Stock Market in a General Equilibrium Model With Technological Uncertainty," *American Economic Review* 57(4), 759–776.

Dubois, Pierre, 2001, "Consumption Insurance with Heterogeneous Preferences. Can Sharecropping Help Complete Markets?" Manuscript, INRA Toulouse.

Dynarski, Susan and Jonathan Gruber, 1997, "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity* 1997(1), 229.

Flavin, Marjorie A., 1981, "The Adjustment of Consumption to Changing Expectations About Future Income," *Journal of Political Economy* 89(5), 974–1009.

Fuchs-Schündeln, Nicola and Matthias Schündeln, 2005, "Precautionary Savings and Self-Selection: Evidence From the German Reunification 'Experiment'," *Quarterly Journal of Economics* 120(3), 1085–1120.

Hayashi, Fumio, Joseph Altonji and Laurence Kotlikoff, 1996, "Risk-sharing Between and Within Families," *Econometrica* 64(2), 261–294.

Health and Retirement Study, 1991-1999, "HRS Core public use datasets; HRS, CODA and War Baby Social Security Covered Earnings and Wage and Self-Employment Income restricted data supplements," Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740). Ann Arbor, Mich.

Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante, 2009, "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," Manuscript, New York University.

Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante, 2008, "Insurance and opportunities: A welfare analysis of labor market risk," *Journal of Monetary Economics* 55, 501–525.

Heckman, James and Edward Vytlacil, 1998, "Instrumental Variables Methods for the Correlated Random Coefficient Model: Estimating the Average Rate of Return to Schooling When the Return is Correlated with Schooling," *Journal of Human Resources* 33(4), 974–987.

Horowitz, Joel L., 2001, "The Bootstrap," in Heckman, James J. and Edward Leamer, editors, *Handbook of Econometrics*, volume V, pp. 3159–3228, Amsterdam: Elsevier.

Kiefer, Nicholas M., 1980, "Estimation of Fixed Effect Models for Time Series of Cross-Sections With Arbitrary Intertemporal Covariance," *Journal of Econometrics* 14(2), 195–202.

Kurosaki, Takashi, 2001, "Consumption Smoothing and the Structure of Risk and Time Preferences: Theory and Evidence from Village India," *Hitotsubashi Journal of Economics* 42(2), 103–117.

Ligon, Ethan, Jonathan P. Thomas and Tim Worrall, 2002, "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *Review of Economic Studies* 69(1), 209–244.

Mace, Barbara J., 1991, "Full Insurance in the Presence of Aggregate Uncertainty," *Journal of Political Economy* 99(5), 928–956.

Mazzocco, Maurizio and Shiv Saini, 2009, "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," Manuscript, Duke University.

Munshi, Kaivan and Mark Rosenzweig, 2009, "Why is Mobility in India So Low? Social Insurance, Inequality, and Growth," manuscript, Brown University.

Nelson, Julie A., 1994, "On Testing for Full Insurance Using Consumer Expenditure Survey Data," *Journal of Political Economy* 102(2), 384–394.

Newey, Whitney K. and Richard J. Smith, 2004, "Higher Order Properties of GMM and Generalized Empirical Likelihood Estimators," *Econometrica* 72(1), 219–255.

Ogaki, Masao and Qiang Zhang, 2001, "Decreasing Relative Risk Aversion and Tests of Risk Sharing," *Econometrica* 69(2), 515–526.

Panel Study of Income Dynamics, 1974-1997, "Family and individual data files," Produced and distributed by the University of Michigan with primary funding from the National Science Foundation, the National Institute on Aging, and the National Institute of Child Health and Human Development. Ann Arbor, Mich.

Ravallion, Martin and Shubham Chaudhuri, 1997, "Risk and Insurance in Village India: Comment," *Econometrica* 65(1), 171–184.

Solon, Gary, Robert Barsky and Jonathan A. Parker, 1994, "Measuring the Cyclicality of Real Wages: How Important is Composition Bias," *Quarterly Journal of Economics* 109(1), 1–25.

Stock, James H., Jonathan H. Wright and Motohiro Yogo, 2002, "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," *Journal of Business and Economic Statistics* 20(4), 518–529.

Townsend, Robert M., 1994, "Risk and Insurance in Village India," *Econometrica* 62(3), 539–591.

Townsend, Robert M., 1995a, "Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies," *Journal of Economic Perspectives* 9(3), 83–102.

Townsend, Robert M., 1995b, "Financial Systems in Northern Thai Villages," *Quarterly Journal of Economics* 110(4), 1011–1046.

Udry, Christopher, 1994, "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria," *Review of Economic Studies* 61(3), 495–526.

Wilson, Robert, 1968, "The Theory of Syndicates," *Econometrica* 36(1), 119–132.

Wooldridge, Jeffrey M., 1997, "On two stage least squares estimation of the average treatment effect in a random coefficient model," *Economics Letters* 56(2), 129–133.

Wooldridge, Jeffrey M., 2003, "Further results on instrumental variables estimation of average treatment effects in the correlated random coefficient model," *Economics Letters* 79(2), 185–191.

Wooldridge, Jeffrey M., 2005, "Fixed-Effects and Related Estimators for Correlated Random-Coefficient and Treatment-Effect Panel Data Models," *Review of Economics and Statistics* 87(2), 385–390.