

Federal Reserve Bank of Minneapolis
Research Department Staff Report 462

September 2011

Heterogeneity and Tests of Risk Sharing*

Sam Schulhofer-Wohl

Federal Reserve Bank of Minneapolis

ABSTRACT

How well do people share risk? Standard risk-sharing regressions assume that any variation in households' risk preferences is uncorrelated with variation in the cyclicalities of income. I combine administrative and survey data to show that this assumption is questionable: 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

*Email: wohls@minneapolisfed.org. I thank Pierre-André Chiappori, James Heckman, Robert Shimer, and especially Robert Townsend for many invaluable discussions. Numerous colleagues, seminar participants, and the editor and referees of the *Journal of Political Economy* also made many 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. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

A large literature starting with Cochrane (1991), Mace (1991), and Townsend (1994) tests whether households are insured against idiosyncratic fluctuations in income by running regressions of the form

$$\Delta \log c_{it} = \Delta d_t + g \Delta \log X_{it} + \Delta \epsilon_{it}^{equal}, \quad (1)$$

where $\Delta \log c_{it}$ and $\Delta \log X_{it}$ are household consumption and income growth and Δd_t is a time indicator variable representing aggregate shocks. Under full risk sharing, consumption should not respond to idiosyncratic shocks after controlling for aggregate shocks, so g should be zero. In practice, however, g is almost always positive and statistically significant.¹

This paper shows that heterogeneity in risk preferences will bias the test in (1) against the null hypothesis of full insurance. When some households are less risk averse than others, a Pareto-efficient consumption allocation puts more aggregate consumption risk on those who are less risk averse. Therefore, under heterogeneity in risk aversion, (1) should be written as

$$\Delta \log c_{it} = b_i \Delta d_t + g \Delta \log X_{it} + \Delta \epsilon_{it}, \quad (2)$$

where b_i is decreasing in i 's risk aversion. Hence, the error term in (1) can be written as

$$\Delta \epsilon_{it}^{equal} = \Delta \epsilon_{it} + (b_i - 1) \Delta d_t. \quad (3)$$

Since Δd_t represents aggregate shocks, income growth $\Delta \log X_{it}$ will be positively correlated with the error term $\Delta \epsilon_{it}^{equal}$ in the standard regression (1) if less-risk-averse households' incomes move more strongly with aggregate shocks. Such a correlation will bias upward the estimated g in (1). In other words, if less-risk-averse households have more procyclical incomes, the standard regression will overestimate the correlation between idiosyncratic income and consumption and can spuriously reject the null of full insurance when insurance is actually

¹Full 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 1995). Mace (1991) does not reject efficiency in U.S. data, but Nelson (1994) overturns this result.

perfect. Section 2 formally derives this bias within a model of full insurance.

There are many reasons to expect that less-risk-averse households *will* have more procyclical incomes. If insurance is imperfect, workers will face a risk-return trade-off, and those who are more risk tolerant will choose jobs that carry more risk, both idiosyncratic and aggregate. Even under full insurance, risk preferences could be correlated with income risk — for example, if preferences for smoothing labor supply are correlated with preferences for smoothing consumption. Of course, other factors such as skills also affect job choices, so the relationship between risk preferences and income risk is ultimately an empirical question. In section 3, I provide suggestive evidence that less-risk-averse people have more procyclical incomes. I classify respondents to the Health and Retirement Study according to how they answer a hypothetical question about taking risky jobs. Earnings are more volatile and more correlated with aggregate consumption for men who report that they are 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 short panels and 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 then measures 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 these 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 heterogeneity 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 literatures on the relationship between preferences and income processes and on risk sharing with heterogeneous preferences. Moskowitz and Vissing-Jørgensen (2002) show that entrepreneurs' incomes are highly correlated with the public

equities market and hence highly procyclical; they also infer from portfolio choices that entrepreneurs may be more risk tolerant than nonentrepreneurs, but they do not have direct measures of preferences. Bonin et al. (2007) show that more-risk-averse people hold jobs with less idiosyncratic risk but do not analyze aggregate risk. Guiso et al. (2002) find that more-risk-averse people have riskier self-reported subjective earnings expectations. Fuchs-Schündeln and Schündeln (2005) measure precautionary savings while accounting for the endogeneity of income risk. 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, so hypotheses can be tested only by assuming a parametric distribution for the regression error. Kurosaki (2001) makes a similar analysis. Dubois (2001) tests for full insurance while allowing risk aversion to vary with observed household characteristics but rules out unobserved heterogeneity. Finally, in a closely related paper, Mazzocco and Saini (forthcoming) 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; estimate each household's preferences, rather than treat preferences as nuisance parameters; and 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. They also do not test whether income processes are correlated with independently measured preferences. 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. 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

²Mazzocco and Saini (forthcoming) suggest applying their methods to long panels of synthetic cohorts, but synthetic cohorts would average out all within-cohort heterogeneity, reducing the need to allow heterogeneity.

allocation and thus may generate spurious rejections of full insurance.

Assume that households' preferences depend only on a single consumption good, c . At each date t , denote the state of the economy by s_t . Household i 's preferences over consumption sequences are represented by discounted expected utility:

$$\mathbb{E}_0 \sum_{t=0}^T \beta^t \frac{[c_{it}(s_t)]^{1-\gamma_i}}{1-\gamma_i}, \quad (4)$$

where the coefficient of relative risk aversion $\gamma_i > 0$ varies across households.³

The optimal allocation can be decentralized, but it is convenient to study a social planner's problem. Let $C_t(s_t)$ be the aggregate consumption available in the economy at date t in state s_t . (If storage is possible, then C_t is net of aggregate storage.) Given Pareto weights α_i , the planner assigns consumption to households to maximize the weighted sum of their discounted expected utilities,

$$\max_{\{c_{it}(s_t)\}} \sum_i \alpha_i \mathbb{E}_0 \sum_{t=0}^T \beta^t \frac{[c_{it}(s_t)]^{1-\gamma_i}}{1-\gamma_i}, \quad (5)$$

subject to the constraint that, for each date and state, the total of households' consumptions is no larger than the aggregate available consumption:

$$\sum_i c_{it}(s_t) \leq C_t(s_t) \quad \forall t, s_t. \quad (6)$$

Assume an interior solution and let $\beta^t \Pr(s_t) \lambda_t(s_t)$ be the Lagrange multiplier on the aggregate resource constraint, where $\Pr(s_t)$ is the probability of state s_t . The first-order condition for consumption is

$$\alpha_i [c_{it}^*(s_t)]^{-\gamma_i} = \lambda_t(s_t). \quad (7)$$

The crucial implication of (7) is the familiar fact that, under full insurance, a household's consumption depends only on the aggregate shock $\lambda_t(s_t)$; conditional on the aggregate shock,

³Heterogeneity in γ_i can represent differences in utility functions or differences in relative risk aversion among households with identical non-CRRA preferences but different consumption levels. Schulhofer-Wohl (2007) shows that CRRA preferences can locally approximate any smooth, concave utility function.

idiosyncratic variables such as the household's income do not affect its consumption at all.

To bring the model to data, assume that consumption is measured with multiplicative error: Observed consumption is $c_{it} = e^{\epsilon_{it}} c_{it}^*$. Then the optimal allocation (7) satisfies

$$\log c_{it} = \frac{\log \alpha_i}{\gamma_i} + \frac{1}{\gamma_i}(-\log \lambda_t) + \epsilon_{it}. \quad (8)$$

Equation (8) says aggregate shocks λ_t have a larger effect on households that have smaller coefficients of relative risk aversion γ_i . Household income does not enter the equation at all, so one could test for full insurance by adding income X_{it} to the equation, as in

$$\log c_{it} = \frac{\log \alpha_i}{\gamma_i} + \frac{1}{\gamma_i}(-\log \lambda_t) + g \log X_{it} + \epsilon_{it}, \quad (9)$$

and 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 (9), however. The studies note that with identical risk preferences — $\gamma_i = \gamma$ — one can estimate

$$\log c_{it} = \frac{\log \alpha_i}{\gamma} + \frac{1}{\gamma}(-\log \lambda_t) + g \log X_{it} + \epsilon_{it}^{equal}, \quad (10)$$

which is simpler than (9) because $(1/\gamma)(-\log \lambda_t)$ is just a time dummy variable.⁴

The central point of this paper is that omitted variable bias makes the estimated coefficient on income in (10) too large. If the true model is (9) but a researcher mistakenly estimates (10), the error term in (10) is

$$\epsilon_{it}^{equal} = (1/\gamma_i - 1/\gamma)(-\log \lambda_t) + \epsilon_{it}. \quad (11)$$

The least squares estimator of the coefficient on income in equation (10) 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 $\text{Cov}(\log X_{it}, \epsilon_{it}^{equal}) < 0$. Suppose that income depends on common and idiosyncratic shocks: $\log X_{it} = q_i m_t + u_{it}$, where m_t is a common shock, q_i is the semielasticity of household i 's income to the common shock, and u_{it} is an idiosyncratic shock to i 's income. For simplicity,

⁴The household fixed effect $\log \alpha_i / \gamma_i$ in (10) is often removed by first-differencing the data to obtain (1).

assume the distributions of q_i and γ_i are stationary. Then, assuming u_{it} and ϵ_{it} are i.i.d.,

$$\text{Cov} \left[\log X_{it}, \epsilon_{it}^{equal} \right] = -\text{Cov} \left[q_i m_t, \left(\frac{1}{\gamma_i} - \frac{1}{\gamma} \right) \log \lambda_t \right] = -\text{Cov} [m_t, \log \lambda_t] \text{Cov} \left[q_i, \frac{1}{\gamma_i} \right]. \quad (12)$$

Aggregate income and aggregate consumption are likely to be positively correlated. Because λ_t is the multiplier on the aggregate resource constraint, it is decreasing in aggregate consumption. Therefore, $\text{Cov}[m_t, \log \lambda_t] < 0$. Hence (12) is positive, and the income coefficient in (10) is biased upward, if $\text{Cov}[q_i, 1/\gamma_i] > 0$. That is, the income coefficient in (10) is *biased upward if income responds more strongly to aggregate shocks for less-risk-averse households*.

A similar bias arises if households have heterogeneous time preferences. If β_i is the discount factor of household i , then the optimal allocation becomes

$$\log c_{it} = \frac{\log \alpha_i}{\gamma_i} + \frac{1}{\gamma_i} (-\log \lambda_t) + t \frac{\log \beta_i}{\gamma_i} + \epsilon_{it}. \quad (13)$$

That is, with heterogeneity in time preference, consumption rises faster for households with larger rates of time preference β_i or larger elasticities of intertemporal substitution $1/\gamma_i$. The heterogeneous trends $t \frac{\log \beta_i}{\gamma_i}$ must be included in the risk-sharing regression to obtain a valid test of full insurance. If they are omitted, as in (10), the estimated coefficient on income will be *biased upward if income rises more quickly for more patient households*.

Other than the Dubois (2001) and Mazzocco and Saini (forthcoming) 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, so relative risk aversion depends on the level of consumption but nothing else. My approach is different. In the next section, I present suggestive evidence that risk 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. My estimators also remove the potential bias from heterogeneity in time preferences.

3. Occupation Choice and Risk Preferences

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 — a crucial test because this hypothesis underpins my argument that standard risk-sharing tests are biased. On entering the study, each HRS respondent is asked the following:

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?⁵

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.⁶ 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 hypothetical jobs take more risks on many real-world dimensions but do not examine income processes.

The HRS includes restricted-access Social Security records of each respondent's annual earnings in jobs and self-employment from 1951 to 1991 or 1997, depending on when the respondent entered the sample. 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 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 omit self-employment earnings.⁷ I deflate earnings by the Consumer Price Index.

⁵To avoid status-quo bias, the question in some years specifies that the respondent must leave the current job and choose between two new jobs, one with a safe income and the other with a risky income.

⁶Some respondents are also asked about jobs that might cut income by 10 percent or 75 percent.

⁷The Web appendix shows that results using Social Security earnings for all years are largely similar.

I restrict my analysis to men ages 23 to 61, since women in the HRS cohorts and older and younger men are less likely to work full time. I drop respondents who did not answer risk-tolerance questions or did not release Social Security records, as well as observations with zero earnings. I assume observations with missing or zero earnings are missing at random.⁸ I also drop men who have fewer than five uncensored earnings observations, because some empirical models I consider are not always identified in such cases. With these restrictions, I have 115,424 annual observations for 4,090 men, of whom 62 percent fall in the low-risk-tolerance group. Appendix A1 gives summary statistics.

Because earnings are top-coded, I measure the relationship between earnings and aggregate shocks in a regression framework, where I can account for censoring with a tobit model. The specification I estimate is

$$\log(\text{earnings}_{it}) = \pi_{0i} + \pi_{1j} \log(PCE_t) + \mathbf{x}'_{it} \boldsymbol{\Pi}_j + v_{ijt}, \quad (14)$$

where PCE_t is real, per capita aggregate personal consumption expenditures from the National Income and Product Accounts; $j = \text{low, high}$ 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 include a quadratic time trend in the controls to account for common trends in aggregate consumption and earnings.

I interpret (14) as an equation that a worker could use to predict his lifetime earnings profile. The coefficient on aggregate consumption shows how income varies with aggregate shocks. The experience and trend coefficients 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 (14) is an admittedly simplistic specification for earnings dynamics, but with top-coded data, I cannot take first differences to compute innovations to income. Also, my goal is only to study the correlation between earnings and aggregate shocks, not to estimate a fully specified model of individual earnings. In the Web appendix, I use data

⁸Treating 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.

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.⁹ The Web appendix also shows that treating the individual-specific intercept as a fixed effect and allowing different coefficients on the regressors for each worker produces qualitatively similar results, though with larger standard errors.

The tobit model assumes the error term is normally distributed and has constant variance over time for each worker. I allow the error term to be correlated across workers within a risk-tolerance group in each year. I employ the Bayesian Markov Chain Monte Carlo method to estimate the posterior distribution of the parameters. The Web appendix gives details.

The top panel of table 1 shows the results. A 1 percent change in aggregate personal consumption raises earnings by 0.1 percentage point more for workers in the high-risk-tolerance group than for workers in the low-risk-tolerance group.¹⁰ It is unclear whether log aggregate personal consumption is the correct aggregate variable to include in this regression — in principle, we should use the multiplier on the aggregate resource constraint, but the multiplier is unobservable — so as a robustness check, I repeat the regression with other aggregate variables. The difference between the risk-tolerance groups is somewhat larger when I use GDP or aggregate wages and salaries in place of personal consumption. 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. For

⁹Although the PSID asked risk-tolerance questions in 1996, most people who answered the questions have too few years of data to precisely estimate the relationship between their earnings and aggregate shocks. The sample size can be increased — though not to nearly the size of the HRS sample — by examining family income and assuming each family has the risk tolerance of that family’s respondent in 1996. In addition, examining family income parallels my use of family income as the idiosyncratic shock in my risk-sharing tests later in the paper. The Web appendix shows that the elasticity of family income to aggregate personal consumption and GDP is higher for families with higher risk tolerance, but the difference is not statistically significant because the standard errors are quite large, and the difference disappears when aggregate wages and salaries are used to measure aggregate shocks.

¹⁰For both groups, the elasticity is surprisingly high, given that the elasticity of aggregate earnings to aggregate consumption has to be near one unless labor income is highly procyclical. Aggregation effects may explain this result: If earnings are more procyclical for low-income workers (Solon et al. 1994), the average elasticity of income to aggregate variables when we weight 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 incomes and reweighting to test this hypothesis. Reweighting is possible in the PSID data in the Web appendix and shows that the coefficient is statistically indistinguishable from one when workers are weighted by permanent income.

all three aggregate variables, the variance of idiosyncratic errors is also higher for workers in the high-risk-tolerance group. In other words, although the estimates are imprecise, the data provide suggestive evidence that 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, more education. To rule out the possibility that these differences in characteristics mean the more-risk-tolerant men never had the opportunity to take less risky jobs, in the bottom panel of table 1 I repeat the regressions on a sample restricted to white, native-born men with exactly 12 years of education (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 aggregate shocks for more risk-tolerant men.

The results support the idea that people share aggregate risk by sorting into jobs according to risk preferences, with several caveats. Most important, 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. Even then, however, the results 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. In addition, differences in reported risk tolerance may reflect something other than differences in preferences. For example, holding preferences 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 find that people who report lower risk tolerance hold safer jobs, without any 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 the measured differences in earnings risk. Other characteristics such as age and household composition or fixed obligations such as mortgages could also affect respondents' willingness to accept risky jobs. This effect could bias my results if the relevant characteristics are systematically related to the time-series correlation of an individual's income with aggregate shocks, as might happen if, say, married men are more likely than never-married men (with the same underlying preferences) to take low-risk

jobs. Otherwise, though, this effect would simply add noise to the measure of preferences and would, if anything, make it more difficult to find a difference between the risk-tolerance groups.

Measurement problems are also a concern. Earnings recorded by Social Security will spuriously move with GDP if workers shift during recessions from covered to uncovered jobs or if underreporting of self-employment income varies with the business cycle. Also, I do not correct for noise in the risk tolerance variable, but unless workers are more likely to be classified incorrectly than to be classified correctly — that is, if true risk tolerance is higher in the group that I classify as having low risk tolerance — misclassification implies that the true difference between the groups is larger than the estimated difference (Aigner 1973).

Finally, nonrandom sample selection could affect the results. Because HRS respondents are surveyed late in life, the results reflect the relationship between earnings and risk tolerance among people who survive to late middle age, rather than among all people. Also, because I exclude anyone with fewer than five uncensored earnings observations, 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 higher mean earnings, this nonrandom selection could cause me to find more variable earnings in the high-risk-tolerance group without any true sorting on the basis of risk. However, appendix table A.1 shows that the mean earnings of the two groups are virtually identical.

4. Econometrics: Robust Risk-Sharing Tests

In section 2, I showed that standard risk-sharing tests are biased if risk preferences are correlated with income processes, and in section 3, I showed that such a correlation indeed exists. Hence, in this section, I derive tests of risk sharing that explicitly account for the correlation of risk preferences with income processes. My tests have two goals. The first goal is to test the null hypothesis of full insurance while allowing risk preferences to be heterogeneous and correlated with income processes. We can do so by estimating equation (9) 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 the heterogeneous-preferences equation (9) and the common-preferences

equation (10) yield the same estimates of g . In particular, if we estimate (9) using a method that nests (10) and reject the null hypothesis that the two estimates of g are equal, then we reject the hypothesis that (10) is correctly specified.

Estimation of (10) is straightforward: The household-specific intercepts $\log \alpha_i/\gamma$ can be removed by a fixed-effects transformation or by first-differencing the data; the time effects $(1/\gamma) \log \lambda_t$ can be represented with time dummy variables; and the coefficient on income can then be estimated by ordinary least squares, or by instrumental variables if income is measured with error. Equation (9) requires more complicated methods, however, because the interaction between preferences γ_i and aggregate shocks $\log \lambda_t$ means that we cannot simply represent aggregate shocks with dummy variables.

Two methods are known for estimating equations like (9) in samples with many households i and a few dates t . The equation can be treated as a factor model (Kiefer 1980; Lawley and Maxwell 1963) and estimated by imposing restrictions on the variance-covariance matrix of the measurement errors ϵ_{it} . Alternatively, Ahn et al. (2001) propose Generalized Method of Moments estimators for a class of equations that includes (9). The two approaches are complementary. The factor approach imposes strong conditions on the distribution of the errors. 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 that the common-preferences regression is misspecified. By contrast, the GMM approach makes fewer assumptions but is valid only when preferences are heterogeneous. It is meaningless to compare a GMM estimate of (9) with an estimate of the common preferences model (10): The estimates are guaranteed to differ since each is valid only under assumptions that make the other invalid.

The key difference between the factor and GMM approaches lies in how the Lagrange multipliers $\log \lambda_t$ are estimated. (Once we know the Lagrange multipliers, we can simply include them with household-specific coefficients on the right-hand side of (9).) To illustrate how the two approaches estimate $\log \lambda_t$, assume temporarily that $\log \alpha_i/\gamma_i$ and g are known

and rewrite (9) as

$$u_{it} = (1/\gamma_i)(-\log \lambda_t) + \epsilon_{it}, \quad (15)$$

where $u_{it} = \log c_{it} - \log \alpha_i/\gamma_i - g \log X_{it}$. Let $\mathbf{u}_i = (u_{i1}, \dots, u_{iT})'$. The factor approach notes that if the measurement error ϵ_{it} is serially uncorrelated and homoskedastic over time for each household and uncorrelated with all other variables in the model, (15) implies

$$\mathbb{E}[\mathbf{u}_i \mathbf{u}_i'] = \mathbb{E}[1/\gamma_i^2] \mathbf{\Lambda} \mathbf{\Lambda}' + \mathbb{E}[\sigma_i^2] \mathbf{I}, \quad (16)$$

where $\mathbf{\Lambda} = -(\log \lambda_1, \dots, \log \lambda_T)'$, \mathbf{I} is the $T \times T$ identity matrix, σ_i^2 is the variance of ϵ_{it} for household i , and $\mathbb{E}[\cdot]$ stands for the mean across households. Equation (16) can be solved for $\mathbf{\Lambda}$ up to a sign and scale normalization.¹¹ Notice that we do not need to estimate each household's individual preferences γ_i in order to solve (16) for $\mathbf{\Lambda}$; preferences are nuisance parameters that have been eliminated from the equation by taking expectations. Although the factor method makes strong assumptions about the consumption measurement error ϵ_{it} , it assumes nothing about whether preferences γ_i are heterogeneous. Thus, under the maintained assumptions on ϵ_{it} , the common preferences model is nested in the factor model; if factor estimates of the heterogeneous-preferences model differ from estimates of the common-preferences model, then the common-preferences model must be misspecified.

The GMM approach instead quasi-differences (15) — along the lines of the quasi-differencing of panel vector autoregressions in Holtz-Eakin et al. (1988) — to obtain

$$u_{it} = \frac{\log \lambda_t}{\log \lambda_{t-1}} u_{i,t-1} + \epsilon_{it} - \frac{\log \lambda_t}{\log \lambda_{t-1}} \epsilon_{i,t-1}. \quad (17)$$

Again, preferences are nuisance parameters that we have eliminated from the equation. We can estimate $\log \lambda_t / \log \lambda_{t-1}$ from the cross-sectional regression of u_{it} on $u_{i,t-1}$; again, we can identify $\log \lambda_t$ up to a sign and scale normalization. Unlike in the factor estimator, we need not assume ϵ_{it} is serially uncorrelated and homoskedastic. But now, because $u_{i,t-1}$ is correlated with $\epsilon_{i,t-1}$, we must use an instrument for $u_{i,t-1}$ and estimate $\log \lambda_t$ from moment

¹¹The normalization is needed because multiplying all of the γ_i 's by a constant and dividing all of the $\log \lambda_t$'s by the same constant does not change the equation. The economic reason why we need a normalization is that doubling the relative risk aversion of all households does not change the efficient allocation.

conditions of the form

$$\mathbb{E} \left[z_{it} \left(u_{it} - \frac{\log \lambda_t}{\log \lambda_{t-1}} u_{i,t-1} \right) \right] = 0, \quad (18)$$

where z_{it} is some instrumental variable orthogonal to ϵ_{it} and $\epsilon_{i,t-1}$. A valid instrument exists only if preferences are heterogeneous. To see why, recall that $u_{i,t-1} = (1/\gamma_i)(-\log \lambda_{t-1}) + \epsilon_{i,t-1}$. A valid instrument must be correlated with $u_{i,t-1}$ but uncorrelated with $\epsilon_{i,t-1}$. Since $\log \lambda_{t-1}$ does not vary in the cross section, an instrument can satisfy both requirements only if it is cross-sectionally correlated with γ_i . If preferences are identical, no variable will be correlated with them and no valid instrument will exist. Therefore, the GMM estimator does not nest the common-preferences model; we cannot compare GMM and common-preferences estimates to determine whether the common-preferences model is correctly specified.

The preceding discussion abstracts from how the household-specific intercepts $\log \alpha_i/\gamma_i$ and income coefficient g are estimated. In the GMM approach, the household-specific intercepts are removed by first-differencing, and the income coefficient is estimated from the same moment conditions that are used to estimate the aggregate shocks based on (17). (Since $u_{it} = \log c_{it} - \log \alpha_i/\gamma_i - g \log X_{it}$, the parameter g appears implicitly in (17).) In the factor approach, the household-specific intercepts are removed by a fixed-effects transformation — that is, by expressing income and consumption as deviations from the household’s mean income and consumption over time. An iterative method is then used to estimate both g and $\log \lambda_t$: First, a value is guessed for g , and the aggregate shocks $\log \lambda_t$ are estimated from (16) based on this guess. Second, the estimated aggregate shocks are used to remove preferences γ_i from (9) by a quasi-fixed-effects transformation that is equivalent to controlling for the aggregate shocks with household-specific coefficients, and a new value for g is then estimated by regressing consumption on income (using instrumental variables to account for measurement error in income). These two steps are iterated until the estimates for g and $\log \lambda_t$ converge. The factor approach must use fixed effects rather than first-differencing to remove the household-specific intercepts because first-differencing would violate the requirement that the errors be serially uncorrelated.

Because the GMM method differences the data from one year to the next while the factor method subtracts from each variable its mean over time, the two methods perform

different experiments. The GMM test analyzes how consumption responds to changes in income from one year to the next. By contrast, the factor test analyzes how consumption responds to deviations of income from its mean over time. Hayashi et al. (1996) note that if insurance is imperfect but households know about income changes in advance and adjust consumption accordingly, consumption growth is uncorrelated with contemporaneous income growth, even though consumption will track income over longer time horizons. A test that compares consumption at date t with its mean over time includes consumption at some long-ago dates in the comparison and thus has a better chance of detecting the long-run changes.¹² Hence, since the factor method measures the relationship between consumption and income over longer periods than the GMM method, the factor method 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. Therefore, in both the factor and

¹²As a simple example, suppose that a household is observed for four years, earning \$100,000 in year 1, \$110,000 in years 2 and 3, and \$120,000 in year 4. Suppose also that the household chooses to consume \$100,000 in year 1, \$110,000 in year 2, and \$115,000 in each of years 3 and 4, perhaps because it has advance knowledge of what it will earn in year 4. The first difference of income is uncorrelated with the first difference of consumption in these data, but the deviation of consumption from its mean is positively correlated with the deviation of income from its mean.

GMM methods, I use leisure as an instrument for income. Leisure is likely to be correlated with income — people who work more hours will have higher income and less leisure, all else equal — but, because leisure is a separate survey item, measurement error in leisure is plausibly uncorrelated with measurement error in income.

One potential problem with leisure as an instrument is that if preferences are non-separable between consumption and leisure, then leisure may affect the marginal utility of consumption. Cochrane (1991) shows that if the planner can freely transfer leisure across households, then nonseparability does not change the optimal allocation and thus does not affect tests of full insurance. However, if leisure is given exogenously and cannot be transferred across households, the efficient allocation depends on leisure, implying that leisure should be controlled for on the right-hand side of (9) and (10). In the GMM tests, I can control for leisure while also using it as an instrument because leisure and lagged leisure can serve as separate instruments, but only a particular linear combination of leisure and its lag — the result of the quasi-differencing in (17) — appears on the right-hand side; thus, I have more instruments than regressors even after controlling for leisure. However, if I both include leisure as a regressor and use it as an instrument, I am implicitly assuming that leisure is not mismeasured, because a mismeasured variable cannot serve as an instrument for itself. In the factor method, I cannot control for leisure at the same time as I use it as an instrument: I found that the iterative procedure was numerically unstable and did not converge when I attempted to use lagged leisure as an instrument, so I restrict myself to using contemporaneous leisure as the instrument, but this leaves me with no extra instruments and prevents me from controlling for leisure.

Appendix A2 gives formulas for the GMM moment conditions and the quasi-fixed-effects transformations for the factor method and states precisely the econometric assumptions each method uses. The appendix also shows how the tests can be extended to allow heterogeneity in time preference by including household-specific trends in (9).

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. In this section, I develop a simple model that shows how coefficients in risk-sharing regressions measure the extent of partial insurance.

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 . There is no household-level storage. If aggregate storage (by a social planner) is feasible, then, as with full insurance, the model should be interpreted as describing the optimal allocation conditional on aggregate storage.

I assume the transfer cost function is

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

This function is convex in the percentage difference between consumption and income. Convex costs will arise if several sources of insurance are available and have different costs; the least costly forms of insurance will be used first, and the marginal cost of a transfer will rise with the magnitude of the transfer. For example, if a consumer holds a buffer stock invested in assets of varying liquidity, it will be less costly to make a small withdrawal from the most liquid assets than to make a large withdrawal. The functional form of h also allows me to express the costs of transfers as a percentage of consumption, and it results in an optimal consumption allocation that takes the same form as the risk-sharing regression (9), so I can use the partial-insurance model to interpret estimates of the risk-sharing regression.

My 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 because it would set aside all other institutions.

My 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 (9) or (10). 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 are attractive for many purposes, they imply that household consumption depends on both contemporaneous and lagged variables, so they cannot be used directly to interpret estimates of an equation such as the frequently estimated risk-sharing regression in which consumption depends only on contemporaneous variables.

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 (5) 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:

$$\sum_i X_{it}(s_t) \geq \sum_i c_{it}(s_t) + \sum_i \phi_i h[X_{it}(s_t), c_{it}(s_t)] \quad \forall t, s_t. \quad (19)$$

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

$$\alpha_i [c_{it}^*(s_t)]^{-\gamma_i} = \lambda_t(s_t) \left(1 + \phi_i \frac{\partial}{\partial c} h[X_{it}(s_t), c_{it}^*(s_t)] \right). \quad (20)$$

Observed consumption thus satisfies

$$\log c_{it} = \frac{\log \alpha_i}{\gamma_i} - \frac{\log \lambda_t}{\gamma_i} - \frac{1}{\gamma_i} \log \left[1 + \phi_i \log \frac{c_{it}}{X_{it}} + \frac{\phi_i}{2} \left(\log \frac{c_{it}}{X_{it}} \right)^2 \right] + \epsilon_{it}. \quad (21)$$

For $\phi_i \log(c_{it}/X_{it})$ close to zero — that is, when the resources lost to imperfect risk sharing

are small — equation (21) is approximately equivalent to

$$\log c_{it} = \frac{\log \alpha_i}{\phi_i + \gamma_i} + \frac{1}{\phi_i + \gamma_i}(-\log \lambda_t) + \frac{\phi_i}{\phi_i + \gamma_i} \log X_{it} + \epsilon_{it}. \quad (22)$$

Equation (22) says income has a larger effect on the consumption of households that are less risk averse (small γ_i) or more difficult to insure (large ϕ_i).¹³ 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: $g \equiv E[\phi_i/(\phi_i + \gamma_i)]$. Then we can rewrite (22) as

$$\log c_{it} = \frac{\log \alpha_i}{\phi_i + \gamma_i} + \frac{1}{\phi_i + \gamma_i}(-\log \lambda_t) + g \log X_{it} + \tilde{\epsilon}_{it},$$

$$\tilde{\epsilon}_{it} = \epsilon_{it} + \left(\frac{\phi_i}{\phi_i + \gamma_i} - g \right) \log X_{it}, \quad (23)$$

which is the same risk-sharing regression proposed in (9), but with a different error term. The estimated income coefficient in risk-sharing regressions such as (9) and (10) thus measures the average relative costs and benefits of risk sharing. 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 (10) instead of the heterogeneous-preferences equation (9). Since the bias in the estimate of g is likely upward, estimates of (10) may lead us to conclude that insurance is worse than it actually is.

Section 4 derived methods for testing the null hypothesis of full insurance based on testing whether $g = 0$ in (9). Because the methods generate estimates of g , I can also use them to measure the extent of partial insurance if I reject full insurance. The only difference between the forms of (9) and (23) is that the error term in (23) potentially contains a random coefficient on income. For simplicity, I assume here that the coefficient on income is common across households — that is, I assume $\phi_i/(\phi_i + \gamma_i) = g$ for all i — so that there is no random

¹³One 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 nonnegativity constraint on measured consumption makes additive errors statistically unattractive.

coefficient. With this assumption, (23) becomes

$$\log c_{it} = \frac{(1-g)\log \alpha_i}{\gamma_i} + \frac{1-g}{\gamma_i}(-\log \lambda_t) + g \log X_{it} + \epsilon_{it}, \quad (24)$$

which is identical to (9) except that the intercept and aggregate shock terms have been multiplied by $1-g$. Rescaling these terms does not change the assumptions required to estimate the equation; thus, given $\phi_i/(\phi_i + \gamma_i) = g$, the methods I use for (9) remain valid for (23). In Schulhofer-Wohl (2007), I show that the estimators also remain valid if the covariance between the instrument (leisure) and income does not depend on $\phi_i/(\phi_i + \gamma_i)$.

Estimates of the partial insurance model require one important caveat. The differencing operators used to remove household fixed effects could amplify any biases from the linear approximation that turns (21) into (22). 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.

6. Robust Risk-Sharing Regressions: Results

This section estimates the effects of idiosyncratic income fluctuations on consumption in equations (9) and (10). I analyze data on consumption, income, and leisure from the Panel Study of Income Dynamics. The PSID 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.¹⁴ 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

¹⁴The 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. I choose not to follow Blundell et al. (2005) in imputing total consumption in the PSID by inverting a food demand equation because that approach would assume that heterogeneity in preferences depends only on observables, contrary to the premise of this paper that unobservable heterogeneity is important.

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 measure leisure by 8,760 (the number of hours in a year) minus hours worked by the head of household.¹⁵ I restrict the sample to households with consumption, income, and leisure data in at least four consecutive years, the minimum number of observations required for the GMM estimators. Appendix A1 describes the data in detail.

Theory describes the optimal allocation of consumption to well-defined households with fixed utility functions. In reality, households change constantly as people are born, die, marry, divorce, and so on. Because these events could change preferences, I create a new household when any household members change. I also put consumption and income in per capita terms by dividing by the number of adult equivalent household members, though this has little effect since the list of household members is constant; the only change over time is in their ages. The Web appendix shows that the differences between tests with and without preference heterogeneity are similar if I use total consumption and income data instead of per capita data, or if I create a new household only when the head or spouse changes.

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. 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.¹⁶

¹⁵The Web appendix shows that the parameter estimates do not change substantially when I define leisure to include the spouse’s leisure for couple-headed households.

¹⁶For 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, consider estimating a parameter vector θ based on the moment conditions $E[\mathbf{g}(\mathbf{x}_i, \theta)] = \mathbf{0}$. Let j index groups of households. An efficient weighting matrix that accounts for clustering is $\mathbf{W} = [\sum_j (\sum_{i \in j} \mathbf{g}(\mathbf{x}_i, \hat{\theta}) - \bar{\mathbf{g}})(\sum_{i \in j} \mathbf{g}(\mathbf{x}_i, \hat{\theta}) - \bar{\mathbf{g}})' / N]^{-1}$, where $\hat{\theta}$ is a first-step consistent estimator and $\bar{\mathbf{g}} = \sum_i \mathbf{g}(\mathbf{x}_i, \hat{\theta}) / N$.

A. 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 biases the income coefficient in a risk-sharing regression away from zero; the second biases the coefficient toward zero. I cannot solve most of the well-known problems here, but I review them as a reminder of important caveats to 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.

Many factors can cause observed income and consumption to be uncorrelated even when insurance is imperfect. For example, because I create a new household when any household member 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 insured. 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, home production, and the like. Classical measurement error in income will also bias the estimated income coefficient toward zero; as I discuss in section 4, I use instrumental variables to remove this attenuation bias.

B. Factor estimates

Table 2 shows factor estimates of the coefficient on income for the common-preferences model (10) and for versions of the heterogeneous-preferences model (9) 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.

In column 1 of the table, when I assume households have identical preferences, a 1 percent increase in income raises food consumption by 0.161 percent, holding aggregate shocks constant. In the remaining columns, allowing variation in preferences reduces the elasticity by a factor of 20 to 40 percent, to between 0.092 and 0.129. 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 models that allow heterogeneity in time preference or in both risk and time preference — columns 3 and 4 — 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. But even with the heterogeneous-preferences estimates, I reject the null hypothesis of full insurance.

C. GMM estimates

Table 3 shows 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 never rejected.

The even-numbered columns of the table allow nonseparability between consumption and leisure, while the odd-numbered columns assume separability. The estimates show that allowing nonseparability always reduces the coefficient on income.¹⁷ Since allowing nonseparability is a key advantage of the GMM approach, I focus on the estimates that allow nonseparability. These estimates show that the elasticity of consumption with respect to income is 0.234 under common preferences (column 2) but falls to between -0.086 and 0.123 when I allow preference heterogeneity (columns 4, 6, and 8). The coefficient on income is never statistically significantly different from zero when I allow for heterogeneity and nonseparability. The GMM results thus show that, when I allow for heterogeneity and nonseparability, the

¹⁷However, the Web appendix shows that under alternative definitions of consumption and household membership, there is no clear pattern in the effect of allowing nonseparability.

data do not reject the hypothesis of full insurance.

D. 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, suggesting that insurance may be imperfect. However, the GMM estimates are not significantly different from zero, and the coefficients are small; the imperfection, if any, is not large. 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 0.092 percent once we allow heterogeneity in risk and time preferences. 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-ninth of what it would be in the absence of all consumption smoothing.

The partial-insurance model of section 5 provides another way to interpret the estimated income coefficient. Recall that, in the model, the elasticity of income to consumption is $E[\phi_i/(\phi_i + \gamma_i)]$, 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 0.092 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 about 0.10 percent of consumption to have a hypothetical 10 percent difference between consumption and income. The cost would be nearly twice as high — 0.18 percent of consumption — if one used the coefficient 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. If households have different risk preferences and if income processes are correlated with preferences, 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 substantially reduces the estimated effect of income on consumption.

The results suggest several directions for further research. 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.¹⁸ The welfare gain from changes in sorting could be calculated by estimating the relationship between preferences, productivity in various occupations, and the time-series properties of individual income. Also, although I have demonstrated that 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. Finally, I study a static model of imperfect risk sharing. The relationship between preferences, income processes, and incomplete insurance could also fruitfully be investigated in a dynamic context.

¹⁸See 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.

Table 1: Pooled random-intercept tobit regressions of log(real annual earnings) on aggregate variables, men ages 23-61 in Health and Retirement Study.

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

High-risk-tolerance individuals are those who prefer a job with a fifty-fifty chance of doubling income or of cutting it by 20 percent to a job that pays the respondent's current income for certain. 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 Web 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. ^aPer capita, chained 2000 dollars. ^bPer capita, deflated by GDP deflator for personal consumption.

Table 2: Factor model estimates of the effect of income on consumption after controlling for aggregate shocks.

	log(consumption per adult equivalent)			
	(1)	(2)	(3)	(4)
log(income)	0.161	0.129	0.105	0.092
95% CI	(0.119, 0.193)	(0.073, 0.184)	(0.056, 0.144)	(0.005, 0.115)
95% CI for diff. from common prefs.	-	(-0.067, 0.010)	(-0.082, -0.028)	(-0.129, -0.033)
Heterogeneity:				
risk aversion	no	yes	no	yes
time preference	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 3: GMM estimates of the effect of income on consumption after controlling for aggregate shocks.

	log(consumption per adult equivalent)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log(income)	0.283 (0.023)	0.234 (0.041)	-0.001 (0.026)	-0.013 (0.055)	0.345 (0.047)	0.123 (0.076)	0.053 (0.029)	-0.086 (0.068)
log(leisure)	-	-0.104 (0.005)	-	-0.079 (0.006)	-	0.083 (0.014)	-	-0.165 (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
Heterogeneity:								
risk aversion	no	no	no	no	yes	yes	yes	yes
time preference	no	no	yes	yes	no	no	yes	yes

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

Appendix

A1. Summary statistics for the HRS and PSID samples

Table A.1 summarizes the earnings data from the HRS.

Table A.2 summarizes the consumption, income, and leisure data from the PSID. 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:

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.3 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 by risk tolerance for Health and Retirement Study earnings samples.

Variable	Low risk tol.		High risk tol.	
	mean	s.d.	mean	s.d.
<i>Variables that are constant for each worker</i>				
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
<i>Variables that change over time</i>				
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. ^aSocial Security earnings through 1979; from 1980, W-2 earnings censored at \$120,000 in nominal terms. Deflated by CPI; 1982–1984 dollars.

Table A.2: Summary statistics for PSID consumption and income sample.

Variable	mean	s.d.
annual food consumption ^a	3840	2271
adult equivalent food consumption ^{a,b}	2200	1206
log(adult equivalent food consumption) ^{a,b}	7.58	0.49
annual income net of transfers ^c	32053	33103
adult equivalent annual income net of transfers ^{b,c}	18357	19695
log(adult equivalent annual income net of transfers) ^{b,c}	9.52	0.90
head's annual hours not at work	7150	1092
log(head's annual hours not at work)	8.86	0.15
Observations	42,740	
Households	5,489	
Years of data per household:		
mean	7.8	
minimum	4	
25th percentile	5	
median	7	
75th percentile	9	
maximum	22	

PSID core sample households with data on head's work hours, family money income, and food consumption in at least four consecutive years. ^aDeflated by food and beverages component of CPI; 1982–1984 dollars. ^bSee equivalence scale in table A.3; scaled so adjustment factor is 1 for a man age 21 to 35 living alone. ^cDeflated by CPI.

Table A.3: 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.

A2. Details of econometric methods

This section gives details of my methods for estimating (9). To keep the notation simple, I consider here a balanced panel of data on N households in T years. Schulhofer-Wohl (2007) shows that the methods remain valid when the panel is unbalanced, as long as missing data are missing at random. Let $y_{it} = \log c_{it}$, $x_{it} = \log X_{it}$, $z_{it} = \log(\textit{leisure}_{it})$, and $d_t = -\log \lambda_t$. I consider the case where both risk and time preferences are heterogeneous, so the equation to be estimated includes household-specific trends:

$$y_{it} = \frac{\log \alpha_i}{\gamma_i} + \frac{1}{\gamma_i} d_t + t \frac{\log \beta_i}{\gamma_i} + g x_{it} + \epsilon_{it}. \quad (\text{A2.1})$$

For any variable ξ_{it} , let $\boldsymbol{\xi}_i = [\xi_{i1}, \dots, \xi_{iT}]'$ denote the column vector of observations on ξ_{it} for household i . Let $\mathbf{1}$ be a column vector of ones. I assume that $T \geq 4$, that each household is observed at least three times, and that some households are observed four times.¹⁹

A2.1. Factor model test

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

Assumption F.1. $E[x_{is}\epsilon_{it}] = E[(1/\gamma_i)\epsilon_{it}] = E[(\log \beta_i/\gamma_i)\epsilon_{it}] = E[\epsilon_{it}] = 0$ for all i, s, t .

Assumption F.2. *Instead of observing actual log income, x_{it} , we observe $x_{it}^* = x_{it} + v_{it}$.*

Assumption F.3. $E[\mathbf{z}_i\epsilon_i'] = E[\mathbf{z}_i\mathbf{v}_i'] = \mathbf{0}$.

Assumption F.1 says measurement error in consumption is uncorrelated with income and preferences. Assumption F.2 says log income is measured with additive error. Assumption F.3 says leisure is uncorrelated with measurement error in consumption and income.

I also assume measurement error in both consumption and income is serially uncorrelated and homoskedastic within each household:²⁰

Assumption F.4. $E[\epsilon_i\epsilon_i'] = \sigma_i^2\mathbf{I}$ and $E[\mathbf{v}_i\mathbf{v}_i'] = \zeta_i^2\mathbf{I}$ for all i .

¹⁹Given only three time periods, we could fit the data perfectly with any g , $\lambda_0 > 0$, $\lambda_1 > 0$, and $\lambda_2 > 0$ by choosing α_i , γ_i , and β_i to solve (A2.1) for $t = 0, 1, 2$.

²⁰We cannot test this assumption in an unbalanced panel unless we assume the pattern of missing data is uncorrelated with the household-specific variance of measurement error.

Let \mathbf{I} be the $T \times T$ identity matrix and $\mathbf{M} = \mathbf{I} - \frac{\mathbf{u}\mathbf{u}'}{T}$. If we multiply any vector by \mathbf{M} , we obtain the residuals from regressing that vector on a constant — in other words, we use fixed effects to remove household-specific means. Similarly, for any $T \times 1$ vector $\tilde{\mathbf{d}}$, define

$$\mathbf{V}(\tilde{\mathbf{d}}) = \mathbf{I} - [\mathbf{t} \tilde{\mathbf{d}}][(\mathbf{t} \tilde{\mathbf{d}})'[\mathbf{t} \tilde{\mathbf{d}}]]^{-1}[\mathbf{t} \tilde{\mathbf{d}}]'. \quad (\text{A2.2})$$

If we multiply any vector by $\mathbf{V}(\tilde{\mathbf{d}})$, we obtain the residuals from regressing that vector on a trend and \tilde{d}_t . That is, multiplying by $\mathbf{V}(\tilde{\mathbf{d}})$ uses a quasi-fixed-effects transformation to remove household-specific trends and, if $\tilde{\mathbf{d}} = \mathbf{d}$, household-specific effects of aggregate shocks.

I require a version of the usual rank condition for instrumental variables:

Assumption F.5. $E[\mathbf{z}'_i \mathbf{M}\mathbf{V}(\tilde{\mathbf{d}})\mathbf{M}\mathbf{x}_i] \neq 0$ for all $\tilde{\mathbf{d}}$ in a neighborhood of \mathbf{d} .

Assumption F.5 says that the instrument is correlated with income after controlling for aggregate shocks, not only for the true aggregate shocks but also for incorrect aggregate shocks in the neighborhood of the true shocks.

A useful result (Lawley and Maxwell 1963, chap. 2) is that the solution to (16) is the same as the solution to the least squares problem

$$\min_{\{\gamma_i\}_{i=1}^N, \{\lambda_t\}_{t=1}^T} \sum_{i,t} [u_{it} - (1/\gamma_i)(-\log \lambda_t)]^2. \quad (\text{A2.3})$$

My estimator of (\mathbf{d}, g) is the pair $(\tilde{\mathbf{d}}, \tilde{g})$ that simultaneously solves the following equations:

$$\begin{aligned} \tilde{\mathbf{d}}(\tilde{g}) &= \arg \min_{\hat{\mathbf{d}}} \frac{1}{N} \sum_{i=1}^N (\mathbf{y}_i - \tilde{g}\mathbf{x}_i^*)' \mathbf{M}\mathbf{V}(\hat{\mathbf{d}})\mathbf{M}(\mathbf{y}_i - \tilde{g}\mathbf{x}_i^*) \\ &\text{subject to } \sum_{t=1}^T d_t = \sum_{t=1}^T t d_t = 0, \quad \sum_{t=1}^T d_t^2 = 1, \quad d_1 > 0. \end{aligned} \quad (\text{A2.4})$$

$$\tilde{g}(\tilde{\mathbf{d}}) = \frac{\sum_{i=1}^N \mathbf{z}'_i \mathbf{M}\mathbf{V}(\tilde{\mathbf{d}})\mathbf{M}\mathbf{y}_i}{\sum_{i=1}^N \mathbf{z}'_i \mathbf{M}\mathbf{V}(\tilde{\mathbf{d}})\mathbf{M}\mathbf{x}_i^*} \quad (\text{A2.5})$$

Equation (A2.4) uses the least squares procedure (A2.3) 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 (A2.3) will continue to provide a consistent estimator of \mathbf{d} under assumption F.4. Hence, $\tilde{\mathbf{d}}(\tilde{g})$ will be consistent for \mathbf{d} if \tilde{g} is a consistent estimator of g . The normalization in (A2.4) is required because, for any nonzero constants k_1 , k_2 , and k_3 , we can replace d_t with $k_1 d_t + k_2 t + k_3$, γ_i with $k_1 \gamma_i$, $\log \beta_i$ with $\log \beta_i - k_2/k_1$, and $\log \alpha_i$ with $\log \alpha_i - k_3/k_1$ without changing the model.

Equation (A2.5) uses z_{it} as an instrument for x_{it}^* to estimate the income coefficient g . If $\tilde{\mathbf{d}} = \mathbf{d}$, then multiplying by $V(\mathbf{d})$ and \mathbf{M} removes the household-specific intercepts, trends, and risk preferences, and \tilde{g} is the estimated slope in a regression of $y_{it} - \log \alpha_i/\gamma_i - d_t/\gamma_i - t \log \beta_i/\gamma_i$ on x_{it} using z_{it} as an instrument; assumptions F.1, F.2, F.3, and F.5 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 (A2.4) and (A2.5) together have a unique solution, this solution is a consistent estimator of (g, \mathbf{d}) . I iterate on (A2.4) and (A2.5) to find a pair $(\tilde{\mathbf{d}}, \tilde{g})$ that solves the two equations.²¹ The estimator can be adapted to allow heterogeneity only in risk preference or only in time preference by omitting \mathbf{t} or \mathbf{d} , respectively, from the \mathbf{V} matrix. Interpreting the estimator as just-identified GMM applied to (A2.5) and the first-order conditions of (A2.4) shows that the estimator is root- N consistent and asymptotically normal.²²

A2.2. GMM tests

The GMM tests allow nonseparability between consumption and leisure, so the equation to be estimated is

$$y_{it} = \frac{\log \alpha_i}{\gamma_i} + \frac{1}{\gamma_i} d_t + t \frac{\log \beta_i}{\gamma_i} + f z_{it} + g x_{it} + \epsilon_{it}. \quad (\text{A2.6})$$

The tests use both long and short time differences of variables. To write these differences concisely, for any ζ , define $\Delta_s \zeta_{it} = \zeta_{it} - \zeta_{i,t-s}$ and $\Delta_s^2 \zeta_{it} = (\zeta_{it} - \zeta_{i,t-1}) - (\zeta_{i,t-s} - \zeta_{i,t-s-1})$.

The tests are based on the following moment condition:

²¹In my data, 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.

²²The model cannot be estimated by GMM because the moment conditions may not have a unique solution. In practice, I construct standard errors and confidence intervals for the parameters by bootstrapping.

Assumption G.1. For all s and t ,

$$E[z_{is}e_{it}] = 0, \quad (\text{A2.7})$$

where e_{it} represents the error term in (9) or (10) depending on which equation is estimated.

I consider only a subset of the moment conditions that (A2.7) could generate. Long leads and lags of leisure are likely to be weak instruments, and not all leads and lags are available for any given observation in an unbalanced panel. I therefore use a small set of relatively strong instruments: leisure and its first lag. Let \mathbf{h}_{it} denote the vector of instruments $[1 \ z_{it} \ z_{i,t-1}]'$.

No heterogeneity: Suppose risk and time preferences are identical across households, so we can normalize $\gamma_i = 1$ and $\log \beta_i = 0$ for all i . If we difference (A2.6) to eliminate the household-specific intercept, we obtain

$$\Delta_3 y_{it} = \Delta_3 d_t + f \Delta_3 z_{it} + g \Delta_3 x_{it} + \Delta_3 \epsilon_{it}.$$

If leisure and its first lag are uncorrelated with ϵ_{it} , the following moment conditions hold:

$$E[\mathbf{h}_{it}(\Delta_3 y_{it} - \Delta_3 d_t - f \Delta_3 z_{it} - g \Delta_3 x_{it})] = \mathbf{0}, \quad t = 4, \dots, T. \quad (\text{A2.8})$$

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 (A2.8) is zero.

Heterogeneity only in time preference: The right-hand side of (A2.6) now includes a household-specific time trend, which can be eliminated by taking second differences:

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

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

$$E[\mathbf{h}_{it}(\Delta_2^2 y_{it} - \Delta_2^2 d_t - f \Delta_2^2 z_{it} - g \Delta_2^2 x_{it})] = \mathbf{0}, \quad t = 4, \dots, T. \quad (\text{A2.9})$$

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 households have identical risk preferences, whether time preferences vary or not.

Heterogeneity only in risk preference: If households have the same rate of time preference, but risk preferences vary, the second difference of (A2.6) is

$$\Delta_2 y_{it} = \frac{1}{\gamma_i} \Delta_2 d_t + f \Delta_2 z_{it} + g \Delta_2 x_{it} + \Delta_2 e_{it}. \quad (\text{A2.10})$$

Equation (A2.10) is equivalent to the model studied by Ahn et al. (2001):²³ $\tilde{y}_{it} = b_i \tilde{d}_t + \tilde{\mathbf{x}}'_{it} \boldsymbol{\theta} + \tilde{\varepsilon}_{it}$. Ahn et al. propose a quasi-differencing estimator. Equation (A2.10) implies that

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

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

$$\mathbf{E} \left[\mathbf{h}_{it} \left\{ \begin{array}{l} \Delta_2 y_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 y_{i,t-1} - f \left(\Delta_2 z_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 z_{i,t-1} \right) \\ -g \left(\Delta_2 x_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 x_{i,t-1} \right) \end{array} \right\} \right] = \mathbf{0}. \quad (\text{A2.11})$$

A test for whether $g = 0$ using moment conditions (A2.11) 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.

Heterogeneity in risk and time preference: Second differences of (A2.6) can be quasi-differenced to produce moment conditions valid when both risk and time preferences vary: For $t = 4, \dots, T$,

$$\mathbf{E} \left[\mathbf{h}_{it} \left\{ \begin{array}{l} \Delta_1^2 y_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 y_{i,t-1} - f \left(\Delta_1^2 z_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 z_{i,t-1} \right) \\ -g \left(\Delta_1^2 x_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 x_{i,t-1} \right) \end{array} \right\} \right] = \mathbf{0}. \quad (\text{A2.12})$$

²³Indeed, Ahn et al. (2001) mention risk-sharing tests as a possible application of their model.

Under assumption G.1, a test for whether $g = 0$ using moment conditions (A2.12) is a valid test for full insurance regardless of variation in risk or time preferences.

Identification in all four sets of moment conditions requires that the instruments be correlated with the right-hand-side variables. Ahn et al. (2001) show that identification in (A2.11) and (A2.12) also requires that an instrument be correlated with γ_i , which in turn implies risk preferences must be heterogeneous. Section 4 provides intuition for this requirement.

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): 219–55.
- Aigner, Dennis J. 1973. "Regression with a Binary Independent Variable Subject to Errors of Observation." *Journal of Econometrics* 1(1): 49–59.
- Attanasio, Orazio, and Steven J. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104(6): 1227–62.
- 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): 537–79.
- 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): 1887–921.
- 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(6): 926–37.
- Cochrane, John H. 1991. "A Simple Test of Consumption Insurance." *Journal of Political Economy* 99(5): 957–76.
- Deaton, Angus. 1997. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. Baltimore: Johns Hopkins University Press for the World Bank.
- Diamond, Peter A. 1967. "The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty." *American Economic Review* 57(4): 759–76.

- 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* 28(1): 229–303.
- Fuchs-Schündeln, Nicola, and Matthias Schündeln. 2005. "Precautionary Savings and Self-Selection: Evidence from the German Reunification 'Experiment.'" *Quarterly Journal of Economics* 120(3): 1085–120.
- Guiso, Luigi, Tullio Jappelli, and Luigi Pistaferri. 2002. "An Empirical Analysis of Earnings and Employment Risk." *Journal of Business & Economic Statistics* 20(2): 241–53.
- Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff. 1996. "Risk-Sharing between and within Families." *Econometrica* 64(2): 261–94.
- 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. 2008. "Insurance and Opportunities: A Welfare Analysis of Labor Market Risk." *Journal of Monetary Economics* 55(3): 501–25.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2009. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." Research Department Staff Report 432, Federal Reserve Bank of Minneapolis.
- Holtz-Eakin, Douglas, Whitney Newey, and Harvey S. Rosen. 1988. "Estimating Vector Autoregressions with Panel Data." *Econometrica* 56(6): 1371–95.
- Horowitz, Joel L. 2001. "The Bootstrap." In *Handbook of Econometrics*, vol. 5, ed. James J. Heckman and Edward Leamer, 3159–228. 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): 103–17.
- Lawley, D. N., and A. E. Maxwell. 1963. *Factor Analysis as a Statistical Method*. London: Butterworth.
- 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): 209–44.
- Mace, Barbara J. 1991. "Full Insurance in the Presence of Aggregate Uncertainty." *Journal of Political Economy* 99(5): 928–56.
- Mazzocco, Maurizio, and Shiv Saini. Forthcoming. "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences." *American Economic Review*.
- Moskowitz, Tobias J., and Annette Vissing-Jørgensen. 2002. "The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle." *American Economic Review* 92(4): 745–78.
- 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): 384–94.
- Ogaki, Masao, and Qiang Zhang. 2001. "Decreasing Relative Risk Aversion and Tests of Risk Sharing." *Econometrica* 69(2): 515–26.
- 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.

Schulhofer-Wohl, Samuel. 2007. "Heterogeneity, Risk Sharing and the Welfare Costs of Risk." Ph.D. dissertation, University of Chicago.

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

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica* 62(3): 539–91.

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

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): 119–32.