Heterogeneity, Risk Sharing and the Welfare Costs of Idiosyncratic Risk

Sam Schulhofer-Wohl*

September 7, 2007

Abstract

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

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

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

1 Introduction

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

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

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

A statistical rejection of the null hypothesis of full insurance is not especially informative, since anecdotal evidence suffices to tell us that at least some idiosyncratic shocks – for example, winning a large lottery jackpot – are not fully insured. It is more interesting to investigate the magnitude of departures from full insurance: Are failures of insurance economically important or not? In this paper, I use a simple model of imperfect insurance to show that the coefficient on income in a risk-sharing regression measures how far the world

[1] Among many examples, full insurance has been rejected in data from the United States (Attanasio and Davis, 1996; Cochrane, 1991; Dynarski and Gruber, 1997; Hayashi et al., 1996), Côte d'Ivoire (Deaton, 1997), India (Munshi and Rosenzweig, 2005; Townsend, 1994), Nigeria (Udry, 1994) and Thailand (Townsend, 1995a,b). Mace (1991) does not reject efficiency in U.S. data, but Nelson (1994) overturns this result.

is from full insurance, but that the coefficient in the standard regression is biased upward. Households in the model trade in complete markets. However, transferring resources between households is costly. The effect of income on consumption, holding aggregate shocks constant, reflects 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. Standard risk-sharing regressions do not measure this effect, though, because they do not fully control for aggregate shocks. Using time indicator variables to represent aggregate shocks, as in the existing literature, assumes that aggregate shocks affect all households equally. But when risk preferences vary across households, even imperfect insurance requires less risk-averse households to bear more aggregate risk; their consumption moves more strongly with aggregate shocks. If less risk-averse households' income also moves more strongly with aggregate shocks, the usual regression will find a spuriously large correlation between income and consumption.

After developing the model in section 2, I demonstrate in section 3 that income *does* vary more with aggregate shocks for less risk-averse households. I classify respondents to the Health and Retirement Study by how they answer hypothetical questions about taking risky jobs. Restricted-access Social Security records show that earnings are more variable over time and more correlated with GDP for men who claim to be less risk averse. Theory suggests that this relationship should be no surprise. If we extend the imperfect insurance model to allow households to choose among risky jobs, choices depend not only on comparative advantage but also on risk preferences, and less risk-averse households take jobs where income carries more aggregate risk because this sorting reduces the need for costly ex post transfers.

The rest of the paper measures the relationship between income and consumption, and uses it to calculate how imperfect insurance is, while accounting for the relationship between risk preferences and income processes. Section 4 describes the econometric methods used. I focus on data with only a few observations per household and therefore treat households' preferences as nuisance parameters that must be eliminated from the equation. Section 5 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 heterogeneous preferences reduces the effect by one-fourth to one-half. The effect is small and, in many specifications, statistically indistinguishable from zero.

Section 6 puts the estimates in context by making some simple welfare calculations. If improving insurance did not change households' job choices, then fully insuring idiosyncratic risk would improve welfare by the same amount as increasing consumption by just a few tenths of one percent in every state of the world. The welfare gain is two to four times larger in estimates that ignore heterogeneous preferences; thus allowing heterogeneity substantially reduces the apparent welfare costs of risk. My results do not, however, imply that risk is unimportant for welfare. The reason I find that income shocks do not greatly affect consumption is that people choose jobs in part based on risk preferences. If these choices run counter to comparative advantage in productivity, they decrease output. Improving insurance could substantially increase welfare by allowing people to choose more productive careers.

This paper is related to small but growing literatures on the relationship between preferences and income processes, and on risk sharing with heterogeneous preferences. Bonin et al. (2006) show that more risk-averse people hold jobs with less idiosyncratic risk, as measured by cross-sectional dispersion in self-reported earnings. They do not analyze macroeconomic risk or the time series properties of individual earnings, and because they use survey data on earnings instead of administrative data, they cannot distinguish measurement error from true variability in earnings. Fuchs-Schündeln and Schündeln (2005) analyze the endogeneity of income risk in the context of measuring the rate of precautionary savings: Ignoring the possibility that more prudent people choose less risky jobs will lead to downward-biased estimates of the importance of precautionary savings. Dynarski and Gruber (1997) test for full insurance while allowing heterogeneity in time preference but not in risk preference.

Townsend (1994) allows heterogeneous risk preferences in part of his seminal paper on testing for full insurance, but he employs a short panel and runs a separate regression for each household, implying that hypotheses can be tested only by assuming a parametric distribution for the regression error term. Kurosaki (2001) makes a similar analysis. Dubois (2001) tests for full insurance with heterogeneous preferences while allowing the coefficient of relative risk aversion to be a function of observed household characteristics, but he rules out unobserved heterogeneity. Finally, Mazzocco and Saini (2007) test for full insurance with heterogeneous preferences in data from rural India. Their dataset has more than 100 observations per household, which allows them to avoid functional form assumptions about utility functions; to estimate each household's preferences, rather than treat preferences as nuisance parameters; and to test for heterogeneity in preferences. However, they do not estimate a model of partial insurance and so cannot say how economically important a rejection of full insurance is. Further, their methods do not apply to short panels and thus cannot be used to study risk sharing using available household-level data from the United States.

2 Interpreting Risk-Sharing Regressions Under Partial Insurance

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 develops a model in which insurance fails because transferring resources between households is costly, and shows how the model can be used to interpret consumption allocations and risk-sharing regressions in a world without perfect insurance.

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

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

where u_i is concave in c. Transferring resources between households to provide insurance is costly. If household i has income X and consumes $c \neq X$, an additional quantity $\phi_i h(X, c)$ of the consumption good is destroyed, where $\phi_i \geq 0$. The parameter ϕ_i measures how difficult it is to insure household i. I assume that $h(X, c) = 0$ if $X = c$ and $h(X, c) > 0$ otherwise, and that h is convex and twice differentiable, with $\partial^2 h/\partial c \partial X < 0$. For now, I assume income and leisure are given exogenously. There is no storage.

This model of transactions costs does not correspond to any particular real-world institution. But 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. It is worth noting that formally modeling only one risk-sharing institution, such as the optimal contract under limited commitment (e.g., Ligon et al., 2002), would in some sense also be a reduced form, since it would set aside all other institutions that are present in reality. The recent literature has also proposed reduced form imperfect-insurance models based on the permanent income hypothesis: Some shocks are fully insured while households use bond-holdings to self-insure against other shocks, and the relative variances of the two kinds of shocks characterize the overall degree of insurance (Blundell et al., 2006; Heathcote et al., 2007). Although such models may feel more realistic, they are not consistent with the excess sensitivity of consumption to anticipated income shocks (e.g., Flavin, 1981). My model is compatible with excess sensitivity since transactions costs keep households from smoothing predictable as well as unpredictable fluctuations.

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

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

subject to the constraint that, for each date and state, aggregate income is at least as large as aggregate consumption plus the cost of transfers between households:

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

Assume an interior solution and let $\Pr(s_t)\lambda_t(s_t)$ be the Lagrange multiplier on the aggregate resource constraint. If there were no transfer costs ($\phi_i = 0$), the first-order condition for consumption would be

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

which is the familiar full-insurance allocation in which idiosyncratic income does not affect consumption. When transfer costs are not zero, the first-order condition for consumption becomes

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

Equation (1) shows that transfer costs distort the allocation away from the first best. In particular, holding leisure constant, consumption rises with income: Differentiating (1) with respect to income and applying the implicit function theorem yields

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

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

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

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

as an approximation to the household's utility function.[2] This approximation can model differences in utility functions or differences in relative risk aversion among households that have identical non-CRRA preferences but varying wealth. Let the cost of transferring resources between households be

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

which satisfies the stated assumptions on h as long as $\log(X/c) < 1$. Assume that consumption is measured with multiplicative error: Observed consumption is $c_{it} = e^{\epsilon_{it}}c^*_{it}$. Then the optimal allocation (1) satisfies

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

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

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

are small – this equation is approximately equivalent to

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

Equation (2) says aggregate shocks λ_t have a larger effect on households that have smaller coefficients of relative risk aversion γ_i or smaller transfer costs ϕ_i. Consumption rises faster for households with a large rate of time preference β_i or a large elasticity of intertemporal substitution $1/\gamma_i$; these differences in patience can also be interpreted as differences in consumption trends for households at different points in the life cycle. Income has a larger effect on households that are less risk averse (small γ_i) or more difficult to insure (large ϕ_i); 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.[3]

I focus on estimating the mean effect of income on consumption, $g \equiv \mathrm{E}[\phi_i/(\phi_i + \gamma_i)]$, so I rewrite (2) as

$$\log c_{it} = \frac{\log \alpha_i}{\phi_i + \gamma_i} + \frac{1}{\phi_i + \gamma_i}(-\log \lambda_t) + \frac{\log \beta_i}{\phi_i + \gamma_i} t + \frac{\xi_i}{\gamma_i} \log \ell_{it} + g \log X_{it} + \epsilon_{it}^{het}. \quad (3)$$

Equation (3) is identical to equation (2) if ϕ_i/γ_i is the same for all households i. Otherwise, the heterogeneous effect of income in (2) is absorbed into the error term in (3). I assume for now that the new error term ϵ_{it}^{het} remains uncorrelated with income X_{it}. This assumption keeps the analysis simple but is not otherwise necessary; in section 4, when I derive methods for estimating (3), I give weaker conditions under which heterogeneous coefficients on income do not bias the results.

One can measure the extent of risk sharing by estimating the parameters of (3). If the coefficient on income is zero, then $\phi_i = 0$ for all households and there is full insurance.

[3] One can derive the same equation in levels by assuming constant absolute risk aversion preferences, $h(X, c) = (c - X)^2/2$ and additive measurement error, but the non-negativity constraint on measured consumption makes additive errors statistically unattractive.

Otherwise, the hypothesis of full insurance is rejected, and the coefficient – the effect of income on consumption, holding aggregate shocks constant – measures the average ratio of costs to benefits of risk sharing, $\mathrm{E}[\phi_i/\gamma_i]$. (In section 6, I use this ratio to perform welfare calculations.) Almost all previous analyses of risk sharing have not estimated equation (3), however.[4] The studies note that with identical risk and time preferences – $\gamma_i = \gamma$ and $\beta_i = \beta$ – one can define $\tilde{\lambda}_t = \lambda_t \beta^{-t}$ and estimate

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

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

The flaw in (4) is omitted variable bias. If the true model is (3) but a researcher mistakenly estimates (4), the error term in (4) is

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

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

For simplicity, consider the case where all households have the same rate of time preference, so the time trend can be absorbed into the aggregate shock $\log \lambda_t$. Suppose that

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

[5] Deaton and Paxson (1994) propose an alternative test of full insurance that also relies crucially on common preferences: According to (4), if consumption is separable from leisure, the cross-sectional variance of log consumption is constant over time under full insurance. If the cross-sectional variance rises over time, insurance is deemed imperfect. Equation (3) shows, however, that when households have different preferences, the cross-sectional variance of log consumption changes over time even under full insurance.

income depends on common and idiosyncratic shocks: $\log X_{it} = q_i m_t + u_{it}$. Assuming u_{it} and ϵ_{it}^{het} are i.i.d.,

$$\mathrm{Cov}\big(\log X_{it}, \epsilon_{it}^{equal}\big) = \mathrm{Cov}\big(q_i m_t, [1/(\phi_i + \gamma_i) - 1/(\phi + \gamma)](-\log \lambda_t)\big). \tag{5}$$

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

Other than the Dubois (2001) and Mazzocco and Saini (2007) papers on heterogeneity and risk sharing mentioned in the introduction, previous studies have addressed the possible bias by finding assumptions under which it does not arise. For example, Cochrane (1991) tests how households smooth shocks other than income that he argues are uncorrelated with preferences, while Ogaki and Zhang (2001) allow decreasing relative risk aversion, $u_i(c) = (c - \theta)^{1-\gamma}/(1 - \gamma)$, so relative risk aversion depends on the level of consumption but nothing else. My approach is different. In the next section, I use theory and data to show that preferences *are* correlated with income processes even after controlling for observed characteristics. Then, in section 4, I derive estimators that eliminate the resulting bias in risk-sharing regressions.

3 Occupation Choice and Risk Preferences

I investigate the relationship between preferences and income processes by extending the imperfect insurance model of the previous section to allow for occupation choice. The economy has a variety of jobs $p \in P$. Each household picks a job at $t = 0$ and keeps it forever. Multiple households can hold the same job. All jobs produce the good c, but in different amounts depending on the state and the household's exogenously given skills. Hours worked are a function of the job, the state and the skills. If a household with skills κ_i has job p, the household's income is $X_{it}(s_t) = X(p, \kappa_i, s_t)$ and its leisure is $\ell_{it}(s_t) = \ell(p, \kappa_i, s_t)$.

Under full insurance, households would choose jobs according to comparative advantage in productivity. Imperfect insurance motivates households also to consider their risk preferences and the riskiness of the job. The planner's problem, which can be decentralized, is now to assign jobs and consumption to households to solve

$$\max_{\{p_i\},\{c_{it}(s_t)\}} \sum_i \alpha_i \mathrm{E}_0 \sum_{t=0}^{T} \beta_i^t u_i[c_{it}(s_t), \ell(p_i, \kappa_i, s_t)] \quad \text{subject to}$$

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

Suppose utility does not depend on leisure. (Similar but less transparent results can be derived when utility depends on leisure.) Let V be the maximized value of the planner's Lagrangian. Consider changing household i's job in a way that increases the household's income only in state s_t at date t, leaving income unchanged for all other dates and states. The change in the planner's Lagrangian from this change in the household's job is

$$\frac{dV}{dX_i(s_t)} = \frac{\partial V}{\partial X_i(s_t)} + \sum_j \alpha_j \mathrm{E}_0 \sum_{t=0}^{T} \frac{\partial V}{\partial c_{jt}(s_t)} \frac{dc_{jt}(s_t)}{dX_i(s_t)} = \frac{\partial V}{\partial X_i(s_t)} = \lambda_t(s_t) \left(1 - \phi_i \frac{\partial h}{\partial X}\right), \quad (6)$$

where $\partial V / \partial c_{jt}(s_t) = 0$ for all households j and dates t by the envelope theorem. The

Lagrange multiplier $\lambda_t(s_t)$ is the social marginal utility of income. Equation (6) therefore shows that transfer costs distort the marginal value of income to i by a factor $1 - \phi_i \partial h / \partial X$. The assumptions on the transfer costs h imply that $\partial h/\partial X$ is negative when consumption exceeds income and positive when consumption falls short of income. Hence, relative to the social marginal utility of income, household i's optimal job choice overweights income in states where i receives transfers from other households. Two examples illustrate how job choices are distorted.

Occupation choice as insurance against aggregate risk: Suppose there is only aggregate risk, and some households are more risk averse than others. Less risk-averse households should insure more risk-averse households against aggregate shocks. Hence if all households have the same skills κ, transfers flow from less risk-averse households when the aggregate shock is bad and to less risk-averse households when the aggregate shock is good. Less risk-averse households put more weight on income earned when the aggregate shock is good; all else equal, *less risk-averse households choose jobs where income varies more with the aggregate shock.*

Occupation choice and intertemporal substitution: Suppose some households discount the future more than others. Relatively patient households should consume less in early periods and more in later periods, so they give more weight to income earned in later periods; all else equal, *more patient people choose jobs where income rises faster over time.*

These informal examples suggest exactly the relationship between preferences and income processes that would produce upward-biased coefficients in the standard risk-sharing regression. More formal results are difficult to obtain because, with background risk, local measures such as the coefficient of absolute risk aversion do not order households' preferences. In any event, theory can shed only limited light on the relationship between risk preferences and income processes, because workers' skills will also affect their job choices and skills could be correlated with preferences in many ways. The relationship between preferences and income processes is ultimately an empirical question, and I now turn to empirical evidence.

3.1 Data

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

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

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

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

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

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

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

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

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

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

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

3.2 Analysis

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

$$\log\left(earnings_{it}\right) = \pi_{0i} + \pi_{1j}\log\left(GDP_t\right) + \mathbf{x}_{it}'\mathbf{\Pi}_j + \upsilon_{ijt}, \tag{7}$$

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

I interpret (7) as an equation that a worker could use to predict his lifetime earnings profile. The coefficient on GDP shows how income varies over the business cycle. The experience and trend variables included as controls reflect how income would vary over time if there were no aggregate shocks. The variance of the error term measures idiosyncratic fluctuations in the worker's earnings. Equation (7) is an admittedly simplistic specification for earnings dynamics – for example, one might prefer to examine income growth rather than the level of income – but with top-coded data, I cannot take first differences to compute innovations to income. In addition, my goal here is only to study the correlation between earnings and aggregate shocks, not to estimate a fully specified model of the time series properties of individual earnings. In the Internet appendix, I use data from the Panel Study of Income Dynamics, where top-coding is minimal but risk-preference data are lacking, to show that regressions in differences would likely produce similar results.[11]

The tobit model for censoring assumes that the error term is normally distributed and has constant variance over time for each worker. These assumptions allow one to write down the likelihood for the censored data. Actually estimating equation (7) by maximum likelihood would be computationally challenging, however, because it would require calculating high-dimensional integrals of normal densities. Instead, I employ the Bayesian Markov Chain Monte Carlo method to find the posterior distribution of the parameters.[12] I allow the error term to be correlated across workers within a risk-tolerance group in each year: $\upsilon_{ijt} = \pi_{jt} + \nu_{it}$, where π_{jt} is a normally distributed random effect. I employ an uninformative prior for the regression coefficients and a standard weakly informative conjugate prior for the variances of the errors. The Internet appendix gives details.

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

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

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

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

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

Equation (7) is misspecified if workers within a risk-tolerance group have different income processes. To see whether this biases my estimates, I estimate regressions in which the coefficients differ for each worker:

$$\log\left(earnings_{it}\right) = \pi_{0i} + \pi_{1i}\log\left(GDP_t\right) + \pi_{2i}experience_{it} + \pi_{3i}experience_{it}^2 + \nu_{it}. \tag{8}$$

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

Table 3 shows the means of the regression coefficients and of the variance of idiosyncratic errors in (8) for workers with high and low risk tolerance. Allowing the coefficients to vary across workers greatly reduces the precision of the results, but the point estimates tell the same story as before: High-risk-tolerance workers face more aggregate and idiosyncratic risk. For example, a 1 percent rise in per capita real GDP, relative to trend, increases earnings by an average of 2.0 percent for high-risk-tolerance workers but by just 1.65 percent for low-risk-tolerance workers. The differences in point estimates persist when I use different aggregate variables and when I restrict the sample to white, native-born men with exactly 12 years of education.

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

3.3 Caveats

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

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

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

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

4 Robust Risk-Sharing Regressions: Econometrics

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

The two approaches are complementary. Given only a few time periods of data, the factor analysis approach requires strong assumptions about the distribution of the error term. It also assumes that consumption and leisure are separable. However, if these assumptions hold, the factor estimator is valid whether preferences are correlated with income processes or not. Thus comparing the factor model estimates with the common-preferences regression can show whether the common-preferences regression is misspecified. The GMM approach makes fewer assumptions, but it is valid only when preferences are correlated with income processes. It is meaningless to compare a GMM estimate of (3) with an estimate of the common preferences model (4): The estimates are guaranteed to differ since each is valid only under assumptions that make the other invalid.

The two estimators also perform different experiments. The factor model, which uses a generalized fixed effects transformation to remove household-specific preference parameters from the equation, analyzes how consumption responds to deviations of household income from the mean of income over time. The GMM method, which uses a generalized differencing transformation, analyzes how consumption responds to changes in income from one year to the next. While the imperfect insurance model in section 2 is static, dynamics are worth keeping in mind in interpreting the empirical results. Hayashi et al. (1996) note that if insurance is imperfect but households know about income shocks in advance and adjust consumption accordingly, consumption growth is uncorrelated with contemporaneous income growth. The factor model measures the relationship between consumption and income over longer time periods and thus may detect more failures of insurance.

In section 5, I will apply the estimators to survey data on income and consumption from the Panel Study of Income Dynamics (PSID). Because households may report income inaccurately, the estimators must account for measurement error. In fact, accounting properly for measurement error is crucial to demonstrating that ignoring heterogeneity in preferences leads to an upward-biased coefficient on income. Classical measurement error in income will bias the coefficient on income toward zero if the regression is estimated by ordinary least squares (OLS). The bias depends on the signal-to-noise ratio in measured income; the less signal, the smaller the estimated coefficient. Adding controls to a regression reduces the signal-to-noise ratio in the regressors. Therefore, adding controls for heterogeneous preferences to an OLS risk-sharing regression will automatically reduce the estimated coefficient on income, whether preferences actually vary or not. To truly demonstrate that accounting for heterogeneous preferences reduces the effect of income on consumption, I have to show that the coefficient changes even after accounting for the consequences of measurement error. Dealing with measurement error is straightforward in the case of GMM, which already uses instrumental variables. For the factor model, the problem is more complicated. The existing literature assumes that regressors are measured without error. I derive an instrumental variables method that is valid when income is measured with error.

The estimators for equation (3) in the literature assume a balanced panel, but the data I analyze form an unbalanced panel. In addition, the estimators assume the coefficient on income is the same for all households, while the model of section 2 has a different coefficient for each household. The rest of this section gives technical details of how I address these problems and can be skipped by readers interested only in the results.

4.1 Heterogeneity in the regression coefficients

I estimate common coefficients on income and leisure in (3). Households have identical coefficients under restrictive assumptions on preferences and technology: $f = \xi_i/(\phi_i + \gamma_i)$ and $g = \phi_i/(\phi_i + \gamma_i)$ for all i. Alternatively, the goal can be to estimate each coefficient's population mean, $f = \mathrm{E}[\xi_i/(\phi_i+\gamma_i)]$ and $g = \mathrm{E}[\phi_i/(\phi_i+\gamma_i)]$. To understand the assumptions necessary for estimating the mean coefficient, some notation is helpful. Let y_{it} stand for log consumption, x_{it} for log income and z_{it} for log leisure. Let $\tilde{a}_i^0 = \log \alpha_i/(\phi_i + \gamma_i)$, $\tilde{b}_i^0 = 1/(\phi_i+\gamma_i)$, $d_t = -\log \lambda_t$, $\tilde{r}_i^0 = \log \beta_i/(\phi_i+\gamma_i)$, $f_i = \xi_i/(\phi_i+\gamma_i) - f$ and $g_i = \phi_i/(\phi_i+\gamma_i) - g$. Then (3) can be written

$$y_{it} = \tilde{a}_i^0 + \tilde{b}_i^0 d_t + \tilde{r}_i^0 t + f z_{it} + g x_{it} + \epsilon_{it}^{het}, \quad \epsilon_{it}^{het} = f_i z_{it} + g_i x_{it} + \epsilon_{it}. \tag{9}$$

Following Wooldridge's (2005) analysis of panel random coefficients models, we can decompose income and leisure into household-specific projections on the aggregate shock and time trend and residuals from these projections:

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

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

The residuals $\ddot{x}_{it}^{het}$ and $\ddot{z}_{it}^{het}$ are the idiosyncratic components of income and leisure. Substituting these decompositions into (9) gives

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

where $a_i = \tilde{a}_i^0 + g_i \tilde{a}_i^x + f_i \tilde{a}_i^z$ and where b_i and r_i are defined similarly.[13]

Consider estimating (3′) by instrumental variables with some variable w_{is} as an instrument. The instrument is strictly exogenous, and hence valid, if

$$\mathrm{E}[w_{is} e_{it}^{het}] = 0 \quad \forall\ s, t. \tag{10}$$

When there are common coefficients on income and leisure – so $f_i = 0$ and $g_i = 0$ for all i – (10) requires only that the instrument be uncorrelated with the consumption measurement error ϵ_{it}. When there are heterogeneous coefficients, (10) also requires that $\mathrm{E}[w_{is}\ddot{x}_{it}^{het} g_i] = \mathrm{E}[w_{is}\ddot{z}_{it}^{het} f_i] = 0$. That is, the covariance of the instrument with the idiosyncratic component of the regressors must not depend on the random coefficient.[14] My instruments are leads and lags of leisure. It is difficult to know whether these instruments satisfy the conditional homoscedasticity requirement; I assume that they do or that the coefficients are identical across households, and for the GMM estimators, I use a test of overidentifying restrictions

[13] This decomposition shows that the household-specific coefficients on aggregate shocks b_i do not equal the elasticity of intertemporal substitution $\tilde{b}_i^0 = 1/\gamma_i$ unless the coefficient on income is the same for all households. Thus, in general, (3′) can be used to measure the quality of risk sharing – given by the income coefficient g – but not the distribution of risk preferences.

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

to check my assumptions.

Stronger assumptions are needed to handle random coefficients on income and leisure in the common-preferences model (4). If we decompose income and leisure into projections on a household fixed effect and a common aggregate shock,

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

then (4) becomes

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

A valid instrument for (4′) must satisfy $\mathrm{E}[w_{is}\ddot{x}_{it}^{equal} g_i] = \mathrm{E}[w_{is}\ddot{z}_{it}^{equal} f_i] = 0$. Since $\ddot{x}_{it}^{equal} = \ddot{x}_{it}^{het} + (\tilde{b}_i^x - \tilde{b}^x)d_t + \tilde{r}_i^x t$, this assumption is stronger than was needed for (3′). Thus, as observed by Wooldridge (2005) for random-coefficients panel data models in general, heterogeneous coefficients on income will cause fewer problems in (3′) than in (4′) even if all households have the same preferences.

4.2 Factor model estimation

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

Assume that preferences are separable in consumption and leisure, so leisure z does not appear on the right-hand side of (3′) and is not part of the error term e_{it}^{het}. Subtracting the within-household mean of each variable from (3′) removes the fixed effect a_i:

$$u_{it}^y = b_i u_{it}^d + r_i u_{it}^t + g u_{it}^x + u_{it}^{e\_het}, \tag{3$''$}$$

where

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

and where u_{it}^x, u_{it}^d, u_{it}^t and $u_{it}^{e\_het}$ are defined similarly. Likewise, (4′) becomes

$$u_{it}^y = d_t + g u_{it}^x + u_{it}^{e\_equal}. \tag{4''}$$

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

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

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

$$\text{such that} \quad \sum_{t=1}^T \tilde{d}_t = \sum_{t=1}^T t \tilde{d}_t = 0, \ \sum_{t=1}^T \tilde{d}_t^2 = 1, \ \tilde{d}_1 > 0. \tag{11}$$

The constraints are normalizations. The first normalization fixes the level and trend of d_t, the second fixes the scale, and the third fixes the sign. To guarantee a unique solution, I also assume that there are at least four time periods, that each household is observed at least three times and that some households are observed four times.[15]

The estimator in (11) appears to suffer from an incidental parameters problem because the number of parameters increases as either N or T increases. We can eliminate the parameters indexed by i to solve this problem. Let $\mathbf{u}_i^x = [u_{i1}^x \cdots u_{iT}^x]'$ and define $\mathbf{u}_i^y$, $\mathbf{u}_i^d$, $\mathbf{u}_i^t$, $\mathbf{u}_i^{e\_het}$ and $\mathbf{d}$ analogously. Let $\boldsymbol{\Delta}_i^d = [\mathbf{u}_i^d \ \mathbf{u}_i^t]$, and let $\tilde{\boldsymbol{\Delta}}_i^d$ denote the matrix in which $\tilde{d}_t$ replaces d_t. Also let $\mathbf{I}_i$ be the $T \times T$ matrix with ι_{it} on the diagonal and zeros elsewhere. The objective function in (11) is strictly convex, and first-order conditions are necessary and sufficient for a minimum. Thus we can replace $\tilde{b}_i$ and $\tilde{r}_i$ in (11) with the corresponding first-order conditions to produce the following equivalent estimator of $\mathbf{d}$ and g:

$$(\hat{\mathbf{d}}, \hat{g}) = \min_{\tilde{\mathbf{d}}, \tilde{g}} \frac{1}{N} \sum_{i=1}^{N} (\mathbf{u}_i^y - \tilde{g}\mathbf{u}_i^x)' \left[\mathbf{I}_i - \tilde{\boldsymbol{\Delta}}_i^d \left(\tilde{\boldsymbol{\Delta}}_i^{d\prime} \hat{\boldsymbol{\Delta}}_i^d \right)^{-1} \tilde{\boldsymbol{\Delta}}_i^{d\prime} \right] (\mathbf{u}_i^y - \tilde{g}\mathbf{u}_i^x), \tag{12}$$

again subject to normalizations on $\mathbf{d}$. The estimator in (12) employs a generalized fixed effects transformation. To see this, let $V(\mathbf{d}) = \left[\mathbf{I}_i - \boldsymbol{\Delta}_i^d (\boldsymbol{\Delta}_i^{d\prime} \boldsymbol{\Delta}_i^d)^{-1} \boldsymbol{\Delta}_i^{d\prime} \right]$. Since $V(\mathbf{d})\boldsymbol{\Delta}_i^d = 0$, we could – if $\mathbf{d}$ were known – multiply (3″) by $V(\mathbf{d})$ to obtain

$$V(\mathbf{d})\mathbf{u}_i^y = gV(\mathbf{d})\mathbf{u}_i^x + V(\mathbf{d})\mathbf{u}_i^{e\_het}, \tag{13}$$

which does not contain the household fixed effects b_i and r_i. Because $\mathbf{d}$ is not known, (12) uses the estimated $V(\hat{\mathbf{d}})$ to transform the model and eliminate the fixed effects.

Kiefer (1980) shows that, if the error term e_{it}^{het} is homoscedastic and serially uncorrelated and if there is no individual-specific intercept or trend, the estimator in (12) is consistent in a balanced panel as long as income is measured without error. The extension to individual-specific intercepts and trends and an unbalanced panel is straightforward.[16] However, if income is measured with error, the estimated income coefficient is biased toward zero. I derive

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

[16] See the Internet appendix.

an instrumental variables version of the estimator to eliminate the attenuation bias. The idea of the estimator is that if $\mathbf{d}$ were known, the income coefficient g could be estimated from (13) using an instrument for x, while if g were known, $\mathbf{d}$ could be estimated by factor analysis regardless of measurement error in x. Because neither g nor $\mathbf{d}$ is known, the procedure iterates between estimating the two parameters.

I rely on a version of the usual exogeneity assumption in linear regression:

Assumption 1. $\mathrm{E}[x_{is}e_{it}^{het}|\boldsymbol{\iota}_i] = \mathrm{E}[b_i e_{it}^{het}|\boldsymbol{\iota}_i] = \mathrm{E}[r_i e_{it}^{het}|\boldsymbol{\iota}_i] = \mathrm{E}[e_{it}^{het}|\boldsymbol{\iota}_i] = 0$ *for all* i, s, t.

Assumption 1 rules out correlations between the error e_{it}^{het} and income that could not be distinguished from the true effect of income on consumption. Writing the expectations conditional on $\boldsymbol{\iota}_i$ says that the exogeneity assumption holds regardless of whether data are missing for a given household at some date; in this sense, missing observations are assumed to be missing at random. If there are heterogeneous coefficients on income, income appears in the error and the assumption requires that 1) $\mathrm{E}[b_i g_i \ddot{x}_{it}^{het}|\boldsymbol{\iota}_i] = \mathrm{E}[r_i g_i \ddot{x}_{it}^{het}|\boldsymbol{\iota}_i] = 0$ and 2) $\mathrm{E}[x_{is}\ddot{x}_{it}^{het} g_i|\boldsymbol{\iota}_i] = 0$. That is, 1) the variance-covariance matrix of preferences and technology (b_i, r_i, g_i) is not correlated with idiosyncratic income shocks, and 2) the variance-covariance matrix of idiosyncratic income shocks is not correlated with the individual-specific effect of income on consumption. The second requirement, in particular, is strong since the model of section 3 shows that households may choose jobs on the basis of risk preferences or risk-sharing costs, both of which enter g_i. However, the model does not predict a sign for the correlation between g_i and the variance of income if the correlation is not zero: g_i is large if a household is difficult to insure or relatively risk tolerant, but difficult-to-insure households will choose less risky jobs, while risk-tolerant households will choose riskier jobs. Recall that even stronger assumptions would be needed to deal with heterogeneous coefficients on income in the common-preferences equation ($4''$).

I also make a standard assumption in factor analysis with fixed T:

Assumption 2. $\mathrm{E}[\mathbf{e}_i^{het}\mathbf{e}_i^{het\prime}] = \sigma_i^2\mathbf{I}$ *for all* i.

Assumption 2 rules out a factor structure in the error term that could not be distinguished from the true factor structure $b_i d_t$. Intuitively, homoscedasticity is required because changing $\hat{\mathbf{d}}$ in the objective function in (12) changes the weights put on observations at various dates. If the variance of e_{it}^{het} were particularly small at some date, the objective function could be minimized by choosing $\hat{\mathbf{d}}$ to put heavy weight on that date, regardless of the true value of $\mathbf{d}$. If there are heterogeneous coefficients on income, this assumption will fail if idiosyncratic income shocks $\ddot{x}_{it}^{het}$ are serially correlated or heteroscedastic over time for a given household.[17]

I assume that instead of income x_{it}, we observe a variable x_{it}^* that satisfies the following assumption:

Assumption 3. $x_{it}^* = x_{it} + v_{it}$, *where* $\mathrm{E}[\mathbf{v}_i\mathbf{v}_i'] = \zeta_i^2\mathbf{I}$ *for all* i.

Assumption 3 is stronger than typical assumptions in measurement error models in that the measurement errors must be homoscedastic and serially uncorrelated. The variance-covariance matrix of the measurement errors in the regressor must be restricted in the same way as the variance-covariance matrix of the regression error term e_{it}^{het} because the measurement errors become part of the error term in the factor analysis part of the procedure.

Finally, I require an instrument. I assume that we observe z_{it} that satisfies:

Assumption 4. $\mathrm{E}[\mathbf{z}_i\mathbf{e}_i^{het\prime}] = \mathrm{E}[\mathbf{z}_i\mathbf{v}_i'] = \mathbf{0}$ *and* $\mathrm{E}[\mathbf{z}_i'V(\tilde{\mathbf{d}})\mathbf{x}_i] \neq 0$ *for all* $\tilde{\mathbf{d}}$ *in a neighborhood of* $\mathbf{d}$.

Assumption 4 says that the instrument is uncorrelated with the regression error term and the measurement errors. It also says that the instrument is correlated with income after controlling for aggregate shocks, not just for the true aggregate shocks but for incorrect aggregate shocks in the neighborhood of the true shocks. The full rank condition must hold

[17] With large-T asymptotics, assumption 2 could be weakened to allow heteroscedasticity and serial correlation, as in the approximate factor model of Chamberlain and Rothschild (1983).

for incorrect aggregate shocks because, in any finite sample, the estimated aggregate shocks will not equal the true aggregate shocks.

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

$$\tilde{\mathbf{d}}(\tilde{g}) = \arg\min_{\hat{\mathbf{d}}} \frac{1}{N} \sum_{i=1}^{N} (\mathbf{u}_i^y - \tilde{g}\mathbf{u}_i^{x^*})' \left[\mathbf{I}_i - \hat{\boldsymbol{\Delta}}_i^d \left(\hat{\boldsymbol{\Delta}}_i^{d\prime} \hat{\boldsymbol{\Delta}}_i^d \right)^{-1} \hat{\boldsymbol{\Delta}}_i^{d\prime} \right] (\mathbf{u}_i^y - \tilde{g}\mathbf{u}_i^{x^*}) \tag{14a}$$

$$\tilde{g}(\tilde{\mathbf{d}}) = \frac{\sum_{i=1}^{N} \mathbf{u}_i^{y\prime} \left[\mathbf{I}_i - \tilde{\boldsymbol{\Delta}}_i^d \left(\tilde{\boldsymbol{\Delta}}_i^{d\prime} \tilde{\boldsymbol{\Delta}}_i^d \right)^{-1} \tilde{\boldsymbol{\Delta}}_i^{d\prime} \right] \mathbf{u}_i^z}{\sum_{i=1}^{N} \mathbf{u}_i^{x^*\prime} \left[\mathbf{I}_i - \tilde{\boldsymbol{\Delta}}_i^d \left(\tilde{\boldsymbol{\Delta}}_i^{d\prime} \tilde{\boldsymbol{\Delta}}_i^d \right)^{-1} \tilde{\boldsymbol{\Delta}}_i^{d\prime} \right] \mathbf{u}_i^z} \tag{14b}$$

Equation (14a) uses the least squares procedure (12) to estimate $\mathbf{d}$, taking as given some value $\tilde{g}$ for the coefficient on income. If $\tilde{g} = g$, the measurement error in x_{it}^* simply increases the variance of the residuals and the least squares procedure will continue to provide a consistent estimator of $\mathbf{d}$. Hence $\tilde{\mathbf{d}}(\tilde{g})$ will be consistent for $\mathbf{d}$ if $\tilde{g}$ is a consistent estimator of g. Equation (14b) uses z_{it} as an instrument for x_{it}^* to estimate the income coefficient g. As in (13), knowledge of $\mathbf{d}$ is required to eliminate the household-specific preference parameters. Assumptions 3 and 4 guarantee that z_{it} is a valid instrument. Thus $\tilde{g}(\tilde{\mathbf{d}})$ will be consistent for the income coefficient g if $\tilde{\mathbf{d}}$ is a consistent estimator of $\mathbf{d}$. It follows that, if (14a) and (14b) together have a unique solution, this solution is a consistent estimator of $(g, \mathbf{d})$.

A pair $(\tilde{\mathbf{d}}, \tilde{g})$ that solves both (14a) and (14b) can be found iteratively: Start with a guess for $\tilde{g}$; use (14a) to compute $\tilde{\mathbf{d}}$; use this new $\tilde{\mathbf{d}}$ to produce a new $\tilde{g}$ using (14b); and repeat until convergence. In the data I analyze, the iteration always converges in 10 or fewer steps. I investigate whether the fixed point is unique by starting the iteration from several places and examining whether all give the same result. The estimator can be adapted to allow heterogeneity only in risk preference or only in time preference by omitting $\mathbf{d}$ or the time trend from the $\boldsymbol{\Delta}$ matrix.

Interpreting the estimator in (14a) as just-identified GMM applied to moment conditions defined by (14b) and the first-order conditions of (14a) shows that the estimator is root-N consistent and asymptotically normal. However, note that the model cannot be estimated by GMM applied to these moment conditions because the first-order conditions of (14a) may not have a unique solution. (For example, in a balanced panel, all eigenvectors of a particular matrix solve the first-order conditions of (14a).)

It may seem desirable to test the homoscedasticity and no serial correlation assumptions on the error term. In a balanced panel, this would be straightforward. For example, to test homoscedasticity, one could test whether the cross-sectional variance of residuals $u_{it}^y - \hat{b}_i \hat{u}_{it}^d - \hat{r}_i u_{it}^t - \hat{g} u_{it}^x$ was the same at each date t. However, in an unbalanced panel, matters become more difficult. Continuing with the example of homoscedasticity, the cross-sectional variance of residuals does not have to be the same at each date even if $\mathrm{Var}[e_{it}^{het}]$ is the same at each date for any given household i, because different households i may be in the sample at different dates. A test of homoscedasticity would actually test the joint hypothesis that the errors are homoscedastic and that there is no selection on the variance of the error term. Since selection on the variance of the error term would not bias the estimates, rejecting this joint hypothesis would not prove that the estimator is invalid, and I do not pursue such tests here.

4.3 GMM estimation

Assume that leisure is not necessarily separable from consumption but is strictly exogenous with respect to the error term:

$$\mathrm{E}[z_{is} e_{it}] = 0 \quad \forall\, s, t, \tag{15}$$

where e_{it} represents the error term in either (3′) or (4′) depending on which equation is to be estimated. Moment condition (15) forms the basis of GMM estimators for the coefficients in

(3′) and (4′) under a variety of assumptions about preferences. It differs in two ways from the assumptions used for the factor model. Condition (15) potentially allows preferences to depend on leisure and says nothing about the variance or serial correlation of the errors. However, if leisure is included on the right-hand side of the equation, (15) also prohibits classical measurement error in leisure – in contrast to the factor model, which assumed measurement error in leisure.

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

Throughout, I consider only a subset of the T^2 moment conditions that could be generated from (15). Long leads and lags of leisure are likely to be weak instruments. For example, leisure at some date is unlikely to be a good predictor of income 20 years in the future. GMM estimators that use many weak instruments can have poor finite sample properties (Stock et al., 2002). In addition, not all leads and lags are available for any given observation in an unbalanced panel. I therefore limit myself to a relatively small set of relatively strong instruments: leisure and its first lag.[18] Let $\mathbf{h}_{it}$ denote the vector of instruments $[1\ z_{it}\ z_{i,t-1}]'$.

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

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

(If preferences are separable in consumption and leisure, leisure can be omitted from the right-hand side.) Suppose leisure and its first lag – $z_{it}, z_{i,t-1}$ – are uncorrelated with the error e_{it}^{equal}. With heterogeneous coefficients on income, this requires $\mathrm{E}[g_i z_{is}\ddot{x}_{it}^{equal}] = \mathrm{E}[g_i z_{is}\ddot{z}_{it}^{equal}] =$

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

0. Then the following moment conditions hold:

$$\mathrm{E}\left[\mathbf{h}_{it}(\Delta_3 y_{it} - \Delta_3 d_t - f\Delta_3 z_{it} - g\Delta_3 x_{it})\right] = \mathbf{0}, \quad t = 2, \ldots, T. \tag{16}$$

In an unbalanced panel, the moment conditions for each date t can be constructed using all observations with data at $t-3$, $t-2$, $t-1$ and t. Not using income as an instrument prevents attenuation bias due to measurement error in income. Equation (16) lists $3(T-3)$ moment conditions in $T-1$ parameters $(\Delta_3 d_4, \ldots, \Delta_3 d_T, f, g)$. The parameters can be estimated by GMM and the $2T-8$ overidentifying restrictions tested using a chi-squared statistic. The estimated coefficient on income from the moment conditions (16) describes the effect of income on consumption if all households have the same risk and time preferences.

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

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

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

$$\mathrm{E}\left[\mathbf{h}_{it}(\Delta_2^2 y_{it} - \Delta_2^2 d_t - f\Delta_2^2 z_{it} - g\Delta_2^2 x_{it})\right] = \mathbf{0}, \quad t = 4, \ldots, T. \tag{17}$$

The parameters can again be estimated by GMM. The estimated coefficient on income using the moment conditions (17) describes the effect of income on consumption if all households have the same risk preferences, whether time preferences vary or not.

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

$$\Delta_2 y_{it} = b_i \Delta_2 d_t + f\Delta_2 z_{it} + g\Delta_2 x_{it} + \Delta_2 e_{it}. \tag{18}$$

Equation (18) is equivalent to the model studied by Ahn et al. (2001):[19] $\tilde{y}_{it} = b_i \tilde{d}_t + \tilde{\mathbf{x}}_{it}' \boldsymbol{\theta} + \tilde{e}_{it}$. Ahn et al. propose a quasi-differencing estimator. Equation (18) implies that

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

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

$$\mathrm{E}\left[\mathbf{h}_{it}\left\{\begin{array}{r} \Delta_2 y_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 y_{i,t-1} - f\left(\Delta_2 z_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 z_{i,t-1}\right) \\ -g\left(\Delta_2 x_{it} - \frac{\Delta_2 d_t}{\Delta_2 d_{t-1}} \Delta_2 x_{i,t-1}\right)\end{array}\right\}\right] = \mathbf{0}. \tag{19}$$

The estimated coefficient on income using (19) describes the effect of income on consumption when risk preferences vary across households but time preferences do not vary.

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

$$\mathrm{E}\left[\mathbf{h}_{it}\left\{\begin{array}{r} \Delta_1^2 y_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 y_{i,t-1} - f\left(\Delta_1^2 z_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 z_{i,t-1}\right) \\ -g\left(\Delta_1^2 x_{it} - \frac{\Delta_1^2 d_t}{\Delta_1^2 d_{t-1}} \Delta_1^2 x_{i,t-1}\right)\end{array}\right\}\right] = \mathbf{0}. \tag{20}$$

The estimated coefficient on income using these moment conditions describes the effect of income on consumption regardless of variation in risk or time preferences.

Identification of the parameters in all four sets of moment conditions – (16), (17), (19) and

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

(20) – requires that the instruments (differences of leisure) be correlated with the right-hand-side variables (differences of income and leisure). Ahn et al. (2001) show that, in addition, identification in (19) and (20) requires that at least one instrument be correlated with b_i. If this condition fails, the moment conditions (19) and (20) have infinitely many solutions and the GMM estimators are not defined. To see why an instrument must be correlated with b_i, consider the case where $f = g = 0$, so income and leisure do not appear on the right-hand side. Then the moment conditions in (19) reduce to the cross-sectional regression of $\Delta_2 y_{it}$ on $\Delta_2 y_{i,t-1}$ using z_{it} and its lags as an instrument for $\Delta_2 y_{i,t-1}$. Identification requires that $\Delta_2 y_{i,t-1}$ be correlated with z_{it} and its lags. But $\Delta_2 y_{i,t-1} = b_i \Delta_2 d_{t-1} + \Delta_2 e_{i,t-1}$; $e_{i,t-1}$ is assumed uncorrelated with z_{it}; and $\Delta_2 d_{t-1}$ does not vary in a cross-section of households at a given date. Hence $\Delta_2 y_{i,t-1}$ is correlated with z_{it} only if b_i is correlated with z_{it}.

The only way to avoid assuming that b_i is correlated with z_{it} is to somehow restrict the distribution of the error term, since then one can identify $\mathbf{d}$ from moment conditions of the form $\mathrm{E}[(\mathbf{u}_i^y - g\mathbf{u}_i^x - r_i\mathbf{u}^t)(\mathbf{u}_i^y - g\mathbf{u}_i^x - r_i\mathbf{u}^t)'] = \mathrm{E}[b_i^2]\mathbf{u}^d\mathbf{u}^{d\prime} + \mathbf{\Sigma}$, where $\mathbf{\Sigma} = \mathrm{E}[\mathbf{u}_i^e\mathbf{u}_i^{e\prime}]$ is known up to a small number of parameters. In the case of an unbalanced panel, there would be different moment conditions for each possible pattern of missing data: $\mathrm{E}[(\mathbf{u}_i^y - g\mathbf{u}_i^x - r_i\mathbf{u}_i^t)(\mathbf{u}_i^y - g\mathbf{u}_i^x - r_i\mathbf{u}_i^t)'|\boldsymbol{\iota}_i] = \mathrm{E}[b_i^2|\boldsymbol{\iota}_i]\mathbf{u}_i^d\mathbf{u}_i^{d\prime} + \mathrm{E}[\mathbf{u}_i^e\mathbf{u}_i^{e\prime}|\boldsymbol{\iota}_i]$. Unless one is willing to assume that the error covariance matrix $\mathrm{E}[\mathbf{u}_i^e\mathbf{u}_i^{e\prime}|\boldsymbol{\iota}_i]$ does not depend on the pattern of missing data $\boldsymbol{\iota}_i$, this implies a vast number of moment conditions each involving only a small number of observations, and GMM estimation would be infeasible.[20] The factor estimator, which in effect assumes $\mathrm{E}[\mathbf{u}_i^e\mathbf{u}_i^{e\prime}|\boldsymbol{\iota}_i] = \mathrm{E}[\sigma_i^2|\boldsymbol{\iota}_i](\mathbf{I} - \boldsymbol{\iota}_i\boldsymbol{\iota}_i'/\boldsymbol{\iota}_i'\boldsymbol{\iota}_i)$, is a special case where it is feasible to use moment conditions based on the distribution of the error term in an unbalanced panel.

Leisure is correlated with the risk preference parameter b_i, and the parameters in (19) and (20) are identified, if and only if the common-risk-preferences models (16) and (17) are

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

misspecified. Moment conditions (19) and (20) therefore may produce different results than (16) and (17) whether risk preferences vary or not; unlike with the factor model, comparing results between GMM estimators will not show whether risk preferences vary. However, if we maintain that heterogeneity in risk preferences is important, the moment conditions in (19) and (20) estimate the effect of income on consumption without using the factor model's assumptions about homoscedasticity, no serial correlation and separability.

All of the moment conditions examine the association between income and consumption over a four-year period. However, the particular income variations that are examined differ. Conditions (16) look at innovations to income, conditions (17) look at the difference between two consecutive innovations, and conditions (19) and (20) consider even more complicated differences. The need for quasidifferencing in (19) and (20) makes it impossible to examine exactly the same innovations in all four cases, but I try to make the results as comparable as possible by using the same instruments and a four-year period in each case.

5 Robust Risk-Sharing Regressions: Results

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

5.1 Data

I analyze data on consumption, income and leisure from the Panel Study of Income Dynamics. The PSID, which has followed thousands of American families since 1968, is among the only panels long enough to permit estimation of a model with multiple household-specific parameters. However, the PSID measures only food consumption.[21] While food is

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

not an ideal proxy for total consumption, it may be more likely to be time separable, as the expected utility formulation assumes.

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

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

Because new households can form from old ones, observations on different households may not be independent. In addition, the PSID uses a clustered sampling design in which many households originally lived in the same geographic areas. I adjust all bootstrap confidence intervals and two-step efficient GMM procedures to account for arbitrary correlation in the error terms across households and over time within each of the 119 original PSID primary sampling units.[22]

5.2 Threats to identification

Many previous researchers have identified ways in which either (1) income is correlated with consumption under full insurance or (2) income is uncorrelated with consumption even though insurance is incomplete or nonexistent. The first problem would render risk-sharing regressions biased; the second would render them useless. My purpose is not to revisit the debate about whether regressions of consumption on income can ever constitute a good test of insurance, nor can I solve most of the well-known problems here. However, I review these problems as a reminder of important caveats in interpreting my results.

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

Because I create a new household when the head or spouse changes, my regressions say nothing about insurance against divorce or death. In addition, a test of risk sharing using

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

data on food consumption lacks power against the alternative that consumption of other goods is not well smoothed. More generally, whatever consumption variable a researcher studies may not be the consumption that households care about. Even given data on consumption of many goods besides food, one would face difficult questions of how to treat durable goods, household production and the like.

5.3 Factor model estimates

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

When I study households' total food consumption and assume households have identical preferences, the elasticity of consumption with respect to income is 0.199 to 0.305, depending on the definition of a household. Thus a 1 percent increase in income raises consumption by 0.199 to 0.305 percent, holding aggregate shocks constant but assuming identical preferences. Allowing variation in preferences reduces the elasticity by one-third to one-half, to between 0.100 and 0.150. I reject at the 5 percent level the hypothesis that the coefficients are equal in models with and without heterogeneity.

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

5.4 GMM estimates

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

The overidentifying restrictions are frequently rejected in the models that assume common preferences, suggesting that these models are misspecified. The overidentifying restrictions are never rejected when heterogeneity in both risk and time preferences is allowed. Allowing heterogeneity in time preferences affects the coefficient on income more than allowing heterogeneity in risk preferences. Depending on the definition of a household, the adjustment for household size and whether I allow nonseparability between consumption and leisure, the elasticity of consumption with respect to income ranges from 0.195 to 0.545 when I assume common preferences. The elasticity ranges from -0.054 to 0.141 when I allow heterogeneity in time preferences, from 0.072 to 0.524 when I allow heterogeneity in risk preferences, and from -0.086 to 0.163 when I allow heterogeneity in both time and risk preferences. There is no clear pattern in how allowing nonseparability between consumption and leisure affects the results. Defining a new household if any household member changes tends to reduce the coefficient on income.

5.5 Discussion

The factor model estimates and GMM estimates both show that allowing heterogeneity tends to reduce the estimated effect of income on consumption, but the magnitudes of the estimated coefficients differ substantially between the factor model and the 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 preferences to be separable in consumption and leisure, 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. In light of the potential problems with the GMM estimates, I view the factor model estimates as more reliable and rely on them in calculating the welfare costs of imperfect insurance in the next section.

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

Besides comparing my estimates to the income elasticity of consumption, one can compare my results to two papers that test risk sharing with PSID data and assume that 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.

6 The Welfare Costs of Idiosyncratic Risk

The model of imperfect risk sharing suggests that imperfect insurance reduces welfare in three ways: insurance transfers are costly; consumption is more volatile than it would be under full insurance, which risk-averse households do not like; and households choose jobs in part to avoid risk rather than to maximize output. This section estimates the first two of these three kinds of cost and finds that both are small.

The calculations aim mainly to put the estimates of the effect of income on consumption in a meaningful context, not to fully characterize the welfare costs of risk. I therefore make several simplifying assumptions. To abstract from issues of intertemporal substitution and serial correlation of shocks, I calculate welfare costs of risk in an economy with one period. Strongly serially correlated shocks could produce larger welfare costs than I calculate, since a single bad shock could result in extremely low consumption for a very long time. However, even with serially correlated shocks, the point would remain that risk sharing looks better when we allow heterogeneous preferences. I also assume that households have CRRA preferences and that utility does not depend on leisure: $u_i(c) = c^{1-\gamma_i}/(1-\gamma_i)$. Further, I do not consider the welfare costs of aggregate risk. In a world with heterogeneous preferences, the effects of aggregate risk depend crucially on whether the right people bear it, which can be determined only given precise data on the preferences and consumption processes of the same households – data I do not have.[23]

I measure welfare costs as percentages of consumption. Imagine offering a household the choice between the real economy and a hypothetical economy where consumption varies less. What fraction of consumption would the household pay to live in the safer economy? If a household's consumption is c_{is}^A in state s in the real world A and c_{is}^B in state s in the less risky world B, then the household's willingness to pay to reduce risk is the fraction k that solves $\mathrm{E}[u_i(c_{is}^A)] = \mathrm{E}\big[u_i\big((1-k)c_{is}^B\big)\big]$.

Throughout, I assume that the mean of consumption does not change in the safer economy. In other words, households do not choose different jobs when insurance arrangements change. If improving insurance causes households to change jobs, welfare will improve further – the third welfare consequence of risk. This welfare gain can be arbitrarily small or large, depending on preferences and technology.[24]

6.1 Transactions costs

The first cost of imperfect insurance is that resources are wasted on transactions. Recall that the income coefficient g in the risk-sharing regression (3) measures relative costs and benefits of moving resources between households. Specifically, the coefficient is the mean ratio of transfer costs to risk aversion plus transfer costs, $g = \mathrm{E}[\phi_i/(\phi_i + \gamma_i)]$, where the cost of making household i's consumption c differ from its income X is $(\phi_i/2)c[\log(c/X)]^2$.

Suppose for simplicity that $\phi_i/\gamma_i = m$ for all households i, so that $g = m/(1+m)$. If the

[23]However, see Schulhofer-Wohl (2007) for estimates of the welfare costs of aggregate risk in an economy with full insurance and heterogeneous risk preferences.

[24]For zero benefits, suppose each household is productive in only one job; people choose the same jobs no matter what. For arbitrarily large benefits, note that in the model of section 2 with log-quadratic transfer costs – $h(X, c) = (c/2)[\log(c/X)]^2$ – a household with income near zero must have consumption near zero to satisfy the aggregate resource constraint. With CRRA utility and coefficient of relative risk aversion at least one, consumption near zero means utility near negative infinity. Hence we can construct jobs in which the expected value of income is large but, because there is a chance of income near zero, no one takes the jobs under partial insurance. Since no one takes these jobs under partial insurance, their existence is consistent with any observed data. Full insurance would allow all households to switch jobs and increase output by an arbitrarily large amount.

average household has log utility, we then have $m = \phi_i/\gamma_i = g/(1-g)$. The estimated income coefficients using the factor model ranged from 0.092 to 0.150. Since food is a necessity and has an income elasticity below unity, these coefficients should be adjusted upward by the elasticity of total spending to food consumption, which Blundell et al. (2005) estimate to be $1/0.88 = 1.14$ in U.S. data. This yields a range of $m = 0.117$ to $m = 0.205$ for the ratio of transactions costs to risk aversion.

Consider the cost of making a household's consumption differ from its income by 10 percent.[25] A 10 percent difference between consumption and income gives $\log(c/X) \approx 0.1$, so the cost of transfers is

$$(\phi/2)c[\log(c/x)]^2 \approx 0.005\phi c = 0.005m\gamma_i c.$$

The estimated values of $m = \phi_i/\gamma_i$ from the models with heterogeneous risk and time preferences therefore imply that, for a household with log utility ($\gamma_i = 1$), it costs between 0.06 percent and 0.1 percent of consumption to have a 10 percent difference between consumption and income. By comparison, the cost is between 0.1 percent and 0.27 percent of consumption if one uses the coefficients estimated assuming common preferences.

6.2 Uninsured idiosyncratic risk

The second cost of imperfect insurance is that consumption varies more than it would under full insurance. To measure this cost, I consider removing all consumption variation due to uninsured idiosyncratic income shocks while leaving in place variation due to aggregate shocks and holding mean consumption constant. With only a few years of data per household, as in the PSID, I cannot precisely estimate each household's income and consumption process.

[25] I choose a 10 percent difference purely for illustration. The mean of $[\log(c/X)]^2$ in the data is larger but depends on assumptions about the share of food in total consumption and – since this mean depends on the variance of measured values of $\log(c/X)$ – is likely dominated by measurement error.

I therefore use a log-normal approximation. Let consumption be $\log c_i^A \sim N(\mu_{iA}, \sigma_{iA}^2)$ under imperfect insurance and $\log c_i^B \sim N(\mu_{iB}, \sigma_{iB}^2)$ under full insurance. (Consumption still varies under full insurance due to aggregate shocks.) If mean consumption is the same in both worlds, a household with CRRA preferences is willing to pay a fraction $k \approx (\gamma_i/2)(\sigma_{iA}^2 - \sigma_{iB}^2)$ of consumption to move from imperfect insurance to full insurance.[26] Thus I can calculate the welfare gain from full insurance by calculating how much full insurance would reduce the variance of consumption.

Full insurance removes only the variation in consumption that is orthogonal to aggregate shocks. From equation (3), the within-household variance of consumption under partial insurance, holding aggregate shocks constant, is

$$\sigma_{iA}^2 - \sigma_{iB}^2 = \text{Var}_{\text{within}}[\log c_{it}^A | \log \lambda_t] = \left(\frac{\phi_i}{\gamma_i}\right)^2 \text{Var}_{\text{within}}[\log X_{it} | \log \lambda_t].$$

Hence the welfare gain from full insurance is

$$k \approx \frac{\gamma_i}{2}\left(\frac{\phi_i}{\gamma_i}\right)^2 \text{Var}_{\text{within}}[\log X_{it} | \log \lambda_t] \tag{21}$$

Although ϕ_i/γ_i may vary across households and may be correlated with the variance of income, some simple estimates of (21) are possible. As shown above, one can derive an estimate of the mean of ϕ_i/γ_i from estimates of equation (3). To estimate the mean within-household variance of log income conditional on the aggregate shock, I run household-by-household regressions of log income on the estimated aggregate shocks and a time trend and calculate the variance of the residuals. Table 7 combines the estimated variance of the residuals and the estimated mean of ϕ_i/γ_i to calculate the welfare cost in (21).

[26] Expected utility under imperfect insurance is $(1-\gamma_i)U^A = \exp[(1-\gamma_i)\mu_{iA} + (1-\gamma_i)^2\sigma_{iA}^2/2]$. Expected utility under full insurance if the household gives up a fraction k of consumption is $(1-\gamma_i)U^B(k) = (1-k)^{1-\gamma_i}\exp[(1-\gamma_i)\mu_{iB} + (1-\gamma_i)^2\sigma_{iB}^2/2]$. Solve $U^B(k) = U^A$ for k, note that holding mean consumption constant requires $\mu_{iA} - \mu_{iB} = (\sigma_{iB}^2 - \sigma_{iA}^2)/2$, and use $1 - \exp(-x) \approx x$ for x near zero to obtain the result.

Since some of the observed variation in income may be due to measurement error, I report results both assuming that all of the observed variation is true variation and assuming that only half of the observed variation is true variation.[27] The welfare gain for log utility is just 0.12 percent to 0.48 percent of consumption before accounting for measurement error, and half that after accounting for measurement error. Even if the coefficient of relative risk aversion is 10, the welfare gain from full insurance after accounting for measurement error is at most 2.4 percent of consumption. Further, accounting for heterogeneity is important: In the factor model results, estimates of ϕ/γ ignoring heterogeneity are 1.6 to 2 times as large as estimates that allow heterogeneity, so ignoring heterogeneity would double to quadruple the apparent welfare gains from full insurance.

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 and can be used to calculate welfare losses from insurance failures. This paper highlights the importance of accounting carefully for preferences when measuring how income affects consumption.

I use a model of imperfect risk sharing to show that if households have different risk and time preferences, income shocks will be correlated with preferences, and a risk-sharing regression that assumes identical preferences will find too large an effect of idiosyncratic income shocks on consumption. Empirical results confirm these predictions. In the HRS, earnings are more variable and more correlated with aggregate shocks among workers who express greater risk tolerance. In the PSID, allowing heterogeneity in preferences substantially re-

[27]Bound et al. (1994) find that the cross-sectional share of measurement error in earnings variance in the PSID is 0.15 to 0.3. The share of measurement error in within-household variation orthogonal to aggregate shocks may be higher since removing fixed effects, trends and aggregate shocks removes true variance.

duces the estimated effect of income on consumption. After accounting for heterogeneity, I find that welfare losses from imperfect insurance of idiosyncratic risk are quite small: on the order of 0.1 to 0.5 percent of consumption for a household with log utility.

The results suggest several directions for further research. First, my model and the HRS data show that when insurance is imperfect, people 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. The welfare gain from changes in sorting could be calculated by using panel data on occupation choice, income and preferences to estimate the relationship between preferences, productivity in various occupations and the time series properties of individual income. Second, although I have demonstrated that – in broad terms – more risk-averse people bear less aggregate risk, I have not tested whether aggregate risk is allocated in precisely the shares required for Pareto efficiency. Detailed data on risk preferences and consumption processes might allow one to measure how well people share aggregate risk. Finally, I study a static model. The relationship between preferences, income processes and risk sharing could also fruitfully be investigated in a dynamic context.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Table 7: Welfare gain from eliminating uninsured idiosyncratic income variation.

	New household definition	
	New head/spouse	Any change
Range of estimates of mean ratio of transfer costs to risk aversion (ϕ/γ)[a]	0.192 to 0.205	0.117 to 0.177
Within-household variance of log(income) conditional on aggregate shock[b]	0.23	0.17
Welfare gain from eliminating uninsured idiosyncratic income variation, assuming no measurement error[c]	$0.42\gamma_i\%$ to $0.48\gamma_i\%$	$0.12\gamma_i\%$ to $0.27\gamma_i\%$
Welfare gain adjusted for measurement error[c,d]	$0.21\gamma_i\%$ to $0.24\gamma_i\%$	$0.06\gamma_i\%$ to $0.13\gamma_i\%$

[a]Calculated from factor model estimates in table 4. [b]Variance of residuals from household-by-household regressions of log(income) on time trend and estimated aggregate shocks; to number of decimal places shown, adjusting income to adult-equivalent units did not change results. [c]As percentage of consumption, given coefficient of relative risk aversion γ_i; from equation (21). [d]Assumes 50 percent of observed within-household variance of income conditional on aggregate shock is due to measurement error.

A PSID consumption, income and leisure data

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

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

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

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

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

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

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

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

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

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

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

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

Table A.2: PSID adult equivalence scale.

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

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

References

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

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

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

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

Blundell, Richard, Luigi Pistaferri and Ian Preston, 2006, "Consumption Inequality and Partial Insurance," Manuscript, University College London and Stanford University.

Bonin, Holger, Thomas Dohmen, Armin Falk, David Huffman and Uwe Sunde, 2006, "Cross-sectional Earnings Risk and Occupational Sorting: The Role of Risk Attitudes," IZA discussion paper No. 1930.

Bound, John, Charles Brown, Greg J. Duncan and Willard L. Rodgers, 1994, "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data," *Journal of Labor Economics* 12(3), 345–368.

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

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

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

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

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

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

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

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

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

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

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

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

Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante, 2007, "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," Manuscript, New York University.

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

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

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

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

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

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

Mazzocco, Maurizio and Shiv Saini, 2007, "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," Manuscript, UCLA.

Munshi, Kaivan and Mark Rosenzweig, 2005, "Why is Mobility in India So Low? Social Insurance, Inequality, and Growth," Working paper No. 097, Bureau for Research in Economic Analysis of Development.

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

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

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

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

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

Schulhofer-Wohl, Sam, 2007, "Heterogeneous Risk Preferences and the Welfare Cost of Business Cycles," Manuscript, Princeton University.

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

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

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

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

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

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

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

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

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

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