

Heterogeneity and Risk Sharing in Thai Villages

Pierre-André Chiappori, Krislert Samphantharak,
Sam Schulhofer-Wohl and Robert M. Townsend*

September 27, 2010

Abstract

We use two distinct methods, drawn from the theories of efficient risk sharing and optimal portfolio choice, to measure heterogeneity in risk aversion among households in Thailand. Our dataset is an unusually long monthly panel that includes information on kinship networks. We demonstrate that among households who have relatives living in the same village, the hypothesis of full insurance cannot be rejected, suggesting that gifts and insurance transfers among relatives provide something close to a complete-markets consumption allocation. We then show that there is substantial heterogeneity in risk preferences according to both of our methods. Estimates from the two methods are positively correlated in most villages. However, we find little relationship between preferences and household wealth or demographic variables. In a computational experiment, we show that the heterogeneity in preferences matters for policy: Although the average household would benefit from eliminating aggregate village-level risk, some relatively risk-tolerant households would experience substantial welfare losses if aggregate risk were eliminated.

*Chiappori: Columbia University; pc2167@columbia.edu. Samphantharak: University of California, San Diego; krislert@ucsd.edu. Schulhofer-Wohl: Federal Reserve Bank of Minneapolis; wohls@minneapolisfed.org. Townsend: Massachusetts Institute of Technology; rtownsen@mit.edu. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

We use two distinct methods to measure heterogeneity in risk aversion across households running running farm and non-farm enterprises in a developing country. One method, drawn from the literature on risk sharing, identifies a household's risk aversion by examining how much its consumption co-moves with aggregate consumption. The intuition – which dates to Wilson (1968) – is that efficient risk sharing allocates more risk to less-risk-averse households, so a household whose consumption strongly co-moves with the aggregate must be relatively less risk averse. The second method comes from portfolio choice theory and uses measures of volatility in the household's consumption growth and in the return on the household's capital assets. The intuition behind this second method – famously exploited with aggregate data to identify the preferences of the representative agent by Mehra and Prescott (1985) – is that the more risk-averse a household is, the safer a portfolio it will choose and the smoother its consumption will be.

The data we analyze are an unusually long monthly panel of households in villages in Thailand, including information on the existence of kinship groups living in the same village. For households who have kin living in the village, we find evidence of nearly complete risk sharing. Indeed, we cannot reject the null hypothesis of full risk sharing, even though we use a powerful test that is biased toward rejection if preferences are heterogeneous. We do reject full insurance among households that have no kin in the village, suggesting strongly that gifts and insurance transfers among family-related households are providing something close to a complete markets allocation. Evidently, informal village institutions in Thailand provide risk sharing similar to what is implicitly assumed when researchers estimate representative-agent models using data from the New York financial markets. Our findings on networks echo the result of Samphantharak and Townsend (2010a, chapter 6) that membership in a kinship network reduces the effect of liquidity constraints on households' financing of fixed

assets and the result of Kinnan and Townsend (2010) that kinship networks are important for households' access to financing and ability to smooth consumption. Since the theory appears to fit best the households with relatives in the village, we restrict the sample to these households for the remainder of the analysis. This restriction is similar to, for example, Vissing-Jørgensen's (2002) method of estimating the elasticity of intertemporal substitution on a restricted sample of households that participate in the stock and bond markets.

Using the sample of households with kin in the village, we establish that there is substantial heterogeneity in risk tolerance according to both of our methods. Further, our two distinct measures of risk tolerance are positively correlated with each other in most villages, giving us some confidence in their validity. However, the correlations are weak, suggesting that each set of estimates may contain substantial amounts of measurement error. Finally, we find that neither of the two measures of risk tolerance is significantly correlated with demographic variables or household wealth. The finding of no correlation between preferences and wealth is consistent with the complete markets hypothesis and, since we are measuring relative risk tolerance, consistent with the finding of Chiappori and Paiella (2008) that the correlation between wealth and relative risk aversion – as estimated from portfolio structures in Italian panel data – is very weak. In addition, the lack of correlation between preferences and demographics is reminiscent of the “massive unexplained heterogeneity” in Italian households' preferences reported by Guiso and Paiella (2008).

Heterogeneity in risk tolerance matters for policy. To make this point, we conduct a hypothetical experiment in which we estimate the welfare gains and losses that would result from eliminating all aggregate, village-level risk. If all households were equally risk averse, all households would benefit from eliminating aggregate risk. Heterogeneity makes the situation more interesting. As demonstrated by Schulhofer-Wohl (2008) using U.S. data, heterogeneity in preferences implies that some sufficiently risk-tolerant households would experience welfare *losses* from eliminating aggregate risk, because these households effectively sell insurance

against aggregate risk to their more risk-averse neighbors and collect risk premia for doing so. In the Thai data, we find that households live with a great deal of aggregate risk – figure 1 shows the volatility of aggregate consumption in each village, with a monthly standard deviation of about 0.14% – and that the average household would be willing to pay to avoid this risk. However, not all households would be willing to pay. In fact, if aggregate risk were eliminated, some relatively risk-tolerant households would suffer welfare losses equivalent to several percent of mean consumption. Heterogeneity in the population is, therefore, substantial.

The paper proceeds as follows. In section 2, we lay out the theory underlying our two methods for estimating preferences. In section 3, we describe the Thai data. Section 4 presents the empirical results, and section 5 concludes. The appendix contains some mathematical derivations.

2 Theory

In this section, we derive two methods to estimate households' risk preferences, one based on measurements of risk sharing among households, and the other based on households' choices of asset portfolios. For both methods, we assume that there is one consumption good, c . We assume that households maximize time-separable discounted expected utility with constant relative risk aversion. We allow each household to have its own rate of time preference and its own coefficient of relative risk aversion. Because we will work with monthly data, we need to distinguish consumption fluctuations that are due to risk from consumption fluctuations that are due to seasonal preferences. Therefore, we also allow each household to have month-specific preferences. That is, household i 's preferences over consumption sequences $\{c_{it}^*(s^t)\}$,

where s^t is the history of states up to date t , are represented by

$$E_0 \left[\sum_{t=0}^T \beta_i^t \xi_{i,m(t)} \frac{[c_{it}^*(s^t)]^{1-\gamma_i}}{1-\gamma_i} \right] \quad (1)$$

where β_i is the household's rate of time preference, γ_i is the household's coefficient of relative risk aversion, $\xi_{i,m}$ is the household's relative preference for consuming in month $m \in \{\text{Jan, Feb}, \dots, \text{Dec}\}$, and $m(t)$ is the month corresponding to date t . We assume $\xi_{i,m}$ is non-stochastic.

For the risk-sharing method, we also assume that consumption is measured with error: We assume that we observe not true consumption c_{it}^* but instead $c_{it} = c_{it}^* \exp(\epsilon_{it})$. Our assumptions on the measurement error ϵ_{it} are relatively weak. We assume that it is mean independent of the date t and of village aggregate consumption $C_{jt}(s^t)$ (defined more precisely below), has mean zero for each household, and is uncorrelated across households:

$$\begin{aligned} E[\epsilon_{it}|i, t, C_{jt}(s^t)] &= 0 \\ E[\epsilon_{it}\epsilon_{i',t'}] &= 0 \quad \forall i \neq i', \forall t, t' \end{aligned} \quad (2)$$

Notice in particular that we are not assuming anything about homoskedasticity or serial correlation of the measurement errors.

After deriving the two methods, we show that we can also use the data to estimate the welfare cost of aggregate risk in the villages in our data, as a function of households' risk preferences.

2.1 Risk-Sharing Method

Let $C_{jt}(s^t)$ be the aggregate consumption available in village j at date t after history s^t . (We take no stand on storage or inter-village risk sharing. If storage is possible, $C_{jt}(s^t)$ is aggregate consumption net of any aggregate storage. If risk is shared between villages, $C_{jt}(s^t)$

is aggregate consumption in village j after any transfers to or from other villages.) Then, following Diamond (1968) and Wilson (1967), any Pareto-efficient consumption allocation satisfies

$$\ln c_{it}^*(s^t) = \frac{\ln \alpha_i}{\gamma_i} + \frac{\ln \beta_i}{\gamma_i} t + \frac{\ln \xi_{i,m(t)}}{\gamma_i} + \frac{1}{\gamma_i} [-\ln \lambda_{j(i),t}(s^t)], \quad (3)$$

where $j(i)$ is household i 's village, α_i is a non-negative Pareto weight, and $\lambda_{jt}(s^t)$ is the Lagrange multiplier on village j 's aggregate resource constraint $\sum_i c_{it}^*(s^t) = C_{jt}(s^t)$ at date t after history s^t . The multiplier $\lambda_{jt}(s^t)$ is a function only of aggregate resources $C_{jt}(s^t)$; for a given village j , any two histories with the same aggregate resources at a particular date will have the same λ at that date.

The first term in (3) is a household-specific fixed effect; some households simply are better off than others and, on average, consume more. The second term is a household-specific trend. Formally, these trends depend on the household's rate of time preference β_i ; informally, the household-specific trends could stand for anything that makes some households want to have different trends in consumption than other households, such as life-cycle considerations. The third term reflects differences in the seasonality of households' preferences. The fourth term shows how consumption depends on aggregate shocks λ_{jt} : Consumption moves more with aggregate shocks for less risk-averse households.

Equation (3) reflects Wilson's (1968) result that doubling every household's coefficient of relative risk aversion will not change the set of Pareto-efficient allocations: The consumption allocation in (3) does not change if, for any non-zero constant m_j specific to village j , we replaced γ_i with $m_j \gamma_i$, replaced $\lambda_{j(i),t}$ with $m_j \lambda_{j(i),t}$, and adjusted α_i , β_i and $\xi_{i,m(t)}$ appropriately. In consequence, when we use a method based on (3) to estimate preferences, we will be able to identify risk preferences only up to scale within each village.

Since consumption is measured with error, an equation for observed consumption under

efficient risk sharing is

$$\ln c_{it} = \frac{\ln \alpha_i}{\gamma_i} + \frac{\ln \beta_i}{\gamma_i} t + \frac{\ln \xi_{i,m(t)}}{\gamma_i} + \frac{1}{\gamma_i} (-\ln \lambda_{j(i),t}) + \epsilon_{it}, \quad (4)$$

where we have suppressed the dependence on the history s^t for convenience.

We do two things. First, we test for efficient risk sharing. Then, under the maintained hypothesis of efficient risk sharing, we use (4) to estimate households' risk preferences.

2.1.1 Test of efficient risk sharing

The standard test for efficient risk sharing in the literature (e.g., Cochrane, 1991; Mace, 1991; Townsend, 1994) can be described as follows. If agents share risk efficiently, then the individual consumption of agent i should depend only on aggregate shocks, as described by equation (4), but not on i 's idiosyncratic income shocks. This result suggests running a regression like:

$$\ln c_{it} = \frac{\ln \alpha_i}{\gamma_i} + \frac{\ln \beta_i}{\gamma_i} t + \frac{\ln \xi_{i,m(t)}}{\gamma_i} + \frac{1}{\gamma_i} (-\ln \lambda_{j(i),t}) + b_j \ln income_{it} + \epsilon_{it}, \quad (5)$$

where household i lives in village $j(i)$. The test, now, would be whether the coefficient b_j is significantly different from zero. Efficient risk sharing would imply $b_j = 0$, whereas any deviation from perfect risk sharing would result in $b_j > 0$.

In practice, however, we use a slightly different test; that is, we run the regression

$$\ln c_{it} = a_i + d_{j(i),t} + b_j \ln income_{it} + u_{it}, \quad (6)$$

where d_{jt} represents the aggregate shock in village j at date t . (Note that, for the sake of precision, we write d_{jt} rather than λ_{jt} in this equation; indeed, if $b_j \neq 0$, our structural model in the previous section is incorrect and d_{jt} is some kind of aggregate shock but need

not be the Lagrange multiplier λ_{jt} .)

The key difference between (5) and (6) is that (6) ignores heterogeneity in preferences and absorbs the household-specific trends and seasonality into the aggregate shocks d_{jt} . The motivation for this variant is that allowing heterogeneity in risk preferences, time preferences or seasonal preferences would make the test *less powerful*. Indeed, several authors have showed that rejecting full insurance is harder in the presence of heterogeneity; therefore, by ignoring heterogeneity, we bias our test toward rejecting full insurance.¹ Since our goal is to show that the evidence against full insurance is weak, we want to strengthen our case by using as powerful a test as we can – hence the choice not to allow heterogeneity. Below, we will show that our test is sufficiently powerful to reject full insurance in some relatively small samples of households that are not in kinship networks. However, in larger samples of households in kinship networks, we do *not* reject full insurance, even though the test is more powerful in larger samples.

2.1.2 Estimating preferences

If there is full insurance, then the data must satisfy (4), and we can use this equation to estimate each household’s risk preferences γ_i . (In principle, we can also estimate each household’s time preferences β_i , but that is not our goal here – primarily because β_i is difficult to interpret since it represents a combination of pure time preference and life-cycle motives.)

The intuition for how we estimate risk preferences is that under full insurance, a household whose consumption moves more with aggregate shocks must be less risk averse. Further, under full insurance, the only reason two households’ consumptions can move together is

¹The intuition is that whenever more risk-tolerant households’ incomes are (weakly) more correlated with the aggregate shock – as seems natural – then the common-preferences test is (weakly) biased against the null of full insurance while the heterogeneous-preferences test is not biased. This, plus the fact that the common-preferences test estimates fewer parameters and hence has more residual degrees of freedom and more power against alternatives, implies that the common-preferences test will reject the null (weakly) more often than the heterogeneous-preferences test if the null is true, and strictly more often if the null is false. The reader is referred to Mazzocco and Saini (2009) and Schulhofer-Wohl (2010) for a precise discussion of these issues.

that both of their consumptions are co-moving with aggregate shocks. Thus, if two households' consumptions are strongly correlated, they both must have consumption that moves strongly with the aggregate shock; they must both be relatively risk tolerant. Similarly, if two households' consumptions are not strongly correlated, at least one must have consumption that does not move strongly with the aggregate shock; at least one must be very risk averse. In consequence, we can identify relatively more and less risk-averse households by looking at the pairwise correlations of their consumption.

Our method uses only the data on households whose consumption is observed in every time period. Suppose that there are J villages, and that for each village j , we have data on N_j households observed in T time periods. These need not be all households in the village for all time periods in which the village has existed.

Let $\{\nu_{it}\}_{t=1}^T$ be the residuals from linearly projecting the time series of log consumption for household i on a household-specific intercept, time trend and month dummies. Log consumption is the left-hand side of (4). Thus, since (4) holds and projection is a linear operator, the log consumption residuals ν_{it} must equal the total of the residuals we would obtain from separately projecting each term on the right-hand side of (4) on a household-specific intercept, time trend and month dummies. There are no residuals from projecting the first three terms on the right-hand side since these terms are equal to a household-specific intercept, trend and month dummies. Thus ν_{it} must equal the total of the residuals from projecting $(-\ln \lambda_{jt})$ and ϵ_{it} . Specifically, suppose that we could observe the Lagrange multipliers λ_{jt} , and let ℓ_{jt} be the residual we would obtain if we hypothetically projected $(-\ln \lambda_{jt})$ on an intercept, a time trend and month dummies.² Also suppose that we could observe the measurement errors ϵ_{it} , and let $\tilde{\epsilon}_{it}$ be the residual we would obtain if we hypothetically projected the time series of ϵ_{it} on a household-specific intercept, time trend and month dummies. Then

²The results of this projection will be the same for all households in a village since the panel is balanced and λ_{jt} is the same for all households in the village.

equation (4) implies

$$\nu_{it} = \frac{1}{\gamma_i} \ell_{j(i),t} + \tilde{\epsilon}_{it}. \quad (7)$$

Since ϵ_{it} is uncorrelated across households, (7) implies that for any two households i and i' in the same village j ,

$$\text{E}[\nu_{it}\nu_{i',t}] = \frac{1}{\gamma_i\gamma_j} \text{E}[\ell_{jt}^2], \quad i \neq i' \quad (8)$$

As discussed above, risk aversion is identified only up to scale within each village; equation (7) would not change if, for any non-zero constant m_j specific to village j , we replaced γ_i with $m_j\gamma_i$ and ℓ_j with $m_j\ell_j$. Since the scale m_j is unidentified, we can normalize $\text{E}[\ell_{jt}^2] = 1$. With this normalization, (8) reduces to

$$\text{E}[\nu_{it}\nu_{i',t}] = \frac{1}{\gamma_i\gamma_j}, \quad i \neq i' \quad (9)$$

Equation (9) applies to each pair of distinct households, so the equation gives us $N_j(N_j-1)/2$ moment conditions in N_j unknowns (the risk aversion coefficients $\{\gamma_i\}_{i=1}^{N_j}$). In principle, we could use these moment conditions to estimate the risk aversion coefficients by the Generalized Method of Moments. However, we would then have many more moment conditions than months of data – for example, in a village with $N_j = 30$ households, which is typical, we would have 435 moment conditions but only 84 months of data – and GMM can perform poorly when there are many moment conditions (Han and Phillips, 2006). We therefore collapse (9) to one moment condition per household by summing over the other households $i' \neq i$, reducing our moment conditions to:

$$\sum_{i' \neq i} \text{E}[\nu_{it}\nu_{i',t}] = \frac{1}{\gamma_i} \sum_{i' \neq i} \frac{1}{\gamma_{i'}} \quad (10)$$

Equation (10) gives us N_j moment conditions in N_j unknowns, so we have a just-identified system. We use these just-identified moment conditions to estimate the parameters by

GMM.³ We can also use GMM to test the null hypothesis that all households in village j have identical preferences, by imposing the restriction that $\gamma_1 = \gamma_2 = \dots = \gamma_{N_j}$ and then testing the $N_j - 1$ overidentifying restrictions with the usual Hansen (1982) χ^2 statistic.

In our GMM estimation, we must impose a sign normalization on the estimated coefficients of relative risk aversion since the moment conditions do not change if we multiply each γ_i by -1 . Since the true coefficients of relative risk aversion must be positive, we impose the normalization that $\sum_{i=1}^{N_j} \gamma_i > 0$.

2.2 Portfolio-Choice Method

We can also use a simple portfolio choice and asset pricing model (Breedon, 1979; Lucas, 1978; Rubinstein, 1976) to recover households' preferences from their asset holdings.⁴

Suppose that household i optimally allocates its assets across a portfolio of assets k that have stochastic gross returns R_{t+1}^k . The household's Euler equation requires that, for every asset k which the household chooses to own,

$$1 = \frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \frac{\beta_i}{u'_i(c_{it}^*)} \text{E}_t[u'_i(c_{i,t+1}^*) R_{t+1}^k]. \quad (11)$$

Taking the unconditional expectation of both sides, applying the law of iterated expectations

³An alternative approach would be to observe that (4) is essentially a factor model – the Lagrange multiplier $\ln \lambda_{jt}$ is an unobserved factor, and risk tolerance $1/\gamma_i$ is the factor loading that specifies how the factor impacts household i – and to estimate the equation by standard factor analysis methods. With a small number of households, as here, the identifying assumption for factor analysis would be that the measurement errors ϵ_{it} are uncorrelated over time and across households and that their variance is constant across households at each date t . Examination of the residuals from the equation suggests, however, that the variance differs across households. Thus we were not confident in the factor analysis assumptions and did not pursue that approach.

⁴The first application of this idea was by Mehra and Prescott (1985), who used asset pricing equations to compute the risk aversion of a representative agent from aggregate consumption data. Mehra and Prescott (1985) concluded that, because the average return on equities in U.S. data is quite high compared with the variances of consumption growth and the return on equities, the representative agent would have to be extremely risk averse to rationalize the data. Below, we find more reasonable values for risk aversion because the variances of returns and consumption growth are higher in our data.

and rearranging terms, the Euler equation requires

$$\begin{aligned}
1 &= \mathbb{E} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} R_{t+1}^k \right] \\
&= \mathbb{E} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right] \mathbb{E}[R_{t+1}^k] + \text{Cov} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k \right]. \quad (12)
\end{aligned}$$

In particular, (12) must hold for any risk-free asset that the household chooses to hold. The households in our sample typically hold inventories of their products; as long as storage is riskless and relative prices are constant, the value of inventory will move one-for-one with inflation, and inventory will have a risk-free gross return of 1. Therefore,

$$1 = \mathbb{E} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right]. \quad (13)$$

We show in appendix A that if we assume (a) households have CRRA preferences, (b) households choose portfolios on the mean-variance frontier, (c) seasonally adjusted consumption growth has a log-normal distribution and (d) the variance of seasonally adjusted consumption growth is “small,” then equations (12) and (13) imply the following equation for risk aversion:

$$\gamma_i = \frac{1}{\sqrt{\text{Var}[\Delta \ln x_{it}^*]}} \left| \frac{\mathbb{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right|, \quad (14)$$

where $R_{i,t+1}^P$ is the gross return on the household’s portfolio, $x_{it}^* = \xi_{i,m(t)}^{-1/\gamma_i} c_{it}^*$ is seasonally adjusted consumption, and all variances and expectations are household-specific.

We cannot directly use (14) to estimate γ_i because consumption may be measured with error and because we observe total consumption, not seasonally adjusted consumption. To estimate $\text{Var}[\Delta \ln x_{it}^*]$, notice that if $(x_{i,t+1}^*/x_{it}^*)$ has a log-normal distribution, then

$$\Delta \ln x_{it}^* = \mu + e_{it} \quad (15)$$

where e_{it} has a normal distribution, and $\text{Var}[\Delta \ln x_{it}^*] = \text{Var}[e_{it}]$. Since $\ln x_{it}^* = -\frac{1}{\gamma_i} \ln \xi_{i,m(t)} + \ln c_{it}^*$ and $\ln c_{it}^* = \ln c_{it} - \ln \epsilon_{it}$, we have

$$\begin{aligned} \Delta \ln c_{it} &= \Delta \ln x_{it}^* + \frac{1}{\gamma_i} \Delta \ln \xi_{i,m(t)} + \Delta \ln \epsilon_{it} \\ &= \frac{1}{\gamma_i} \Delta \ln \xi_{i,m(t)} + \mu + e_{it} + \epsilon_{i,t+1} - \epsilon_{it}. \end{aligned} \tag{16}$$

Let $\hat{V}_{e,i}$ be the variance of residuals from a household-specific regression of $\Delta \ln c_{it}$ on month dummy variables. According to (16), if we had an infinitely long time series for the household, these residuals would equal $e_{it} + \epsilon_{i,t+1} - \epsilon_{it}$ and their variance would be $\text{Var}[e_{it}|i] + 2\text{Var}[\epsilon_{it}|i] - 2\text{Cov}[\epsilon_{i,t+1}, \epsilon_{i,t}|i]$. Therefore, $\hat{V}_{e,i}$ converges in probability to $\text{Var}[e_{it}|i] + 2\text{Var}[\epsilon_{it}|i] - 2\text{Cov}[\epsilon_{i,t+1}, \epsilon_{i,t}|i]$ as $T \rightarrow \infty$. If we had estimates of the household-specific variance and serial correlation of measurement error, we could use them to adjust $\hat{V}_{e,i}$ and obtain an estimate of $\text{Var}[\Delta \ln x_{it}^*] (= \text{Var}[e_{it}])$. Since we do not have a good way to estimate the variance and serial correlation of measurement error, however, we make no adjustment and use $\hat{V}_{e,i}$ as our estimate of the variance of seasonally adjusted consumption. Likewise, we make no adjustment for measurement error in calculating the variance of returns.

Because we are not accounting for measurement error, our estimates of both of the variances in the denominator of (14) will be biased upward. (The estimated variances include both the true variances and the variance of measurement error, so the estimates are higher than the true variances.) Thus our estimates of households' risk aversion coefficients γ_i will be biased downward. In examining variation in the estimated γ_i across households, we are implicitly assuming that the bias due to measurement error is the same for all households.

A further problem in using (14) to estimate γ_i is that although the time series average of a household's actual investment returns $R_{i,t+1}^P$ will converge in a sufficiently long sample to the household's expected return $E[R_{i,t+1}]$, the time series average may differ substantially from the expected return in our finite sample. Thus, for some households, we may estimate

a negative return on assets even though no household would rationally choose assets with a negative expected return. If the estimated return on assets is negative, we will estimate $\gamma_i < 0$, which does not make sense. Therefore, we calculate our estimate of γ_i only for those households that have positive estimated return on assets.

We test for heterogeneity in preferences under the portfolio-choice approach as follows. Let $\hat{\gamma}_i^{PC}$ be the estimate of household i 's risk aversion obtained by using finite-sample means and variances in (14). Let $\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})$ be the associated standard error of this estimate. Let $\bar{\gamma}_j$ be the mean risk aversion of the observed households in village j , and let $\hat{\gamma}_j$ be the estimate of this mean obtained by averaging the estimates $\hat{\gamma}_i^{PC}$ in village j . (Because we have defined $\bar{\gamma}_j$ as the mean for the observed households, it differs from $\hat{\gamma}_j$ only because of estimation error in $\hat{\gamma}_i^{PC}$; there is no discrepancy arising from using data on a finite number of households in the village. It follows that $\hat{\gamma}_j$ converges in probability to $\bar{\gamma}_j$ as the number of time periods goes to infinity, which will be important in the analysis that follows.) Under the null hypothesis that all households in village j have the same risk preferences, $\gamma_i = \bar{\gamma}_j$, we have that

$$\frac{\hat{\gamma}_i^{PC} - \bar{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \xrightarrow{d} N(0, 1) \quad (17)$$

where the convergence in distribution is as the number of time periods goes to infinity. Assume for now that the estimation errors $\hat{\gamma}_i^{PC} - \gamma_i$ are independent across households.

Then

$$\sum_{i \in j} \left(\frac{\hat{\gamma}_i^{PC} - \bar{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 \xrightarrow{d} \chi^2(0, N_j). \quad (18)$$

We cannot calculate the χ^2 statistic in (18) because we do not observe $\bar{\gamma}_j$ but only the estimate $\hat{\gamma}_j$. Observe that

$$\sum_{i \in j} \left(\frac{\hat{\gamma}_i^{PC} - \hat{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 = \sum_{i \in j} \left(\frac{\hat{\gamma}_i^{PC} - \bar{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 + \sum_{i \in j} \left(\frac{\bar{\gamma}_j - \hat{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 + 2 \sum_{i \in j} \frac{(\hat{\gamma}_i^{PC} - \bar{\gamma}_j)(\bar{\gamma}_j - \hat{\gamma}_j)}{[\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})]^2} \quad (19)$$

Since $\hat{\gamma}_j \xrightarrow{P} \bar{\gamma}_j$, the second and third terms on the right-hand side of (19) converge in probability to zero. Therefore,

$$\sum_{i \in j} \left(\frac{\hat{\gamma}_i^{PC} - \hat{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 \xrightarrow{d} \chi^2(0, N_j). \quad (20)$$

We use (20) to test the null hypothesis of no heterogeneity in preferences within each village. We obtain $\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})$ by bootstrapping. To account for possible serial correlation in consumption growth and return on assets, we use a block bootstrap and draw blocks of 12 months of data with replacement from the original sample, then recalculate $\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})$ in each bootstrap sample; $\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})$ is the standard deviation of the bootstrap estimates obtained for household i . Note that the test can also be implemented using risk tolerance instead of risk aversion by substituting $1/\hat{\gamma}_i^{PC}$ for $\hat{\gamma}_i^{PC}$, $\widehat{1/\gamma_j}$ for $\hat{\gamma}_j$ and $\widehat{\text{s.e.}}(1/\hat{\gamma}_i^{PC})$ for $\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})$ in (20). There is no particular reason to prefer one of these tests over the other, so we perform both tests.

If the estimation errors $\hat{\gamma}_i^{PC} - \gamma_i$ are correlated across households – for example because of common shocks to consumption or returns – the above analysis is not precisely correct. We think it would be difficult to account for possible correlation in the estimation errors without a detailed statistical model of asset returns, and even then, inference would be conditional on assuming the model was correct. However, we note that correlated estimation errors would make the estimated preferences $\hat{\gamma}_i^{PC}$ similar across households in finite sample even if there is heterogeneity. Therefore, correlated estimation errors would reduce the power of our test. If we reject common preferences while assuming uncorrelated estimation errors, then we can be confident the rejection would be even stronger if we accounted for correlation in the estimation errors.

2.3 Relationship between risk-sharing and portfolio-choice approaches

The risk-sharing and portfolio-choice methods differ in two important respects. First, the risk-sharing approach identifies risk preferences only up to scale, and the scale is village-specific. Thus, we cannot use the risk-sharing method to determine whether average risk aversion differs across villages. The portfolio-choice method identifies risk preferences exactly, not just up to scale, so we can use it to determine whether the average household is more risk averse in some villages than others.

Second, the portfolio-choice method differs from the risk-sharing approach both in the form of the equation estimated and in the data used. The portfolio-choice method uses the relationship between an individual household's consumption risk and the household's asset risk and returns to find the household's risk preferences, on the assumption that the household has chosen its portfolio optimally. The risk-sharing method ignores asset returns and uses the correlation between each household's consumption and aggregate consumption to find the risk preferences of all households at once. Because the two methods differ, each serves as a check on the other; if the two methods give similar results, we can have more confidence that our estimates accurately reflect households' actual preferences.

We note that there is no contradiction in assuming full insurance for the risk-sharing estimation method but using each household's idiosyncratic asset returns to estimate preferences in the portfolio-choice method. Even if households are fully insured against idiosyncratic shocks to asset returns, the Euler equation (11) must hold – Samphantharak and Townsend (2010b) show that it is the first-order condition in a social planner's problem – and, therefore, the portfolio-choice method remains valid.

2.4 The welfare cost of aggregate risk

We follow the method of Schulhofer-Wohl (2008) to estimate the welfare cost of aggregate risk. The basic idea, following Lucas (1987), is to calculate a household's expected utility from a risky consumption stream and compare it to the amount of certain consumption that would yield the same utility.

In essence, we will compare three economies. Economy 1 is the real economy; the aggregate endowment in it is risky. Economy 2 is a hypothetical economy in which the aggregate endowment is constant and equal to the expected aggregate endowment from economy 1. Some households would be better off in economy 2 than economy 1, while others are worse off, depending on their risk aversion: In economy 1, a nearly risk-neutral household can sell insurance against aggregate risk to more risk-averse households, and this nearly risk-neutral household would be worse off if it lived in economy 2 and had no opportunity to sell insurance. We would like to estimate how much better off or worse off households would be in economy 2. To do so, we introduce economy 3, which has a constant aggregate endowment equal to $(1 - k)$ times the aggregate endowment in economy 2. For each household, we find the value of k such that the household would be indifferent between living in economy 1 and living in economy 3. If $k > 0$, then the household is indifferent between the real economy 1 and a hypothetical economy where consumption is certain but smaller by the fraction k ; thus, the household is willing to give up a fraction k of its consumption to eliminate aggregate risk. If $k < 0$, aggregate risk gives the household a welfare gain equal to a fraction k of consumption.

We briefly outline the method here and refer interested readers to Schulhofer-Wohl (2008) for details.

We assume the world consists of a sequence of one-period economies indexed by date t . (By considering one-period economies, we avoid the problem that households with different risk preferences also have different preferences for intertemporal substitution and thus will

make intertemporal trades even in the absence of aggregate risk. The assumption of a one-period economy means we are treating shocks as serially uncorrelated. We think this assumption is reasonable in the context of rural villages where many shocks are related to weather.) Each economy can be in one of several states s , each with probability π_s . The states and their probabilities are the same for all dates t , and households know the probabilities. Before the state is known, the households trade a complete set of contingent claims.

We assume aggregate income in economy t in state s is $g_t m_s$, where g_t is a non-random sequence and m_s represents the shock in state s . We normalize the shocks such that $\sum_s \pi_s m_s = 1$, i.e., the expected value of aggregate income in economy t is g_t . There is no storage (or, if there is storage, “aggregate income” refers to aggregate income net of aggregate storage).

Each household is described by a coefficient of relative risk aversion γ_i and an endowment share w_i : Household i 's endowment in economy t in state s is $w_i g_t m_s$, so there is only aggregate risk and no idiosyncratic risk. We assume the joint distribution of endowment shares and risk preferences is the same at each date.⁵

Because markets are complete, the welfare theorems apply, and the consumption allocation will be the same as we derived for the risk-sharing method. One can use the allocation to derive household i 's expected utility in economy t before the state is realized. Let U_{it}^* denote this expected utility. (This is expected utility in economy 1.) Now suppose the household gave up a fraction k of its endowment but eliminated all aggregate risk, receiving consumption equal to $w_i(1 - k)g_t$ in every state in economy t . Let $\hat{U}_{it}(k)$ be the utility of a household that gave up a fraction k of its endowment but eliminated all aggregate risk. (This is expected utility in economy 3.) The welfare cost of aggregate risk, expressed as a

⁵Since the economy lasts only one period, we do not need to consider heterogeneity in households' discount factors or in their seasonal preferences as in (1).

fraction of consumption, is value of k that solves

$$\hat{U}_{it}(k) = U_{it}^*. \quad (21)$$

Schulhofer-Wohl (2008) shows that the welfare cost depends only on the household's risk aversion γ_i , not on its endowment share or the size of the economy g_t , and can be written

$$k(\gamma_i) = 1 - \left(\sum_s \pi_s (p_s^*)^{-(1-\gamma_i)/\gamma_i} \right)^{\gamma_i/(1-\gamma_i)} \quad (22)$$

where $\pi_s p_s^*$ is the equilibrium price of a claim to one unit of consumption in state s and where the prices are normalized such that $\sum_s \pi_s p_s^* m_s = 1$. It is worth noting that for γ_i sufficiently close to zero, $k(\gamma_i)$ is negative, which means the household has a welfare *gain* from aggregate risk. The gain arises because the household is selling so much insurance to more risk-averse households that the resulting risk premiums more than offset the risk the household faces.

We estimate the welfare cost of aggregate risk separately for each village j in the data, but to simplify the notation, we suppress the dependence on j in what follows. Our objective is to estimate the function $k(\gamma_i)$ giving welfare costs of aggregate risk as a function of a household's risk aversion. To do so, we must estimate village j 's prices p_s^* , which appear in the welfare cost formula (22), and village j 's aggregate shocks m_s , which do not appear in the formula but are required to normalize the prices correctly. Schulhofer-Wohl (2008) proposes the following procedure, which we follow here.

We have data on a random sample of households in village j for a sequence of dates $\tau = 1, \dots, T$. Since the model is stationary, we can use the data at different dates to recover information about the states realized at those dates; averages over many dates will be the same as averages over the possible states.

The following notation is useful: For any variable ξ , let $\hat{E}_\tau[\xi]$ be the sample mean of ξ across the households in village j at date τ . Also, let $\theta_i = 1/\gamma_i$ be household i 's risk tolerance, and let $\bar{\theta}$ be the mean of θ_i for all households in village j , including households that are not in our sample.

First, we estimate the mean risk tolerance $\bar{\theta}$ as follows. We use the portfolio-choice method to obtain an estimate $\hat{\gamma}_i$ of the risk aversion of each household i . We then estimate $\bar{\theta}$ by $\hat{\bar{\theta}}$, the sample mean of $1/\hat{\gamma}_i$ among the households in village j . The law of large numbers implies that $\hat{\bar{\theta}}$ is a consistent estimator of $\bar{\theta}$.

Second, we estimate the aggregate shocks m_s as follows. Let $\widehat{\ln m_\tau}$ be the residual from a time-series regression of the log of the sample average of observed consumption $[\ln(\hat{E}_\tau[c_{i\tau}])]$ on an intercept, a time trend and month dummies. Let $\hat{m}_\tau = \exp(\widehat{\ln m_\tau})$ be the estimated aggregate shock at date τ ; Schulhofer-Wohl (2008) shows that, in the limit as the numbers of households and time periods go to infinity, \hat{m}_τ is a consistent estimator of the aggregate shock m_s for the state s that was realized at date τ .

Third, we estimate the prices p_s^* as follows. Given $\bar{\theta}$, let $\widehat{\ln p_\tau^*(\bar{\theta})}$ be $(-1/\bar{\theta})$ times the residual from a time series regression of the sample average of observed log consumption $[(\hat{E}_\tau[\ln c_{i\tau}])]$ on an intercept, time trend and month dummies. (The regression here is the same as that used to estimate aggregate shocks, except that for aggregate shocks, the dependent variable was the log of mean consumption, while for prices, the dependent variable is the mean of log consumption.) If we knew $\bar{\theta}$, we could estimate the price by $\tilde{p}_\tau^*(\bar{\theta}) = \exp[\widehat{\ln p_\tau^*(\bar{\theta})}]$, which Schulhofer-Wohl (2008) shows is a consistent estimator of the price p_s^* for the state s that was realized at date τ . However, we do not know the exact value of the mean risk tolerance but instead must use our estimate $\hat{\bar{\theta}}$. Our estimated prices are thus $\hat{p}_\tau^* = \tilde{p}_\tau^*(\hat{\bar{\theta}})$. We impose the normalization that $\sum_s \pi_s p_s^* m_s = 1$ by scaling the estimated prices such that $T^{-1} \sum_{\tau=1}^T \hat{p}_\tau^* \hat{m}_\tau = 1$. The arguments in Schulhofer-Wohl (2008) imply that \hat{p}_τ^* is a consistent estimator of p_s^* in the limit as the numbers of households and time periods go to infinity.

Finally, given the estimated prices, we estimate the welfare cost of aggregate risk, as a function of the household's risk aversion γ_i , by replacing averages over states with averages over dates and replacing actual with estimated prices in (22):

$$\hat{k}(\gamma_i) = 1 - \left(\frac{1}{T} \sum_{\tau=1}^T (\hat{p}_\tau^*)^{-(1-\gamma_i)/\gamma_i} \right)^{\gamma_i/(1-\gamma_i)} \quad (23)$$

The results in Schulhofer-Wohl (2008) and the consistency of $\hat{\theta}$ imply that $\hat{k}(\gamma_i)$ is a consistent estimator of the welfare cost $k(\gamma_i)$ in the limit as the numbers of households and time periods go to infinity.

Although \hat{k} is a consistent estimator of the true welfare cost, \hat{k} is biased away from zero. The reason is that the estimated aggregate shocks and prices vary over time both because actual shocks hit the economy and because, in a finite sample, measurement error causes the average of households' observed consumption to fluctuate more than the average of their true consumption. In consequence, the data make the economy appear riskier than it really is. Following Schulhofer-Wohl (2008), we solve this problem with a bootstrap bias correction. Let \hat{k} be the estimated willingness to pay in the original sample, and let k_1, \dots, k_Q be estimates calculated using Q different samples of the same size as the original sample, drawn from the original data with replacement. A bias-corrected estimate of k is $2\hat{k}^* - \sum_{q=1}^Q k_q/Q$.

Our bootstrap procedure must deal with two sources of sampling variation: We have data on only some households in the village and on only some time periods from the entire history of the world. To address these two sources of variation, we resample both households and time periods in our bootstrap procedure. Specifically, we first draw households from the original data with replacement, generating a list of households to include in the bootstrap sample. Next, we resample with replacement 12-month blocks of time (to account for serial correlation in shocks) and generate a list of months to include in the bootstrap sample. The bootstrap sample then consists of data points corresponding to each household on the list

of households, for each month on the list of months. However, we do not bootstrap our estimate of the mean risk tolerance $\bar{\theta}$; rather, we estimate this parameter using the original sample, and then employ the same estimate when we calculate the welfare cost in each bootstrap sample. We experimented with bootstrapping our estimate of $\bar{\theta}$ but found that this resulted in extremely large standard errors. In consequence, the bootstrapping accounts for our uncertainty about the welfare cost conditional on the mean risk tolerance, but not for our uncertainty about the mean risk tolerance.

3 Data

We apply the estimation methods described in the previous section to the households in the Townsend Thai Monthly Survey. Several features of this survey make it useful for our study. First, the survey provides us with consumption and asset return data at the household level, allowing us to apply the portfolio-choice method to each household, rather than relying on the aggregate consumption data commonly used in the asset-pricing literature. Also, having household-level data allows us to compare the results from the risk-sharing method, household by household, to those from the portfolio-choice method. Second, the survey has relatively high frequency over many years, providing us with a relatively long time series on consumption and asset return fluctuations. Finally, the survey also has information on household demography that we can use to define kinship networks, which are one potential mechanism resulting in the full insurance assumed by the risk-sharing method.

This section presents a brief background on the survey and descriptive statistics of the variables we analyze. Detailed description of the survey, construction of financial variables, and additional descriptive statistics can be found in Samphantharak and Townsend (2010a,b).

3.1 The Townsend Thai Monthly Survey and Sample Selection

The Townsend Thai Monthly Survey is an on-going intensive monthly survey initiated in 1998 in four provinces of Thailand. Chachoengsao and Lopburi are semi-urban provinces in a more developed central region near the capital city, Bangkok. Buriram and Sisaket provinces, on the other hand, are rural and located in a less developed northeastern region by the border of Cambodia. In each of the four provinces, the survey is conducted in four villages. This monthly survey began with an initial village-wide census. Every structure and every household was enumerated and the defined household units were created based on sleeping and eating patterns. Further, all individuals, households, and residential structures in each of the 16 villages can be identified in subsequent, monthly responses. From the village-wide census, approximately 45 households in each village were randomly sampled to become survey respondents. The survey itself began in August 1998 with a baseline interview on initial conditions of sampled households. The monthly updates started in September 1998 and track inputs, outputs, and changing conditions of the same households over time.

Sample selection for households included in this paper deserves special attention. First, the data used in this paper is based on the 84 months starting from month 5, from January 1999 through December 2005. These months are the entire sample available at the time of the initial writing of this draft and reflect the fact that data for analysis are received from the field survey unit with a considerable lag. Second, we include only the households that were present in the survey throughout the 84 months, dropping households that moved out of the village before month 88 as well as households that were later added to the survey to replace the drop-out households. This criterion also ensures that consumption for each household is strictly positive in every month, allowing us to have a balanced panel of the monthly change in consumption. Third, we drop households whose income is zero in any month. Fourth, since we compute our returns on assets from net income generated from cultivation, livestock, fish and shrimp farming, and retail business, we exclude from this study the households whose

entire income in every period during the 84 months was from wage earnings and not directly from asset-utilizing production activities. Finally, we include in our sample only households that belong to kinship networks, because full insurance is rejected when we include those households not in the network. Kinship networks are important economic features of the villages we study. Relatives usually live in the same village and engage in both financial and non-financial transactions among themselves, potentially leading to full insurance. We construct kinship networks from the information on close familial relatives that are not a part of the household. A household is defined as in a network if it has at least one familial relative living in the same village. There are 369 households in the sample: 72 from Chachoengsao, 82 from Buriram, 93 from Lopburi, and 122 from Sisaket.

3.2 Construction of Variables

3.2.1 Consumption

Our consumption variable includes both monthly consumption of food and monthly expenditure on nonfood items and utilities. Food consumption includes the consumption of outputs such as crops produced by the household, the consumption of food from inventories, and expenditures on food provided by nonhousehold members. Unlike other modules of the Townsend Thai Monthly Survey, several consumption items are collected weekly during months 1-25 and bi-weekly afterward, in order to minimize recall errors. We convert consumption to per capita units by dividing by the number of household members present during the month to which the consumption refers.

We put consumption in real terms by deflating the data with the monthly Consumer Price Index (CPI) at the national level from the Bank of Thailand. Although we realize that inflation in each village could differ from national inflation, we must rely on the national statistics because we do not have a reliable village-level price index at the time of writing

this paper.

3.2.2 Return on Assets

Consistent with the consumption data, we use a household as our unit of analysis and consider the return on households total fixed assets rather than returns on individual assets. Specifically, we define the rate of return on assets (ROA) as the households accrued net income divided by households average total fixed assets over the month in which the income was generated. Since we want to measure the real rate of return rather than the nominal rate of return, we use real accrued net income and the real value of the households fixed assets in our calculation, again deflating the data using the monthly national-level CPI from the Bank of Thailand.

Our simple calculation of ROA raises one obvious problem. In our data, households' net incomes embed contributions from both physical capital and human capital, but we are interested in the risks and returns to physical assets. ROA is therefore overestimated. As a remedy, we calculate the compensation to household labor and subtract this labor compensation from total household income. Compensation to household labor includes both the explicit wage earnings from external labor markets and the implicit shadow wage from labor spent on the households own production activities. The calculation also takes into account the fact that households select into different occupations, as described in detail in Samphantharak and Townsend (2010a, chapter 5) and Townsend and Yamada (2008).

3.2.3 Summary statistics

Table 1 presents descriptive statistics for household consumption and ROA. Mean per capita real consumption is 1294.6 Thai baht per month (in April 2001 baht). According to the Penn World Table (Heston et al., 2009), purchasing power parity in 2001 was 12.35 baht per U.S. dollar, so on average, households in the sample live on the equivalent of about U.S. \$3.50

per person, per day. The table shows that consumption grows slowly, on average, but that consumption growth is quite variable. Asset returns are high on average but also quite variable.

4 Results

Table 2 presents the tests of efficient risk sharing based on (6).⁶ The coefficient on income is statistically significant at the 5% level in only one of the 16 villages. In the other 15 villages, we cannot reject the null hypothesis of full insurance. Further, the effect of idiosyncratic income shocks on consumption is small in practical terms: Even in the village where we reject full insurance, a 1% increase in income is associated with only an 0.013% increase in consumption after controlling for aggregate shocks. When we estimate a common coefficient on income across all villages, we gain statistical power and come close to rejecting the null of full insurance at the 5% level, but the coefficient is small; a 1% increase in income is associated with an 0.003% increase in consumption. We note that the evidence against full insurance is weak even though we have not allowed for nonseparability between consumption and leisure or for heterogeneity in risk preferences, both of which would lead our test to over-reject full insurance.⁷ We think, therefore, that there is little evidence against full insurance in the villages we study, and that it is reasonable to proceed to estimate risk preferences under the maintained hypothesis of full insurance.

We would consider the assumption of full insurance less reasonable if we included households not in kinship networks in our sample. Table 3 tests for full insurance using data on

⁶For these tests only, but not for the rest of the paper, we use total consumption and income rather than per capita variables because converting to per capita units would produce a mechanical correlation between measured per capita income and measured per capita consumption if there is any measurement error in household size.

⁷Classical measurement error in income would lead our test to under-reject full insurance. However, unless the signal-to-noise ratio is very small – which we think is unlikely given the detailed nature of the survey questionnaire – the true elasticity of consumption with respect to idiosyncratic income must still be close to zero in practical terms.

the 136 households that have measured consumption and income in every month but are not in kinship networks. In most villages, this sample is much smaller than the sample of households in networks, so the test is less powerful. Nonetheless, we reject full insurance at the 5% level in three of the 14 villages where we have enough data to perform the test (albeit with the wrong sign – a negative effect of income on consumption – in one of these three villages). In addition, the coefficient on income is larger in the pooled regression when we look at non-network households.

Table 4 presents the tests of the null hypothesis of identical risk preferences, based on the risk-sharing method and the GMM overidentification statistic for moment conditions (10). We reject the null of identical preferences at the 5% level in eight of the 16 villages and at the 10% level in nine of the 16. When we pool the data from all villages, we gain statistical power and strongly reject the null that preferences are identical within each village. (Our pooled test makes no assumptions about whether there is heterogeneity across villages.)

Table 5 presents the estimates of risk preferences in each village from the portfolio-choice method. The table shows the mean risk aversion and mean risk tolerance in each village and the tests for heterogeneity based on the test statistic in (20). The average estimated risk aversion across the entire sample is about 1, though average risk aversion is about half that in some villages and twice as high in others. When we construct the test statistic for heterogeneity using estimates of households' risk tolerance, we strongly reject the null hypothesis of identical preferences. The rejection is less strong when we construct the test statistic using estimates of households' risk aversion (the inverse of risk tolerance); in that case, we reject identical preferences in 10 of the 16 villages.

Table 6 investigates the relationship between the two methods for measuring risk aversion. For each household i that has positive estimated return on assets,⁸ we have two estimates of

⁸Recall that we cannot use the portfolio-choice method when the household's estimated return on assets is negative.

risk tolerance: $1/\hat{\gamma}_i^{RS}$, the estimate from the risk-sharing method, and $1/\hat{\gamma}_i^{PC}$, the estimate from the portfolio-choice method. (We examine $1/\hat{\gamma}_i$ rather than $\hat{\gamma}_i$ because our moment conditions for the risk-sharing method in equation (10) are linear in $1/\gamma_i$ but not in γ_i .) We calculate the correlation of $1/\hat{\gamma}_i^{RS}$ and $1/\hat{\gamma}_i^{PC}$ within each village.⁹ We then use a Monte Carlo permutation test to see whether the correlation is statistically significantly different from zero.¹⁰ The table shows that our two estimates of preferences are positively correlated in 11 of the villages and negatively correlated in the other five. In four of the six villages in which the correlation is statistically significantly different from zero at least at the 10% level, the correlation is positive. We conclude that there is some weak evidence of a positive relationship between our two estimates of each household’s preferences.

Tables 7 and 8 examine the relationship of risk tolerance to observed demographic characteristics of the household. Table 7 uses preferences estimated from the risk-sharing method, while table 8 uses preferences estimated from the portfolio-choice method. For comparability, in each table we restrict the sample to households for which both methods could be used to estimate preferences. (This requires the household to have positive estimated asset returns

⁹Recall that the risk-sharing method identifies preferences only up to a village-specific scale factor. That is, $1/\hat{\gamma}_i^{RS} \approx m_j/\gamma_i$, where m_j is an unknown number. Therefore, while our two estimates should be positively correlated within each village if they are both accurate estimates of the true risk tolerance $1/\gamma_i$, it is not worthwhile to compare the levels of $1/\hat{\gamma}_i^{RS}$ and $1/\hat{\gamma}_i^{PC}$ because the levels can differ even if our methods are correct. We cannot pool the data from all villages and then use village fixed effects to account for m_j (for example, by regressing $1/\hat{\gamma}_i^{RS}$ on $1/\hat{\gamma}_i^{PC}$ and a set of village dummy variables) because m_j multiplies γ_i but the village fixed effects would be additive. Further, we cannot take logs of the risk preference estimates – so that the scale factor m_j would enter additively – because for some households, our estimated $1/\hat{\gamma}_i^{RS}$ is negative.

¹⁰The permutation test randomly re-orders the list of $1/\hat{\gamma}_i^{PC}$ across households in 100,000 ways and computes the correlation of each re-ordered list with the original list of $1/\hat{\gamma}_i^{RS}$. (We use a Monte Carlo test with 100,000 draws rather than an exact test with all possible permutations because some villages in our sample have too many possible permutations – in a village of 34 households, there are $34! \approx 3 \times 10^{38}$ possible re-orderings – to compute all of the possible correlations in a reasonable amount of time.) This procedure gives us the sampling distribution of the correlation coefficient when there is no actual correlation between the two measures of risk tolerance, since by re-ordering the list of $1/\hat{\gamma}_i^{PC}$ we are re-assigning risk tolerance estimates to different households at random. The two-sided p-value for the null hypothesis of no correlation is the fraction of re-ordered correlations that are larger in absolute value than the actual value of $\text{corr}(1/\hat{\gamma}_i^{RS}, 1/\hat{\gamma}_i^{PC})$ in the original data. If $\text{corr}(1/\hat{\gamma}_i^{RS}, 1/\hat{\gamma}_i^{PC}) > 0$, the one-sided p-value is the fraction of re-ordered correlations that are greater than $\text{corr}(1/\hat{\gamma}_i^{RS}, 1/\hat{\gamma}_i^{PC})$; if $\text{corr}(1/\hat{\gamma}_i^{RS}, 1/\hat{\gamma}_i^{PC}) < 0$, the one-sided p-value is the fraction of re-ordered correlations that are less than $\text{corr}(1/\hat{\gamma}_i^{RS}, 1/\hat{\gamma}_i^{PC})$.

so that the portfolio-choice method is feasible.) Demographics are measured in the initial round of the survey. We find little evidence that estimated risk preferences are related to demographics. There is a positive, marginally statistically significant relationship between risk tolerance and the head's age using either method to measure preferences. Education, net wealth, and the numbers of men, women and children in the household are not associated with either measure of risk tolerance. These patterns persist whether we include village fixed effects in the regressions or not. In addition, observed demographics explain only a few percent of the variation in estimated risk tolerance. Theory provides little guidance as to whether we should expect observable variables to be related to preferences. For example, net wealth may depend in large part on a household's initial endowment when the economy began, and – since initial endowments – theory has little to say about whether the initial endowment, and thus wealth, will be related to preferences. Recall also that, under complete markets, wealth *per se* has nothing to do with risk aversion: Complete markets lead to a complete separation between consumption and production, so there is no reason why risk preferences in themselves should affect how much wealth a household accumulates.¹¹

In table 9 and figure 2, we turn to estimating the welfare costs of aggregate risk. Table 9 shows the estimated mean risk tolerance in each village and the willingness to pay to eliminate aggregate risk for a household with that risk tolerance. In every village, the estimated mean risk tolerance is above 1. The estimates do not necessarily imply, though, that the mean risk aversion is below 1: Risk aversion is the inverse of risk tolerance, so by Jensen's inequality, mean risk aversion is greater than the inverse of mean risk tolerance. In each village, a household whose risk tolerance matches the mean risk tolerance of that village is estimated to face welfare losses from aggregate risk. The welfare losses for households with the mean

¹¹Note, however, that households with a higher elasticity of intertemporal substitution *will* accumulate more wealth if the economy is growing over time (Dumas, 1989; Wang, 1996). With time-separable expected utility, the coefficient of relative risk aversion is the inverse of the elasticity of intertemporal substitution. Thus there is some reason, in a time-separable expected utility model, to expect a relationship between wealth and risk aversion. A model with, e.g., recursive utility could break this link.

risk tolerance are on the order of one-half to one percent of mean consumption, or about 10 times what has been estimated for the United States (Lucas, 1987; Schulhofer-Wohl, 2008). However, owing to the small sample size for each village, our estimates are imprecise and the 95% confidence intervals for the mean household’s welfare loss include zero. Figure 2 shows the importance of heterogeneity for understanding the welfare cost of risk. In each village, the less risk-averse a household is, the smaller its welfare cost, and households that are sufficiently close to risk neutral have welfare gains from aggregate risk. For example, in village 7 in Chachoengsao, some very risk-averse households have welfare losses from aggregate risk equivalent to about four percent of consumption, but sufficiently risk-tolerant households could have welfare *gains* in excess of six percent of consumption.

5 Conclusion

Conclusion goes here.

A Derivation of portfolio-choice method

We show here how to obtain (14) from (12) and (13) under the assumptions that (a) households choose portfolios on the mean-variance frontier, (b) households have CRRA preferences, (c) seasonally adjusted consumption growth has a log-normal distribution and (d) the variance of seasonally adjusted consumption growth is “small.” Much of our exposition parallels Cochrane (2001) and Samphantharak and Townsend (2010b).

Recall that for any two random variables A and B , $\text{Cov}(A, B) = \text{Corr}(A, B)\sqrt{\text{Var}(A)}\sqrt{\text{Var}(B)}$.

Hence (12) can be written

$$1 = \mathbb{E} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right] \mathbb{E}[R_{t+1}^k] \\ + \text{Corr} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k \right] \sqrt{\text{Var} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right]} \sqrt{\text{Var} [R_{t+1}^k]} . \quad (24)$$

Substituting (13) into (24) yields

$$1 = \mathbb{E}[R_{t+1}^k] + \text{Corr} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k \right] \sqrt{\text{Var} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right]} \sqrt{\text{Var} [R_{t+1}^k]} . \quad (25)$$

Then, rearranging terms in (25), we obtain

$$\frac{\mathbb{E}[R_{t+1}^k] - 1}{\sqrt{\text{Var}[R_{t+1}^k]}} = -\sqrt{\text{Var} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right]} \text{Corr} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k \right] . \quad (26)$$

Since $-1 \leq \text{Corr} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k \right] \leq 1$, (26) implies

$$\left| \frac{\mathbb{E}[R_{t+1}^k] - 1}{\sqrt{\text{Var}[R_{t+1}^k]}} \right| \leq \sqrt{\text{Var} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right]} , \quad (27)$$

which is the Hansen and Jagannathan (1991) bound on the risk premium.

Inequality (27) applies not just to any single asset k but to any combination of assets – in particular, to the household's actual portfolio. Thus if $R_{i,t+1}^P$ is the gross return on the household's portfolio, then

$$\left| \frac{\mathbb{E}[R_{t+1}^P] - 1}{\sqrt{\text{Var}[R_{t+1}^P]}} \right| \leq \sqrt{\text{Var} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right]} , \quad (28)$$

Portfolios on the mean-variance frontier maximize the expected return for any given variance.

Therefore, if we apply assumption (a) – that the household chooses a portfolio on the mean-variance frontier – (28) must hold with equality. (Otherwise, there would be a portfolio with the same variance as the household’s portfolio but higher expected return, contradicting the assumption that the household’s portfolio is on the mean-variance frontier.) Applying assumption (b) – CRRA preferences – (28) at equality becomes

$$\left| \frac{\mathbb{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \sqrt{\text{Var} \left[\beta_i \frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \left(\frac{c_{i,t+1}^*}{c_{it}^*} \right)^{-\gamma_i} \right]}. \quad (29)$$

Using (13), it is convenient to rewrite (29) as

$$\left| \frac{\mathbb{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \frac{\sqrt{\text{Var} \left[\beta_i \frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \left(\frac{c_{i,t+1}^*}{c_{it}^*} \right)^{-\gamma_i} \right]}}{\mathbb{E} \left[\beta_i \frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \left(\frac{c_{i,t+1}^*}{c_{it}^*} \right)^{-\gamma_i} \right]}. \quad (30)$$

We define seasonally adjusted consumption $x_{it}^* = \xi_{i,m(t)}^{-1/\gamma_i} c_{it}^*$. Then (30) can be written as

$$\left| \frac{\mathbb{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \frac{\sqrt{\text{Var} \left[\beta_i \left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right]}}{\mathbb{E} \left[\beta_i \left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right]}. \quad (31)$$

We now use assumption (c), that seasonally adjusted consumption growth $(x_{i,t+1}^*/x_{it}^*)$ has a log-normal distribution. Note that if

$$\Delta \ln x_{it}^* = \ln(x_{i,t+1}^*/x_{it}^*) \sim N(\mu_{dx}, \sigma_{dx}^2), \quad (32)$$

then

$$\ln(x_{i,t+1}^*/x_{it}^*)^{-\gamma_i} = -\gamma_i \ln(x_{i,t+1}^*/x_{it}^*) \sim N(-\gamma_i \mu_{dx}, \gamma_i^2 \sigma_{dx}^2). \quad (33)$$

Further, for any random variable A , if $\ln A \sim N(\mu_A, \sigma_A^2)$, then $E(A) = \exp(\mu_A + \sigma_A^2/2)$ and $\text{Var}(A) = [\exp(\sigma_A^2) - 1] \exp(2\mu_A + \sigma_A^2)$. Thus under the log-normality assumption,

$$E \left[\beta_i \left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right] = \beta_i E \left[\left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right] = \beta_i \exp \left[-\gamma_i \mu_{dx} + \frac{\gamma_i^2 \sigma_{dx}^2}{2} \right] \quad (34)$$

and

$$\text{Var} \left[\beta_i \left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right] = \beta_i^2 \text{Var} \left[\left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right] = \beta_i^2 [\exp(\gamma_i^2 \sigma_{dx}^2) - 1] \exp(-2\gamma_i \mu_{dx} + \gamma_i^2 \sigma_{dx}^2). \quad (35)$$

Substituting (34) and (35) into the right-hand side of (31) gives

$$\left| \frac{E[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \frac{\beta_i \sqrt{\exp(\gamma_i^2 \sigma_{dx}^2) - 1} \exp(-\gamma_i \mu_{dx} + \gamma_i^2 \sigma_{dx}^2/2)}{\beta_i \exp \left[-\gamma_i \mu_{dx} + \frac{\gamma_i^2 \sigma_{dx}^2}{2} \right]} = \sqrt{\exp(\gamma_i^2 \sigma_{dx}^2) - 1}. \quad (36)$$

Finally, using assumption (d), if σ_{dx}^2 is close to zero, then $\exp(\gamma_i^2 \sigma_{dx}^2) - 1 \approx \gamma_i^2 \sigma_{dx}^2$. Thus (approximately)

$$\left| \frac{E[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \gamma_i \sqrt{\text{Var}[\Delta \ln x_{it}^*]} \quad (37)$$

and hence

$$\gamma_i = \frac{1}{\sqrt{\text{Var}[\Delta \ln x_{it}^*]}} \left| \frac{E[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right|. \quad (38)$$

References

- Breeden, Douglas T., 1979, "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics* 7(3), 265–296.
- Chamberlain, Gary and Michael Rothschild, 1983, "Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets," *Econometrica* 51(5), 1281–1304.
- Chiappori, Pierre-André, and Monica Paiella, 2008, "Relative Risk Aversion Is Constant: Evidence from Panel Data," manuscript, Columbia University and University of Naples Parthenope.
- Cochrane, John H., 1991, "A Simple Test of Consumption Insurance," *Journal of Political Economy* 99(5), 957–976.
- Cochrane, John, 2001, *Asset Pricing*, Princeton, N.J.: Princeton University Press.
- Deaton, Angus, 1997, *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*, Baltimore: Johns Hopkins University Press for the World Bank.
- Dumas, Bernard, 1989, "Two-Person Dynamic Equilibrium in the Capital Market," *Review of Financial Studies* 2(2), 157-188.
- Guiso, Luigi, and Monica Paiella, 2008, "Risk Aversion, Wealth, And Background Risk," *Journal of the European Economic Association* 6(6), 1109–1150.
- Han, Chirok and Peter C.B. Phillips, 2006, "GMM With Many Moment Conditions," *Econometrica* 74(1), 147–192.
- Hansen, Lars Peter, 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica* 50(4), 1029–1054.
- Hansen, Lars Peter and Ravi Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy* 99(2), 225–262.
- Hayashi, Fumio, Joseph Altonji and Laurence Kotlikoff, 1996, "Risk-sharing Between and Within Families," *Econometrica* 64(2), 261–294.
- Heston, Alan, Robert Summers and Bettina Aten, 2009, "Penn World Table Version 6.3," Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.
- Kinnan, Cynthia, and Robert M. Townsend, 2010, "Kinship and Financial Networks, Formal Financial Access and Risk Reduction," manuscript, Northwestern University and MIT.
- Lucas, Robert E., Jr., 1978, "Asset Prices in an Exchange Economy," *Econometrica* 46(6), 1429–1446.

- Lucas, Robert E., Jr., 1987, *Models of Business Cycles*, New York: Basil Blackwell.
- Mace, Barbara J., 1991, "Full Insurance in the Presence of Aggregate Uncertainty," *Journal of Political Economy* 99(5), 928–956.
- Mazzocco, Maurizio and Shiv Saini, 2009, "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," manuscript, UCLA.
- Mehra, Rajnish and Edward C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15(2), 145–161.
- Rubinstein, Mark, 1976, "The Valuation of Uncertain Income Streams and the Price of Options," *Bell Journal of Economics* 7(2), 407–425.
- Samphantharak, Krislert, and Robert M. Townsend, 2010a, *Households as Corporate Firms: An Analysis of Household Finance Using Integrated Household Surveys and Corporate Financial Accounting*, Econometric Society Monograph 46. Cambridge, U.K.: Cambridge University Press.
- Samphantharak, Krislert and Robert M. Townsend, 2010b, "Risk and Return in Village Economies," manuscript, UCSD and MIT.
- Scheinkman, José Alexandre, 1984, "General Equilibrium Models of Economic Fluctuations: A Survey of Theory," manuscript, University of Chicago.
- Schulhofer-Wohl, Sam, 2008, "Heterogeneous risk preferences and the welfare cost of business cycles," *Review of Economic Dynamics* 11(4), 761–780.
- Schulhofer-Wohl, Sam, 2010, "Heterogeneity and Tests of Risk Sharing," manuscript, Princeton University.
- Townsend, Robert M., 1994, "Risk and Insurance in Village India," *Econometrica* 62(3), 539–591.
- Townsend, Robert M., and Hiroyuki Yamada, 2008, "The Summary of the Monthly Household Survey of Thailand on Labor Issues," manuscript, University of Chicago.
- Vissing-Jørgensen, Annette, 2002, "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy* 110(4), 825–853.
- Wang, Jiang, 1996, "The term structure of interest rates in a pure exchange economy with heterogeneous investors," *Journal of Financial Economics* 41, 75–110.
- Wilson, Robert, 1968, "The Theory of Syndicates," *Econometrica* 36(1), 119–132.

Table 1: Descriptive statistics.

Variable	mean	std. dev.	observations
real total consumption per capita	1298.5	2367.8	30,576
ln(real total consumption per capita)	6.85	0.69	30,576
one-month consumption growth	0.0027	0.5048	30,212
return on assets	7.32	121.78	30,576

The table reports descriptive statistics for key variables. The unit of analysis is the household-month. Consumption (in Thai baht) is monthly household food consumption and monthly household expenditure on nonfood consumption items. Consumption is adjusted to real per capita units using monthly household size data and nationwide Consumer Price Index (base month April 2001). One-month consumption growth is the log change in real total consumption per capita from month $t - 1$ to month t , and is calculated for all months except the first month in the sample. Return on assets is in percent.

Table 2: Tests of efficient risk sharing.

village	coeff.	std. err.	p-value	obs.	HH	R^2
<i>Chachoengsao</i>						
2	0.0090	0.0054	0.112	1764	21	0.103
4	-0.0028	0.0070	0.692	2016	24	0.176
7	0.0032	0.0085	0.718	924	11	0.211
8	-0.0011	0.0048	0.827	1260	15	0.239
<i>Buriram</i>						
2	-0.0081	0.0113	0.479	2352	27	0.130
10	0.0051	0.0047	0.304	1008	11	0.199
13	0.0098	0.0057	0.105	1596	18	0.250
14	0.0073	0.0060	0.237	1932	22	0.246
<i>Lop Buri</i>						
1	0.0099	0.0068	0.160	1932	23	0.074
3	-0.0102	0.0132	0.452	1260	15	0.137
4	-0.0008	0.0062	0.898	2520	30	0.080
6	0.0004	0.0057	0.945	2100	25	0.146
<i>Sisaket</i>						
1	0.0057	0.0041	0.177	2352	28	0.180
6	0.0004	0.0034	0.899	3276	39	0.185
9	0.0132*	0.0033	0.000	3024	36	0.217
10	0.0128	0.0058	0.041	1596	19	0.157
<i>pooled</i>						
-	0.0034	0.0017	0.052	30576	364	0.167

The table reports the effect of idiosyncratic income shocks on consumption. Unit of observation is household-month. Consumption is monthly household food consumption and monthly household expenditure on nonfood consumption items. Income is monthly accrued income. Consumption and income are adjusted for inflation using national Consumer Price Index. Each row reports a separate regression using data from one village. Column labeled “coeff.” reports the coefficient on log income in an OLS regression of log consumption on household fixed effects, time fixed effects and log income (6); “std. err.” is the standard error, clustered by household; p-value is for a test of the null hypothesis that the coefficient on log income is zero; “obs.” is the number of household-month observations; and “hh” is the number of households. Pooled regression uses data from all villages and interacts time effects with village effects to allow different aggregate shocks by village. * indicates coefficient is statistically significantly different from zero at 5% level.

Table 3: Tests of efficient risk sharing, for households not in kinship networks.

village	coeff.	std. err.	p-value	obs.	HH	R^2
<i>Chachoengsao</i>						
2	-0.0031	0.0071	0.669	1680	20	0.104
4	0.0195	0.0167	0.264	1176	14	0.194
7	-0.0019	0.0067	0.780	1596	19	0.150
8	-0.0059	0.0081	0.480	1344	16	0.206
<i>Buriram</i>						
2	0.0143	0.0133	0.317	672	8	0.183
10	-	-	-	84	1	.
13	0.0101	0.0049	0.108	420	5	0.504
14	0.0079*	0.0022	0.036	336	4	0.465
<i>Lop Buri</i>						
1	-0.0045	0.0106	0.681	756	9	0.164
3	0.0216	0.0278	0.474	504	6	0.220
4	0.0001	0.0071	0.991	588	7	0.180
6	-0.0125	0.0062	0.101	504	6	0.228
<i>Sisaket</i>						
1	0.0278*	0.0065	0.002	840	10	0.213
6	0.0129	0.0126	0.383	336	4	0.413
9	-	-	-	84	1	.
10	-0.0004*	0.0000	0.000	168	2	0.615
<i>pooled</i>						
-	0.0032	0.0032	0.315	11088	132	0.210

The table reports the effect of idiosyncratic income shocks on consumption among households not in kinship networks. Unit of observation is household-month. Consumption is monthly household food consumption and monthly household expenditure on nonfood consumption items. Income is monthly accrued income. Consumption and income are adjusted for inflation using national Consumer Price Index. Each row reports a separate regression using data from one village; test is not performed if village has only one household not in a network. Column labeled “coeff.” reports the coefficient on log income in an OLS regression of log consumption on household fixed effects, time fixed effects and log income (6); “std. err.” is the standard error, clustered by household; p-value is for a test of the null hypothesis that the coefficient on log income is zero; “obs.” is the number of household-month observations; and “hh” is the number of households. Pooled regression uses data from all villages and interacts time effects with village effects to allow different aggregate shocks by village. * indicates coefficient is statistically significantly different from zero at 5% level.

Table 4: Tests for heterogeneity in risk preferences (risk-sharing method).

village	χ^2	d.f.	p-value
<i>Chachoengsao</i>			
2	38.70	20	0.007
4	35.46	23	0.047
7	21.20	10	0.020
8	33.55	14	0.002
<i>Buriram</i>			
2	34.32	26	0.127
10	9.93	10	0.446
13	28.41	17	0.040
14	37.17	21	0.016
<i>Lop Buri</i>			
1	25.64	22	0.267
3	29.80	14	0.008
4	37.89	29	0.125
6	44.79	24	0.006
<i>Sisaket</i>			
1	40.05	27	0.051
6	44.33	38	0.222
9	37.45	35	0.357
10	22.33	18	0.218
<i>pooled</i>			
-	521.03	348	5.0×10^{-9}

The table reports tests of the null hypothesis that all households in a given village have the same coefficient of relative risk tolerance. χ^2 is the overidentification test statistic for the null hypothesis that all households in the village have the same risk tolerance, obtained by estimating moment condition (10) by two-step efficient GMM under the restriction $\gamma_1 = \gamma_2 = \dots = \gamma_{N_j}$; “d.f.” is the degrees of freedom of the χ^2 statistic, equal to the number of households in the village minus one. Pooled test is for the hypothesis that risk tolerance is constant within each village, without assuming anything about heterogeneity across villages.

Unit of observation is household-month. Consumption is monthly household food consumption and monthly household expenditure on nonfood consumption items. Consumption is adjusted to real per capita units using monthly household size data and nationwide Consumer Price Index.

Table 5: Tests for heterogeneity in risk preferences (portfolio-choice method).

village	households	risk aversion γ_i			risk tolerance $1/\gamma_i$		
		mean	χ^2	p-value	mean	χ^2	p-value
<i>Chachoengsao</i>							
2	13	2.00	277.29	0.0000	1.56	3543.60	0.0000
4	21	0.79	78.44	0.0000	2.47	1646.42	0.0000
7	6	0.98	6.69	0.3509	1.28	32.21	0.0000
8	14	0.61	31.11	0.0053	5.11	7986.64	0.0000
<i>Buriram</i>							
2	18	0.62	12.54	0.8184	2.97	368.59	0.0000
10	8	0.34	5.87	0.6618	4.02	147.64	0.0000
13	10	0.41	14.27	0.1610	7.61	2255.00	0.0000
14	15	0.84	73.55	0.0000	3.55	4209.49	0.0000
<i>Lop Buri</i>							
1	19	1.20	96.08	0.0000	1.36	1011.17	0.0000
3	8	2.12	348.07	0.0000	1.33	3981.73	0.0000
4	27	1.40	173.59	0.0000	1.29	2061.54	0.0000
6	24	1.82	485.27	0.0000	1.29	3074.97	0.0000
<i>Sisaket</i>							
1	22	0.43	21.94	0.4633	3.78	457.10	0.0000
6	34	0.78	117.07	0.0000	1.85	2010.67	0.0000
9	22	0.76	33.96	0.0495	3.24	2141.48	0.0000
10	13	0.47	9.68	0.7199	2.90	36.03	0.0006
<i>pooled</i>							
-	274	0.98	1358.43	0.0000	2.64	77568.89	0.0000

The table reports tests of the null hypothesis that all households in a given village have the same coefficient of relative risk aversion or coefficient of relative risk aversion. The sample includes only households with positive estimated return on assets (so portfolio-choice method is feasible). “households” is the number of households in the village included in the sample; $\chi^2 = \sum_{i \in j} [(\hat{\gamma}_i^{PC} - \hat{\gamma}_j) / \widehat{s.e.}(\hat{\gamma}_i^{PC})]^2$ is the test statistic for the null that all households in the village have the same preferences. Standard errors are obtained by bootstrapping, using 1,000 draws from the original sample with replacement; bootstrap draws are of 12-month blocks to account for serial correlation. Pooled test is for the hypothesis that risk tolerance is constant within and across villages.

Unit of observation is household-month. Consumption is monthly household food consumption and monthly household expenditure on nonfood consumption items. Consumption is adjusted to real per capita units using monthly household size data and nationwide Consumer Price Index.

Table 6: Correlation between estimated risk tolerance from risk-sharing and portfolio-choice methods.

village	corr.	p-value		HH
		1-sided	2-sided	
<i>Chachoengsao</i>				
2	0.180	0.278	0.557	13
4	0.254	0.145	0.266	21
7	-0.652	0.083	0.167	6
8	-0.416	0.005	0.120	14
<i>Buriram</i>				
2	0.337	0.074	0.159	18
10	0.522	0.073	0.200	8
13	-0.003	0.536	0.994	10
14	0.179	0.291	0.563	15
<i>Lop Buri</i>				
1	0.118	0.314	0.632	19
3	0.129	0.400	0.745	8
4	0.049	0.390	0.810	27
6	0.790	0.000	0.000	24
<i>Sisaket</i>				
1	-0.178	0.225	0.455	22
6	0.279	0.062	0.112	34
9	0.034	0.419	0.853	22
10	-0.014	0.559	0.959	13

The table reports correlations between estimates of households' preferences obtained by the two different methods developed in the text. The unit of observation is the household. The sample includes only households with positive estimated return on assets (so portfolio-choice method is feasible). P-values are from a Monte Carlo permutation test in which we randomly re-order the list of $1/\hat{\gamma}_i^{PC}$ in 100,000 ways and compute the correlation of each re-ordered list with the original list of $1/\hat{\gamma}_i^{RS}$. The p-value for the null hypothesis of no correlation is the percentile of $\text{corr}(1/\hat{\gamma}_i^{RS}, 1/\hat{\gamma}_i^{PC})$ in the distribution of re-ordered correlations.

Table 7: Association between household demographics and estimated risk tolerance from risk-sharing method.

	estimated risk tolerance						
	<i>A. Without village fixed effects</i>						
adult men	0.012						0.011
	(0.009)						(0.011)
adult women		0.008					0.001
		(0.008)					(0.008)
children			0.003				0.005
			(0.005)				(0.004)
head's age				0.001*			0.001
				(0.000)			(0.000)
highest education					0.000		-0.000
					(0.002)		(0.002)
net wealth (millions of baht)						-0.000	-0.001
						(0.002)	(0.002)
joint signif. p-value							0.147
R-squared	0.011	0.004	0.002	0.020	0.000	0.000	0.032
	<i>B. With village fixed effects</i>						
adult men	0.009						0.005
	(0.007)						(0.009)
adult women		0.008					0.006
		(0.008)					(0.008)
children			-0.006				-0.006
			(0.005)				(0.005)
head's age				0.000			0.000
				(0.000)			(0.001)
highest education					0.001		0.000
					(0.002)		(0.002)
net wealth (millions of baht)						0.001	0.001
						(0.002)	(0.002)
joint signif. p-value							0.281
R-squared	0.215	0.214	0.215	0.214	0.211	0.213	0.223
Observations	274	274	274	274	274	274	274

The table reports the association between demographic variables and households' estimated preferences obtained by the risk-sharing method. The unit of observation is the household. The sample includes only households with positive estimated return on assets (so portfolio-choice method is feasible). Heteroskedasticity-robust standard errors clustered by village in parentheses. Demographics are measured in the initial survey. Net wealth is in millions of baht. "Joint signif. p-value" is the p-value for the null hypothesis that the coefficients on all of the demographic variables are zero in a regression including all the variables at once. * indicates coefficient is statistically significantly different from zero at the 5% level.

Table 8: Association between household demographics and estimated risk tolerance from portfolio-choice method.

	estimated risk tolerance						
	<i>A. Without village fixed effects</i>						
adult men	0.158						0.102
	(0.395)						(0.465)
adult women		-0.015					-0.215
		(0.257)					(0.183)
children			0.077				0.134
			(0.243)				(0.251)
head's age				0.024			0.030
				(0.016)			(0.015)
highest education					0.031		0.053
					(0.061)		(0.063)
net wealth (millions of baht)						-0.031	-0.055
						(0.056)	(0.071)
joint signif. p-value							0.285
R-squared	0.001	0.000	0.001	0.006	0.001	0.001	0.012
	<i>B. With village fixed effects</i>						
adult men	0.172						0.070
	(0.409)						(0.518)
adult women		-0.111					-0.318
		(0.202)					(0.178)
children			-0.001				0.049
			(0.270)				(0.271)
head's age				0.017			0.025
				(0.016)			(0.015)
highest education					0.068		0.092
					(0.055)		(0.080)
net wealth (millions of baht)						-0.001	-0.022
						(0.073)	(0.090)
joint signif. p-value							0.095
R-squared	0.125	0.124	0.124	0.126	0.127	0.124	0.133
Observations	274	274	274	274	274	274	274

The table reports the association between demographic variables and households' estimated preferences obtained by the portfolio-choice method. The unit of observation is the household. The sample includes only households with positive estimated return on assets (so portfolio-choice method is feasible). Heteroskedasticity-robust standard errors clustered by village in parentheses. Demographics are measured in the initial survey. Net wealth is in millions of baht. "Joint signif. p-value" is the p-value for the null hypothesis that the coefficients on all of the demographic variables are zero in a regression including all the variables at once.

Table 9: Estimated mean risk tolerance and welfare cost of aggregate risk for a household with the mean risk tolerance, by village.

village	mean risk tolerance	WTP to eliminate aggregate risk	
		bias-corrected estimate	95% confidence interval
<i>Chachoengsao</i>			
2	1.56	0.4%	(-0.7%,1.1%)
4	2.47	0.2%	(-0.6%,0.6%)
7	1.28	2.9%	(-0.5%,5.3%)
8	5.11	0.3%	(-0.1%,0.6%)
<i>Buriram</i>			
2	2.97	0.6%	(-0.4%,1.3%)
10	4.02	0.4%	(-1.7%,1.4%)
13	7.61	0.2%	(-0.1%,0.4%)
14	3.55	0.7%	(-0.3%,1.3%)
<i>Lop Buri</i>			
1	1.36	0.2%	(-1.4%,0.9%)
3	1.33	0.9%	(-2.8%,2.2%)
4	1.29	0.4%	(-0.7%,1.1%)
6	1.29	0.4%	(-0.8%,1.1%)
<i>Sisaket</i>			
1	3.78	0.3%	(-0.2%,0.6%)
6	1.85	0.9%	(-0.4%,1.8%)
9	3.24	0.5%	(0.0%,1.0%)
10	2.90	1.0%	(-1.2%,2.1%)

The table reports the mean risk tolerance among households in each village, as estimated with the portfolio-choice method, and the estimated willingness to pay to eliminate aggregate risk for a household with the mean risk tolerance. The mean risk tolerance is estimated based only on households with positive estimated return on assets (so portfolio-choice method is feasible). The willingness to pay is reported as a percentage of mean consumption; the table shows the bootstrap bias-corrected point estimate and the 95% equal-tailed percentile confidence interval, calculated from 1,000 bootstrap samples drawn from the original sample with replacement. To construct each bootstrap sample, we first draw households from the original data with replacement, generating a list of households to include in the bootstrap sample; next, we resample with replacement 12-month blocks of time (to account for serial correlation in shocks) and generate a list of months to include in the bootstrap sample; finally, the bootstrap sample consists of data points corresponding to each household on the list of households, for each month on the list of months.




Figure 1: Volatility of aggregate consumption, by village.

Each graph shows the time series of seasonally adjusted, detrended aggregate consumption for a given village. The unit of observation is the village-month. We compute seasonally adjusted, detrended aggregate consumption as follows. For each household in the village, we find the residuals from an OLS regression of the time series of the household's log consumption on a household-specific intercept, household-specific trend, and household-specific month indicator variables. Seasonally adjusted, detrended aggregate consumption for a given village and month is the mean of the log consumption residuals for that village's households in that month.




Figure 2: Welfare costs of aggregate risk.

Each graph shows the bootstrap bias-corrected estimate of the willingness to pay to eliminate aggregate risk, as a function of the household's risk tolerance, for a given village. Positive numbers mean the household has a welfare gain from aggregate risk and is willing to pay to eliminate risk; negative numbers mean the household has a welfare loss from aggregate risk.