

Federal Reserve Bank of Minneapolis
Research Department

Portfolio Choices and Risk Preferences in Village Economies*

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

Working Paper 706

May 2013

ABSTRACT

We use a model of optimal portfolio choice to measure heterogeneity in risk aversion among households in Thai villages. There is substantial heterogeneity in risk preferences, positively correlated in most villages with alternative estimates based on a full risk-sharing model.

Keywords: Risk preferences; Heterogeneity; Portfolio choice
JEL classification: D12, D14, D53, D81, D91, G11, O16

*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. We thank Joan Gieseke for editorial assistance. Portions of this paper previously circulated as part of Research Department Working Paper 683, Federal Reserve Bank of Minneapolis, "Heterogeneity and Risk Sharing in Village Economies." 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

In recent research (Chiappori et al., 2013), we showed how to use a full risk-sharing model to measure heterogeneity in risk aversion among households running farm and non-farm enterprises in a developing country. This paper complements that analysis with a measure of risk preferences based on portfolio choice. The intuition behind the portfolio choice 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. This argument holds even if markets are incomplete; hence, our estimates here provide a robustness check on the strong assumption of full insurance in Chiappori et al. (2013).

We apply the method to a long panel of households in 16 villages in Thailand. We find that the estimates based on portfolio choice are positively correlated in most villages with the estimates based on full insurance. However, the correlations are weak, suggesting that each set of estimates may contain substantial amounts of measurement error. As with the results in Chiappori et al. (2013) on full insurance, and consistent with studies in other contexts such as Chiappori and Paiella (2011) and Guiso and Paiella (2008), we find that risk tolerance is not significantly correlated with demographic variables or household wealth.

One advantage of the portfolio-choice method for estimating risk preferences is that it identifies the mean level of risk tolerance in a village, unlike the risk-sharing approach, which identifies risk preferences only up to an unknown village-specific scale parameter. Thus, the portfolio-choice method allows us to determine whether average risk aversion differs across villages.

In addition, the portfolio-choice method provides a valuable check on the risk-sharing

method because the two methods differ 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.

The paper proceeds as follows. Section 2 details the theory behind our method, section 3 describes the data, section 4 presents the results, and section 5 concludes. The appendix contains some mathematical derivations.

2. Theory

Following Chiappori et al. (2013), we assume that household i 's preferences over consumption sequences $\{c_{it}^*(s^t)\}$ 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 s^t is the history of states up to date t , 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 non-stochastic 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 use a simple portfolio choice and asset pricing model (Breedon, 1979; Lucas, 1978; Rubinstein, 1976) to recover households' preferences from their asset holdings.¹ 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 that the household chooses to own,

$$1 = \frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \frac{\beta_i}{u'_i(c_{it}^*)} \mathbb{E}_t[u'_i(c_{i,t+1}^*) R_{t+1}^k]. \quad (2)$$

Taking the unconditional expectation of both sides, applying the law of iterated expectations and rearranging terms, the Euler equation requires

$$\begin{aligned} 1 &= \mathbb{E} \left[\frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \frac{\beta_i u'_i(c_{i,t+1}^*)}{u'_i(c_{it}^*)} R_{t+1}^k \right] \\ &= \mathbb{E} \left[\frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \frac{\beta_i u'_i(c_{i,t+1}^*)}{u'_i(c_{it}^*)} \right] \mathbb{E}[R_{t+1}^k] + \text{Cov} \left[\frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \frac{\beta_i u'_i(c_{i,t+1}^*)}{u'_i(c_{it}^*)}, R_{t+1}^k \right]. \quad (3) \end{aligned}$$

In particular, (3) 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,

$$1 = \mathbb{E} \left[\frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \frac{\beta_i u'_i(c_{i,t+1}^*)}{u'_i(c_{it}^*)} \right]. \quad (4)$$

¹The 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.

We show in appendix A that if we assume (a) households have constant relative risk aversion (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 (3) and (4) imply the following equation for risk aversion:

$$\gamma_i = \frac{1}{\sqrt{\text{Var}[\Delta \ln x_{it}^*]}} \left| \frac{\text{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right|, \quad (5)$$

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 (5) 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

$$\Delta \ln x_{it}^* = \mu + e_{it}, \quad (6)$$

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

$$\begin{aligned} \Delta \ln c_{it} &= \Delta \ln x_{it}^* + \frac{1}{\gamma_i} \Delta \ln \xi_{i,m(t)} + \Delta \ln \epsilon_{it} \\ &= \frac{1}{\gamma_i} \Delta \ln \xi_{i,m(t)} + \mu + e_{it} + \epsilon_{i,t+1} - \epsilon_{it}. \end{aligned} \quad (7)$$

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 (7), if we had an infinitely long time series for the house-

hold, 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 (5) 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 (5) 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 (5). 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

$$\frac{\hat{\gamma}_i^{PC} - \bar{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \xrightarrow{d} N(0, 1), \quad (8)$$

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

$$\sum_{i \in j} \left(\frac{\hat{\gamma}_i^{PC} - \bar{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 \xrightarrow{d} \chi^2(0, N_j). \quad (9)$$

We cannot calculate the χ^2 statistic in (9) because we do not observe $\bar{\gamma}_j$ but only the estimate $\hat{\gamma}_j$. Observe that

$$\sum_{i \in j} \left(\frac{\hat{\gamma}_i^{PC} - \hat{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 = \sum_{i \in j} \left(\frac{\hat{\gamma}_i^{PC} - \bar{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 + \sum_{i \in j} \left(\frac{\bar{\gamma}_j - \hat{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 + 2 \sum_{i \in j} \frac{(\hat{\gamma}_i^{PC} - \bar{\gamma}_j)(\bar{\gamma}_j - \hat{\gamma}_j)}{[\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})]^2}. \quad (10)$$

Since $\hat{\gamma}_j \xrightarrow{p} \bar{\gamma}_j$, the second and third terms on the right-hand side of (10) converge in proba-

bility to zero. Therefore,

$$\sum_{i \in j} \left(\frac{\hat{\gamma}_i^{PC} - \hat{\gamma}_j}{\widehat{\text{s.e.}}(\hat{\gamma}_i^{PC})} \right)^2 \xrightarrow{d} \chi^2(0, N_j). \quad (11)$$

We use (11) 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 . 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 (11). 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 that the rejection would be even stronger if we accounted for correlation in the estimation errors.

Note that using each household’s idiosyncratic asset returns to estimate preferences in the portfolio-choice method is not inconsistent with the full insurance assumption that we

used in Chiappori et al. (2013): Even if households are fully insured against idiosyncratic shocks to asset returns, the Euler equation (2) 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.

3. Data

We use data from the Townsend Thai Monthly Survey. The survey provides us with 84 monthly observations on consumption and asset returns 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. We refer readers to Chiappori et al. (2013) for an overview of the survey, our sample selection criteria, and descriptions of most of the variables we analyze. This section describes the measure of return on assets (ROA) that we need for the portfolio-choice method and some additional selection criteria that the method requires.

Consistent with the consumption data, we use a household as our unit of analysis and consider the return on the household’s total fixed assets rather than returns on individual assets. Specifically, we define ROA as the household’s accrued net income divided by the household’s 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 household’s fixed assets in our calculation, 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 household's own production activities. The calculation also takes into account that households select into different occupations, as described in detail in Samphantharak and Townsend (2010a, chapter 5) and Townsend and Yamada (2008).

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. There are 369 households in the sample: 72 from Chachoengsao, 82 from Buriram, 93 from Lopburi, and 122 from Sisaket.

Table 1 presents descriptive statistics for ROA. Asset returns are high on average but also quite variable.

4. Results

Table 2 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 (11). 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 3 investigates the relationship between the estimates from the portfolio-choice and risk-sharing methods for measuring risk aversion. The estimates based on the risk-sharing method are from Chiappori et al. (2013). For each household i that has positive estimated return on assets,² we have two estimates of 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 the moment conditions for the risk-sharing method 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.³ We then use a Monte Carlo permutation test to see whether the correlation is statistically significantly different from zero.⁴ 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

²Recall that we cannot use the portfolio-choice method when the household's estimated return on assets is negative.

³Chiappori et al. (2013) show 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, although our two estimates should be positively correlated within each village if they are both accurate estimates of the true risk tolerance $1/\gamma_i$, comparing the levels of $1/\hat{\gamma}_i^{RS}$ and $1/\hat{\gamma}_i^{PC}$ is not worthwhile 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.

⁴The permutation test randomly reorders the list of $1/\hat{\gamma}_i^{PC}$ across households in 100,000 ways and computes the correlation of each reordered 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 reorderings — 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 reordering 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 reordered 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 reordered 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 reordered correlations that are less than $\text{corr}(1/\hat{\gamma}_i^{RS}, 1/\hat{\gamma}_i^{PC})$.

of the six villages in which the correlation is statistically significantly different from zero at least at the 10 percent 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.

Table 4 examines the relationship of risk tolerance to observed demographic characteristics of the household as measured in the initial round of the survey. As in Chiappori et al. (2013), we find little relationship between risk preferences and demographics.

5. Conclusion

This paper uses a model of household-level portfolio choice to measure the risk preferences of Thai households. Although preferences are measured with a substantial amount of noise, the estimates are correlated with alternative results based on full insurance in most villages. Thus, methods heretofore applied to data from industrialized countries with deep financial markets are also useful for understanding the behavior of households in a quite different economic environment.

Appendix

A1. Derivation of Portfolio-Choice Method

We show here how to obtain (5) from (3) and (4) 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 ,

$$\text{Cov}(A, B) = \text{Corr}(A, B) \sqrt{\text{Var}(A)} \sqrt{\text{Var}(B)}. \quad (\text{A1})$$

Hence, (3) can be written as

$$\begin{aligned} 1 = & \text{E} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right] \text{E}[R_{t+1}^k] \\ & + \text{Corr} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k \right] \sqrt{\text{Var} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right]} \sqrt{\text{Var} [R_{t+1}^k]} . \quad (\text{A2}) \end{aligned}$$

Substituting (4) into (A2) yields

$$1 = \text{E}[R_{t+1}^k] + \text{Corr} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k \right] \sqrt{\text{Var} \left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)} \right]} \sqrt{\text{Var} [R_{t+1}^k]} . \quad (\text{A3})$$

Then, rearranging terms in (A3), we obtain

$$\frac{\mathbb{E}[R_{t+1}^k] - 1}{\sqrt{\text{Var}[R_{t+1}^k]}} = -\sqrt{\text{Var}\left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}\right]} \text{Corr}\left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k\right]. \quad (\text{A4})$$

Since $-1 \leq \text{Corr}\left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}, R_{t+1}^k\right] \leq 1$, (A4) implies

$$\left| \frac{\mathbb{E}[R_{t+1}^k] - 1}{\sqrt{\text{Var}[R_{t+1}^k]}} \right| \leq \sqrt{\text{Var}\left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}\right]}, \quad (\text{A5})$$

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

Inequality (A5) 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

$$\left| \frac{\mathbb{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| \leq \sqrt{\text{Var}\left[\frac{\xi_{i,m(t+1)} \beta_i u'_i(c_{i,t+1}^*)}{\xi_{i,m(t)} u'_i(c_{it}^*)}\right]}. \quad (\text{A6})$$

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 — (A6) must hold with equality. (Otherwise, there would be a portfolio with the same variance as the household's portfolio but a higher expected return, contradicting the assumption that the household's portfolio is on the mean-variance frontier.) Applying assumption (b) — CRRA preferences — (A6) at equality becomes

$$\left| \frac{\mathbb{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \sqrt{\text{Var}\left[\beta_i \frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \left(\frac{c_{i,t+1}^*}{c_{it}^*}\right)^{-\gamma_i}\right]}. \quad (\text{A7})$$

Using (4), it is convenient to rewrite (A7) as

$$\left| \frac{\mathbb{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \frac{\sqrt{\text{Var} \left[\beta_i \frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \left(\frac{c_{i,t+1}^*}{c_{it}^*} \right)^{-\gamma_i} \right]}}{\mathbb{E} \left[\beta_i \frac{\xi_{i,m(t+1)}}{\xi_{i,m(t)}} \left(\frac{c_{i,t+1}^*}{c_{it}^*} \right)^{-\gamma_i} \right]}. \quad (\text{A8})$$

We define seasonally adjusted consumption $x_{it}^* = \xi_{i,m(t)}^{-1/\gamma_i} c_{it}^*$. Then (A8) can be written as

$$\left| \frac{\mathbb{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \frac{\sqrt{\text{Var} \left[\beta_i \left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right]}}{\mathbb{E} \left[\beta_i \left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right]}. \quad (\text{A9})$$

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

$$\Delta \ln x_{it}^* = \ln (x_{i,t+1}^*/x_{it}^*) \sim N(\mu_{dx}, \sigma_{dx}^2), \quad (\text{A10})$$

then

$$\ln (x_{i,t+1}^*/x_{it}^*)^{-\gamma_i} = -\gamma_i \ln (x_{i,t+1}^*/x_{it}^*) \sim N(-\gamma_i \mu_{dx}, \gamma_i^2 \sigma_{dx}^2). \quad (\text{A11})$$

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

$$\mathbb{E} \left[\beta_i \left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right] = \beta_i \mathbb{E} \left[\left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right] = \beta_i \exp \left[-\gamma_i \mu_{dx} + \frac{\gamma_i^2 \sigma_{dx}^2}{2} \right] \quad (\text{A12})$$

and

$$\text{Var} \left[\beta_i \left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right] = \beta_i^2 \text{Var} \left[\left(\frac{x_{i,t+1}^*}{x_{it}^*} \right)^{-\gamma_i} \right] = \beta_i^2 [\exp(\gamma_i^2 \sigma_{dx}^2) - 1] \exp(-2\gamma_i \mu_{dx} + \gamma_i^2 \sigma_{dx}^2). \quad (\text{A13})$$

Substituting (A12) and (A13) into the right-hand side of (A9) gives

$$\left| \frac{\text{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \frac{\beta_i \sqrt{\exp(\gamma_i^2 \sigma_{dx}^2) - 1} \exp(-\gamma_i \mu_{dx} + \gamma_i^2 \sigma_{dx}^2 / 2)}{\beta_i \exp\left[-\gamma_i \mu_{dx} + \frac{\gamma_i^2 \sigma_{dx}^2}{2}\right]} = \sqrt{\exp(\gamma_i^2 \sigma_{dx}^2) - 1}. \quad (\text{A14})$$

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,

$$\left| \frac{\text{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right| = \gamma_i \sqrt{\text{Var}[\Delta \ln x_{it}^*]} \quad (\text{A15})$$

and hence

$$\gamma_i = \frac{1}{\sqrt{\text{Var}[\Delta \ln x_{it}^*]}} \left| \frac{\text{E}[R_{i,t+1}^P] - 1}{\sqrt{\text{Var}[R_{i,t+1}^P]}} \right|. \quad (\text{A16})$$

References

- Breeden, Douglas T., 1979, “An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities,” *Journal of Financial Economics* 7(3), 265–296.
- Chiappori, Pierre-André, and Monica Paiella, 2011, “Relative Risk Aversion Is Constant: Evidence from Panel Data,” *Journal of the European Economic Association* 9(6), 1021–1052.
- Chiappori, Pierre-André, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M. Townsend, 2013, “Heterogeneity and Risk Sharing in Village Economies,” Research Department Working Paper 683, Federal Reserve Bank of Minneapolis.
- Cochrane, John, 2001, *Asset Pricing*, Princeton, NJ: Princeton University Press.
- Guiso, Luigi, and Monica Paiella, 2008, “Risk Aversion, Wealth, and Background Risk,” *Journal of the European Economic Association* 6(6), 1109–1150.
- 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.
- Lucas, Robert E., Jr., 1978, “Asset Prices in an Exchange Economy,” *Econometrica* 46(6), 1429–1446.
- 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 Pricing 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, UK: Cambridge University Press.

Samphantharak, Krislert, and Robert M. Townsend, 2010b, “Risk and Return in Village Economies,” manuscript, UCSD and MIT.

Townsend, Robert M., and Hiroyuki Yamada, 2008, “The Summary of the Monthly Household Survey of Thailand on Labor Issues,” manuscript, University of Chicago.

Table 1: Descriptive statistics.

Variable	mean	std. dev.	observations
return on assets	7.32	121.78	30,576

The unit of analysis is the household-month. Return on assets is in percentages.

Table 2: Tests for heterogeneity in risk preferences (portfolio-choice method).

village	households	risk aversion γ_i			risk tolerance $1/\gamma_i$		
		mean	χ^2	<i>p</i> -value	mean	χ^2	<i>p</i> -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>Lopburi</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). The column labeled “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 3: Correlation between estimated risk tolerance from risk-sharing and portfolio-choice methods.

village	corr.	<i>p</i> -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>Lopburi</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 reorder the list of $1/\hat{\gamma}_i^{PC}$ in 100,000 ways and compute the correlation of each reordered 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 reordered correlations.

Table 4: 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 are 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.