

Federal Reserve Bank of Minneapolis
Research Department

Heterogeneity and Risk Sharing in Village Economies*

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

Working Paper 683

Revised May 2012

ABSTRACT

We show how to use panel data on household consumption to directly estimate households' risk preferences. Specifically, we measure heterogeneity in risk aversion among households in Thai villages using a full risk-sharing model, which we then test allowing for this heterogeneity. There is substantial, statistically significant heterogeneity in estimated risk preferences. Full insurance cannot be rejected. As the risk sharing, as-if-complete-markets theory might predict, estimated risk preferences are unrelated to wealth or other characteristics. The heterogeneity matters for policy: Although the average household would benefit from eliminating village-level risk, less-risk-averse households who are paid to absorb that risk would be worse off by several percent of household consumption.

Keywords: Risk preferences; Heterogeneity; Complete markets; Insurance

JEL: 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. 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 measure heterogeneity in risk aversion among households running farm and non-farm enterprises in a developing country using a full risk-sharing model. From the literature on risk sharing, a household's risk aversion is identified up to scale 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 data we analyze are an unusually long monthly panel of households in villages in Thailand. We test the null hypothesis of full risk sharing in these data, incorporating our measure of heterogeneity in preferences, and find that the null is difficult to reject. This finding suggests that local institutions provide something close to a complete markets allocation. Our goal here is not to identify the specific institutions that aid in risk sharing, but we conjecture that they may include gifts between households as well as implicitly state-contingent informal or formal loans, such as those described in Udry (1994).¹

Further, we find that estimated risk tolerance is not significantly correlated with demographic variables or household wealth. The finding of no correlation between preferences and wealth is consistent, however, with the complete markets hypothesis and, since we are measuring relative risk tolerance, consistent with the finding of Chiappori and Paiella (2011)

¹Using the same data set that we analyze, Sripakdeevong and Townsend (2012) find that use of borrowing products is associated with lower coefficients of relative risk aversion. However, this relationship is limited to the subset of borrowers who do not roll over their loans. Other borrowers defer repayment when circumstances are bad by refinancing across lenders (typically informal lenders and the semi-formal village fund). The combined credit contracts resemble insurance, and these borrowers do not have significantly different risk aversion from non-borrowers.

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 13 percent — 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.

Our study is far from the first to measure heterogeneity in risk preferences. In development economics, efforts to measure risk aversion date at least to Binswanger (1980), who used experiments and hypothetical questions to measure the risk aversion of households in India. Many subsequent authors have used similar methods; recent examples from develop-

ing countries include Cole et al. (2012), Harrison et al. (2010), and Liu (forthcoming). The innovation in our work is that we estimate preferences from data on households' everyday behavior, rather than from their behavior in experiments that may or may not correspond to decisions the households would face if the experimenter were not present. (Our work is thus similar in spirit that of Chiappori and Paiella, 2011, who estimate Italian households' preferences from their portfolio choices.) However, we cannot directly compare the magnitude of our risk aversion estimates to those in the experimental literature because our estimates are identified only up to scale. We also contribute to a newer literature on tests of risk sharing with heterogeneous preferences. Schulhofer-Wohl (2011) and Mazzocco and Saini (2012) develop tests that do not require estimates of preferences but, as a result, require either large numbers of households or complex nonparametric methods. We provide an alternative test that uses our estimates of preferences to test for full insurance in a simple linear regression.

The paper proceeds as follows. In section 2, we lay out the theory underlying our methods for estimating preferences and for testing for full insurance. In section 3, we describe the Thai data. Section 4 presents the empirical results, and section 5 concludes.

2. Theory

In this section, we derive a method to estimate households' risk preferences based on measurements of risk sharing among households; we derive a test of full insurance based on our preference estimates; and we show how to estimate the welfare cost of aggregate risk in the villages in our data, as a function of households' risk preferences.

We assume that there is one consumption good, c , and 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 household-specific and aggregate states of nature up to date t , are represented by

$$\mathbb{E}_0 \left[\sum_{t=0}^T \beta_i^t \xi_{i,m(t)} \frac{[c_{it}^*(s^t)]^{1-\gamma_i}}{1-\gamma_i} \right], \quad (1)$$

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.

We 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:

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

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

A. Estimating Preferences

Let $C_{jt}(s^t)$ be the aggregate consumption available in village j at date t after history s^t .² Then, following Diamond (1967) and Wilson (1968), 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 $\lambda_{j(i),t}(s^t)$ 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 $\lambda_{j(i),t}(s^t)$ 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. To be concise, we henceforth let λ_{jt} denote this multiplier.

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 coef-

²We 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.

efficient 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. Under the maintained hypothesis of full insurance, the data must satisfy (4), and we can use this equation to estimate each household's risk preferences γ_i .³ 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 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

³In 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.

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, time 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.⁴ 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 equation (4) implies

$$\nu_{it} = \frac{1}{\gamma_i} \ell_{j(i),t} + \tilde{\epsilon}_{it}. \quad (5)$$

⁴The 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.

Since ϵ_{it} is uncorrelated across households, (5) implies that for any two households i and i' in the same village j ,

$$\mathbb{E}[\nu_{it}\nu_{i',t}] = \frac{1}{\gamma_i\gamma_{i'}}\mathbb{E}[\ell_{jt}^2], \quad i \neq i'. \quad (6)$$

As discussed above, risk aversion is identified only up to scale within each village; equation (5) would not change if, for any non-zero constant η_j specific to village j , we replaced γ_i with $\eta_j\gamma_i$ and ℓ_j with $\eta_j\ell_j$. Since the scale η_j is unidentified, we can normalize $\mathbb{E}[\ell_{jt}^2] = 1$. With this normalization, (6) reduces to

$$\mathbb{E}[\nu_{it}\nu_{i',t}] = \frac{1}{\gamma_i\gamma_{i'}}, \quad i \neq i'. \quad (7)$$

Equation (7) 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 (7) to one moment condition per household by summing over the other households $i' \neq i$, reducing our moment conditions to

$$\sum_{i' \neq i} \mathbb{E}[\nu_{it}\nu_{i',t}] = \frac{1}{\gamma_i} \sum_{i' \neq i} \frac{1}{\gamma_{i'}}. \quad (8)$$

Equation (8) 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.⁵ 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 $\sum_{i=1}^{N_j} \gamma_i > 0$.

The assumption that measurement error is uncorrelated across households is crucial for our results, because we identify household i 's risk aversion from how its consumption moves relative to aggregate consumption. If measurement error were correlated across households and the correlated measurement errors had the same proportional effect on all households — which would happen, for example, if the price index were measured with error — the estimated correlation of each household's consumption with aggregate consumption would be biased toward 1. In that case, all of our risk aversion estimates would be biased toward 1 and our test could fail to reject the null of identical preferences even when preferences are heterogeneous. Of course, under other assumptions on the measurement error, the bias could go in other directions.

⁵An 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.

B. 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

$$\ln c_{it} = \frac{\ln \alpha_i}{\gamma_i} + \frac{\ln \beta_i}{\gamma_i} t + \frac{\ln \xi_{i,m(t)}}{\gamma_i} + \frac{1}{\gamma_i} (-\ln \lambda_{j(i),t}) + b_j \ln income_{it} + \epsilon_{it}, \quad (9)$$

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, most of the literature uses a slightly different test and runs the regression

$$\ln c_{it} = a_i + d_{j(i),t} + b_j \ln income_{it} + u_{it}, \quad (10)$$

where d_{jt} represents the aggregate shock in village j at date t .⁶

The key difference between (9) and (10) is that (10) ignores heterogeneity in both risk and time preferences and absorbs the household-specific trends and seasonality into the aggregate shocks d_{jt} . Recently, Schulhofer-Wohl (2011) and Mazzocco and Saini (2012) have shown that tests based on (10) may be biased against the null of full insurance when risk preferences are heterogeneous, because (10) assumes aggregate shocks affect all households'

⁶For the sake of precision, we write d_{jt} rather than λ_{jt} in this equation; indeed, if $b_j \neq 0$, our model in the previous section is incorrect and d_{jt} need not be the Lagrange multiplier λ_{jt} .

consumption equally even though, under heterogeneous preferences, aggregate shocks have a larger effect on the consumption of less-risk-averse households. Both Schulhofer-Wohl (2011) and Mazzocco and Saini (2012) propose alternative tests that do not assume identical preferences. Here, we present another alternative test that also allows heterogeneity, based on our method for estimating risk preferences.

If we multiply both sides of (4) by γ_i , we find that efficient risk sharing implies

$$\gamma_i \ln c_{it} = \ln \alpha_i + t \ln \beta_i + \ln \xi_{i,m(t)} + (-\ln \lambda_{j(i),t}) + \gamma_i \epsilon_{it}. \quad (11)$$

Thus, although aggregate shocks do not affect all households' consumption equally, aggregate shocks do have an equal effect on consumption scaled by risk aversion. To test whether risk sharing is efficient, we test whether income is excluded from (11). That is, we run the regression

$$\hat{\gamma}_i \ln c_{it} = \ln \alpha_i + t \ln \beta_i + \ln \xi_{i,m(t)} + (-\ln \lambda_{j(i),t}) + b_j \ln income_{it} + \gamma_i \epsilon_{it} + u_{it}, \quad (12)$$

where

$$u_{it} = (\hat{\gamma}_i - \gamma_i) \ln c_{it} \quad (13)$$

and where, because we do not observe the household's actual risk aversion, we use the estimated risk aversion $\hat{\gamma}_i$ to construct the dependent variable. Given estimates of risk aversion, the regression in (12) is straightforward to estimate because the right-hand side includes only time dummies rather than the interaction of risk aversion with time dummies. Under the null hypothesis of full insurance, $\hat{\gamma}_i$ is a consistent estimator of γ_i and therefore \hat{u}_{it} converges

in probability to 0 for all i, t in the limit as $T \rightarrow \infty$. Since ϵ_{it} is independent of all variables in the model, it follows that the ordinary least squares estimator of b_j is consistent; thus, under the null of full insurance, the ordinary least squares (OLS) estimator of b_j should not be statistically significantly different from zero.

Thus, our test for full insurance is as follows. First, obtain risk preference estimates $\hat{\gamma}_i$. Second, regress $\hat{\gamma}_i \ln c_{it}$ on $\ln income_{it}$, an aggregate time dummy, and household-specific intercepts, trends and season dummies as shown in (12). Finally, test whether the estimated coefficient on $\ln income_{it}$ is zero. Because $\hat{\gamma}_i$ appears only in the dependent variable, we do not need to correct the point estimates or standard errors in (12) to account for the estimation of $\hat{\gamma}_i$ in a previous step.

C. 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 .⁷ 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

⁷By 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. If shocks are actually persistent, our results will underestimate the welfare costs of risk. However, as Schulhofer-Wohl (2008) notes, persistent shocks in a dynamic model would be similar to i.i.d. shocks with a large variance; thus, allowing for persistence might raise the magnitude of our welfare estimates but would not change the key result that heterogeneity in risk preferences reduces welfare costs compared with an economy where there is no heterogeneity.

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.⁸

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 fraction of consumption, is the value of k that solves

$$\hat{U}_{it}(k) = U_{it}^*. \tag{14}$$

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 as

$$k(\gamma_i) = 1 - \left(\sum_s \pi_s (p_s^*)^{-(1-\gamma_i)/\gamma_i} \right)^{\gamma_i/(1-\gamma_i)}, \tag{15}$$

⁸Since 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).

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 (15), 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, since our method for estimating preferences identifies households' preferences only up to scale, we cannot estimate the mean risk tolerance $\bar{\theta}$. Instead, we assume the mean risk tolerance is $\bar{\theta} = 1$, corresponding to logarithmic utility for the average household. When

we describe our results, we discuss how they would change if the mean risk tolerance were different.

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.⁹ Schulhofer-Wohl (2008) shows that $\hat{p}_\tau^*(\bar{\theta}) = \exp[\widehat{\ln p}_\tau^*(\bar{\theta})]$ is a consistent estimator of the price p_s^* for the state s that was realized at date τ , in the limit as the numbers of households and time periods go to infinity. We impose the normalization that $\sum_s \pi_s p_s^* m_s = 1$ by scaling the estimated prices such that $T^{-1} \sum_{\tau=1}^T \hat{p}_\tau^* \hat{m}_\tau = 1$.

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 (15):

$$\hat{k}(\gamma_i) = 1 - \left(\frac{1}{T} \sum_{\tau=1}^T (\hat{p}_\tau^*)^{-(1-\gamma_i)/\gamma_i} \right)^{\gamma_i/(1-\gamma_i)}. \quad (16)$$

⁹The 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.

The results in Schulhofer-Wohl (2008) imply that, conditional on the mean risk tolerance $\bar{\theta}$, $\hat{k}(\gamma_i)$ 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.

It is interesting to consider how the welfare estimates would change if we ignored heterogeneity. Schulhofer-Wohl (2008) shows that, generically, each household has strictly lower welfare costs if it lives in a heterogeneous-agent economy than if it is a representative agent in an economy where all agents have the same preferences. Furthermore, the welfare cost for a representative agent with risk aversion γ_i can be estimated by

$$\hat{k}^{rep}(\gamma_i) = 1 - \left(\frac{1}{T} \sum_{\tau} \hat{m}_{\tau}^{1-\gamma_i} \right)^{1/(1-\gamma_i)}. \quad (17)$$

When we turn to the data, we will compare the estimates allowing heterogeneity from (16) with estimates that incorrectly ignore heterogeneity from (17) and show that allowing heterogeneity leads to quantitatively important reductions in the estimated welfare cost of aggregate risk.

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 $2\hat{k}^* - \sum_{q=1}^Q k_q/Q$.¹⁰

3. Data

We apply the estimation methods described in the previous section to the households in the Townsend Thai Monthly Survey. The survey has relatively high frequency over many years, providing us with a relatively long time series on consumption fluctuations. 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).

A. The Townsend Thai Monthly Survey and Sample Selection

The Townsend Thai Monthly Survey is an ongoing 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

¹⁰Our 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.

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 are 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 data are missing in any month. Overall, our restrictions eliminate 29 percent of the households in the initial survey, leaving a final sample of 505 households: 141 from Chachoengsao, 102 from Buriram, 122 from Lopburi, and 140 from Sisaket. (Of the households dropped from the data, about half are dropped because their consumption data are missing in some months, and about half for other reasons.)

B. Construction of Consumption Variable

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 biweekly 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 regional level from the Thailand Ministry of Commerce. Although we realize that inflation in each village could differ from regional inflation, we must rely on the regional statistics because we do not have a reliable village-level price index at the time of writing this paper.

Table 1 presents descriptive statistics for household consumption. Mean per capita real consumption is 1607.2 Thai baht per month (in 2007 baht). According to the Penn World Table (Heston et al., 2011), purchasing power parity in 2007 was 15.39 baht per U.S. dollar, so on average, households in the sample live on the equivalent of about U.S. \$3.43 per person, per day.

C. Construction of Income Variable

The income data come from the underlying survey instruments which distinguish several potential income sources: crops, livestock, fish and shrimp, and wage earning. For each source, both revenues and expenses are measured to calculate a net profit. Further, income and all other variables are cross-checked with each other via the creation of standard financial accounts, treating the households as if they were firms; see Samphantharak and Townsend

(2010a). These accounts include a statement of income, a statement of cash flow, and a balance sheet. In our analysis, we use the accrual notion of income, in which expenses are booked at the time of sale of product, rather than the the cash-flow notion in which expenses are booked when paid. (We have conducted robustness checks to ensure our results do not depend on using the accrual notion of income.)

4. Results

This section presents the results from applying our methods for testing efficiency, estimating preferences and estimating the welfare costs of aggregate risk to the Thai data.

A. Test of Efficient Risk Sharing

Table 2 presents the tests of efficient risk sharing based on (12).¹¹ The coefficient on income is statistically significant at the 5 percent level in only one of the 16 villages. When we estimate a common coefficient on income across all villages, we gain statistical power but nonetheless only barely reject the null of full insurance. We note that the evidence against full insurance is weak even though we have not allowed for nonseparability between consumption and leisure, which would lead our test to over-reject full insurance.¹² In results not reported here, we failed to reject full insurance when we ignored preference heterogeneity and estimated (10). In addition, in table 3 we show tests of the hypothesis that there is efficient risk sharing both within and across the four study villages in each province and fail

¹¹For 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.

¹²Classical measurement error in income would lead our test to under-reject full insurance. However, unless the signal-to-noise ratio is small — which we think is unlikely given the detailed nature of the survey questionnaire — this bias is small.

to reject full risk sharing in any of the four provinces.¹³ 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.

B. Estimation of Risk Preferences

Table 4 presents the tests of the null hypothesis of identical risk preferences, based on the GMM overidentification statistic for moment conditions (8). We reject the null of identical preferences at the 5 percent level in four of the 16 villages and at the 10 percent level in eight 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 examines the relationship of risk tolerance to observed demographic characteristics of the household in the initial round of the survey. We find little evidence that estimated risk preferences are related to demographics. There is a positive, statistically significant relationship between risk tolerance and the head's age. Education, net wealth, and the numbers of men, women, and children in the household are not associated with risk tolerance. These patterns persist whether or not we include village fixed effects in the regressions. 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, but theory has little to say about whether the initial endowment, and thus wealth, will be related to preferences. Recall also that, under

¹³For these tests, we re-estimate households' preferences under the maintained hypothesis of full risk sharing among the four study villages in each province.

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.¹⁴

C. Welfare Costs of Aggregate Risk

In table 6 and figure 2, we turn to estimating the welfare costs of aggregate risk. Table 6 shows the willingness to pay to eliminate aggregate risk for a household with the mean risk tolerance, under our benchmark assumption that the mean risk tolerance is 1. Welfare losses for households with the mean risk tolerance are on the order of 1 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 percent confidence intervals for the mean household's welfare loss include zero in all but two villages. The estimates also show that allowing heterogeneity matters dramatically for the results: the welfare costs are typically two to three times as large if we assume all households have identical risk preferences.

Figure 2 shows the importance of heterogeneity for understanding the welfare cost of risk. The figure contains a separate graph for each village. The graph shows the welfare cost of aggregate risk, on the vertical axis, as a function of a household's risk tolerance, on the horizontal axis. In each village, the more risk tolerant a household is, the smaller its welfare cost, as evidenced by the downward slope of the welfare cost function as we move to the right

¹⁴Note, 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.

on the graphs. Furthermore, households that are sufficiently close to risk neutral have welfare gains from aggregate risk: The welfare cost is less than zero. For example, in village 9 in Sisaket, some very risk-averse households — those with risk tolerance close to zero — have welfare losses from aggregate risk equivalent to about 3 percent of consumption; however, households with risk tolerance of 5, equivalent to a coefficient of relative risk aversion of 0.2, have welfare *gains* of about 4 percent of consumption.

5. Conclusion

This paper uses the benchmark of full risk sharing, but incorporates heterogeneity in risk preferences, and presents a novel way to test the null hypotheses of full insurance and of homogeneous risk preferences. The first hypothesis, perfect risk sharing, cannot be rejected even in pooled data with some power, while the second hypothesis, homogeneity, is soundly rejected (even though common unobserved measurement error would bias this test toward failing to reject the null). Our method uses an unusually long panel data set on households in each of four diverse regions of an emerging market country, Thailand, treating villages and then counties as the risk-sharing unit. We also use the data to quantify the welfare impact of (counterfactual) insurance against village-level aggregate shocks. Overall, the welfare costs of aggregate risk, and the gains from insurance, are less than they would be under homogeneous preferences. Further, relatively risk-tolerant households would actually lose under policy interventions that remove risk, because when risk is present these households benefit from providing *de facto* insurance to their more risk-averse neighbors.

References

- Binswanger, Hans P., 1980, "Attitudes toward Risk: Experimental Measurement in Rural India," *American Journal of Agricultural Economics* 62(3), 395–407.
- 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.
- Cochrane, John H., 1991, "A Simple Test of Consumption Insurance," *Journal of Political Economy* 99(5), 957–976.
- Cole, Shawn, Xavier Giné, Jeremy Tobacman, Robert Townsend, Petia Topalova, and James Vickery, 2012, "Barriers to Household Risk Management: Evidence from India," manuscript, Harvard Business School.
- Diamond, Peter A., 1967, "The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty," *American Economic Review* 57(4), 759–776.
- 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.
- Harrison, Glenn W., Steven J. Humphrey, and Arjan Verschoor, 2010, “Choice Under Uncertainty: Evidence from Ethiopia, India, and Uganda,” *Economic Journal* 120(543), 80–104.
- Heston, Alan, Robert Summers, and Bettina Aten, 2011, “Penn World Table Version 7.0,” Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.
- Liu, Elaine, forthcoming, “Time to Change What to Sow: Risk Preferences and Technology Adoption Decisions of Cotton Farmers in China,” *Review of Economics and Statistics*.
- 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, 2012, “Testing Efficient Risk Sharing with Heterogeneous Risk Preferences,” *American Economic Review* 102(1), 428–468.
- Samphantharak, Krislert, and Robert M. Townsend, 2010, *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, 2011, “Risk and Return in Village Economies,” manuscript, UCSD and MIT.

- 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, 2011, “Heterogeneity and Tests of Risk Sharing,” *Journal of Political Economy* 119(5), 925–958.
- Sripakdeevong, Parit, and Robert M. Townsend, 2012, “Informal Networks and Shadow Banking,” unpublished project document, MIT.
- Townsend, Robert M., 1994, “Risk and Insurance in Village India,” *Econometrica* 62(3), 539–591.
- Udry, Christopher, 1994, “Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria,” *Review of Economic Studies*
- Wang, Jiang, 1996, “The Term Structure of Interest Rates in a Pure Exchange Economy with Heterogeneous Investors,” *Journal of Financial Economics* 41(1), 75–110.
- Wilson, Robert, 1968, “The Theory of Syndicates,” *Econometrica* 36(1), 119–132.

Table 1: Descriptive statistics.

Variable	mean	std. dev.
real total consumption per capita	1607.2	2928.7
ln(real total consumption per capita)	7.07	0.68
real net income	12,351.2	45,157.4
ln(real net income)	6.94	3.50
adult men	1.43	0.83
adult women	1.61	0.80
children	1.47	1.23
head's age	52.1	13.6
highest education (years)	8.3	3.7
net wealth (millions of baht)	2.26	12.4
households	505	
monthly observations	42,420	

The table reports descriptive statistics for household consumption and demographics. 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 regional Consumer Price Index (base year 2007). Demographics and net wealth are measured in the initial survey.

Table 2: Tests of efficient risk sharing at the village level.

village	coeff.	std. err.	<i>p</i> -value	obs.	HH	<i>R</i> ²
<i>Chachoengsao</i>						
2	-0.0313	0.0490	0.527	3444	41	0.039
4	-0.0949	0.8099	0.907	3192	38	0.020
7	0.1197	0.4118	0.773	2520	30	0.038
8	0.0976	0.1115	0.388	2688	32	0.057
<i>Buriram</i>						
2	0.2156	0.2992	0.476	3024	36	0.029
10	0.0170	0.0882	0.851	1092	13	0.092
13	0.0703	0.0492	0.166	2184	26	0.128
14	0.1845	0.1105	0.107	2268	27	0.122
<i>Lopburi</i>						
1	-0.1059	0.1984	0.597	2688	32	0.038
3	1.6043	1.2067	0.199	1764	21	0.074
4	0.2643	0.3172	0.410	3192	38	0.031
6	0.7524	0.5867	0.209	2604	31	0.025
<i>Sisaket</i>						
1	0.1297*	0.0529	0.019	3276	39	0.085
6	-0.0462	0.0874	0.600	3612	43	0.031
9	0.0493	0.0406	0.232	3108	37	0.088
10	0.2012	0.1563	0.213	1764	21	0.063
<i>pooled</i>						
-	0.1662*	0.0828	0.045	42420	505	0.025

The table reports the effect of idiosyncratic income shocks on the product of log consumption and estimated risk aversion. 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 regional 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 the product of estimated risk aversion and log consumption on household fixed effects, time fixed effects, and log income (12); “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 percent level.

Table 3: Tests of efficient risk sharing at the province level.

province	coeff.	std. err.	p -value	obs.	HH	R^2
Chachoengsao	0.3398	0.2424	0.163	11,844	141	0.007
Buriram	0.1127	0.0626	0.075	8,568	102	0.009
Lopburi	-0.3633	0.1950	0.065	10,248	122	0.008
Sisaket	0.1354	0.0821	0.101	11,760	140	0.010
pooled	0.0604	0.0833	0.469	42,420	505	0.007

The table reports the effect of idiosyncratic income shocks on the product of log consumption and estimated risk aversion. 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 regional Consumer Price Index. Each row reports a separate regression using data from one province. Column labeled “coeff.” reports the coefficient on log income in an OLS regression of the product of estimated risk aversion and log consumption on household fixed effects, time fixed effects, and log income (12); “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 provinces and interacts time effects with province effects to allow different aggregate shocks by province.

Table 4: Tests for heterogeneity in risk preferences.

village	χ^2	d.f.	p -value
<i>Chachoengsao</i>			
2	50.68	40	0.120
4	51.91	37	0.053
7	54.93	29	0.003
8	53.19	31	0.008
<i>Buriram</i>			
2	48.03	35	0.070
10	8.72	12	0.726
13	34.19	25	0.104
14	44.15	26	0.015
<i>Lopburi</i>			
1	42.23	31	0.086
3	29.32	20	0.082
4	46.76	37	0.130
6	47.98	30	0.020
<i>Sisaket</i>			
1	48.57	38	0.117
6	46.3	42	0.299
9	37.81	36	0.387
10	23.09	20	0.284
<i>pooled</i>			
-	667.86	489	1.22×10^{-7}

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 (8) 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 regional Consumer Price Index.

Table 5: Association between household demographics and estimated risk tolerance.

	estimated risk tolerance						
	<i>A. Without village fixed effects</i>						
adult men	0.010						0.008
	(0.007)						(0.007)
adult women		0.006					0.001
		(0.006)					(0.006)
children			0.004				0.005
			(0.003)				(0.003)
head's age				0.001*			0.001*
				(0.000)			(0.000)
highest education					-0.000		-0.001
					(0.001)		(0.001)
net wealth (millions of baht)						-0.000	-0.000
						(0.000)	(0.000)
joint signif. p -value							0.0618
R -squared	0.009	0.003	0.003	0.016	0.000	0.001	0.027
	<i>B. With village fixed effects</i>						
adult men	0.009						0.007
	(0.005)						(0.007)
adult women		0.006					0.003
		(0.006)					(0.006)
children			-0.003				-0.002
			(0.003)				(0.003)
head's age				0.001*			0.001
				(0.000)			(0.000)
highest education					0.000		-0.000
					(0.001)		(0.001)
net wealth (millions of baht)						-0.000	-0.000
						(0.000)	(0.000)
joint signif. p -value							0.0315
R -squared	0.134	0.130	0.129	0.139	0.128	0.128	0.144
Observations	505	505	505	505	505	505	505

The table reports the association between demographic variables and households' estimated preferences. The unit of observation is the household. 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 percent level.

Table 6: Estimated welfare cost of aggregate risk for a household with the mean risk tolerance, by village.

village	willingness to pay to eliminate aggregate risk			
	allowing heterogeneous preferences		assuming identical preferences	
	estimate	95% c.i.	estimate	95% c.i.
<i>Chachoengsao</i>				
2	0.8%	(-0.3%,1.7%)	1.7%	(-0.2%,3.0%)
4	0.9%	(-0.8%,2.3%)	2.3%	(-1.7%,4.4%)
7	0.5%	(-1.1%,2.6%)	2.6%	(-1.0%,4.7%)
8	1.4%	(0.1%,2.1%)	2.1%	(0.0%,3.5%)
<i>Buriram</i>				
2	1.5%	(-1.0%,3.3%)	3.3%	(-1.0%,6.2%)
10	2.4%	(-3.4%,8.0%)	8.0%	(-0.3%,14.1%)
13	1.5%	(-0.4%,2.5%)	2.5%	(-0.2%,4.4%)
14	2.3%	(-0.1%,4.2%)	4.2%	(0.3%,7.1%)
<i>Lopburi</i>				
1	0.4%	(-0.8%,1.4%)	1.4%	(-1.1%,2.7%)
3	1.1%	(-1.8%,2.8%)	2.8%	(-2.0%,5.3%)
4	0.4%	(-0.5%,0.8%)	0.8%	(-0.7%,1.9%)
6	0.4%	(-0.6%,1.2%)	1.2%	(-0.5%,2.5%)
<i>Sisaket</i>				
1	0.3%	(-0.9%,1.8%)	1.8%	(-1.1%,3.3%)
6	1.5%	(-0.5%,3.4%)	3.4%	(-0.4%,5.8%)
9	1.6%	(0.1%,2.5%)	2.5%	(0.4%,4.2%)
10	2.4%	(-1.8%,6.6%)	6.6%	(0.7%,10.9%)

The table reports the estimated willingness to pay to eliminate aggregate risk in each village for a household with the mean risk tolerance, assuming the mean risk tolerance is 1. Willingness to pay is reported as a percentage of mean consumption. Results under “allowing heterogeneous preferences” show the estimates from our heterogeneous-preferences formula (16); results under “assuming identical preferences” show the estimates from the formula (17) that assumes all households have the mean risk tolerance. “Estimate” is the bootstrap bias-corrected point estimate of willingness to pay and “95% c.i.” is the 95 percent 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.

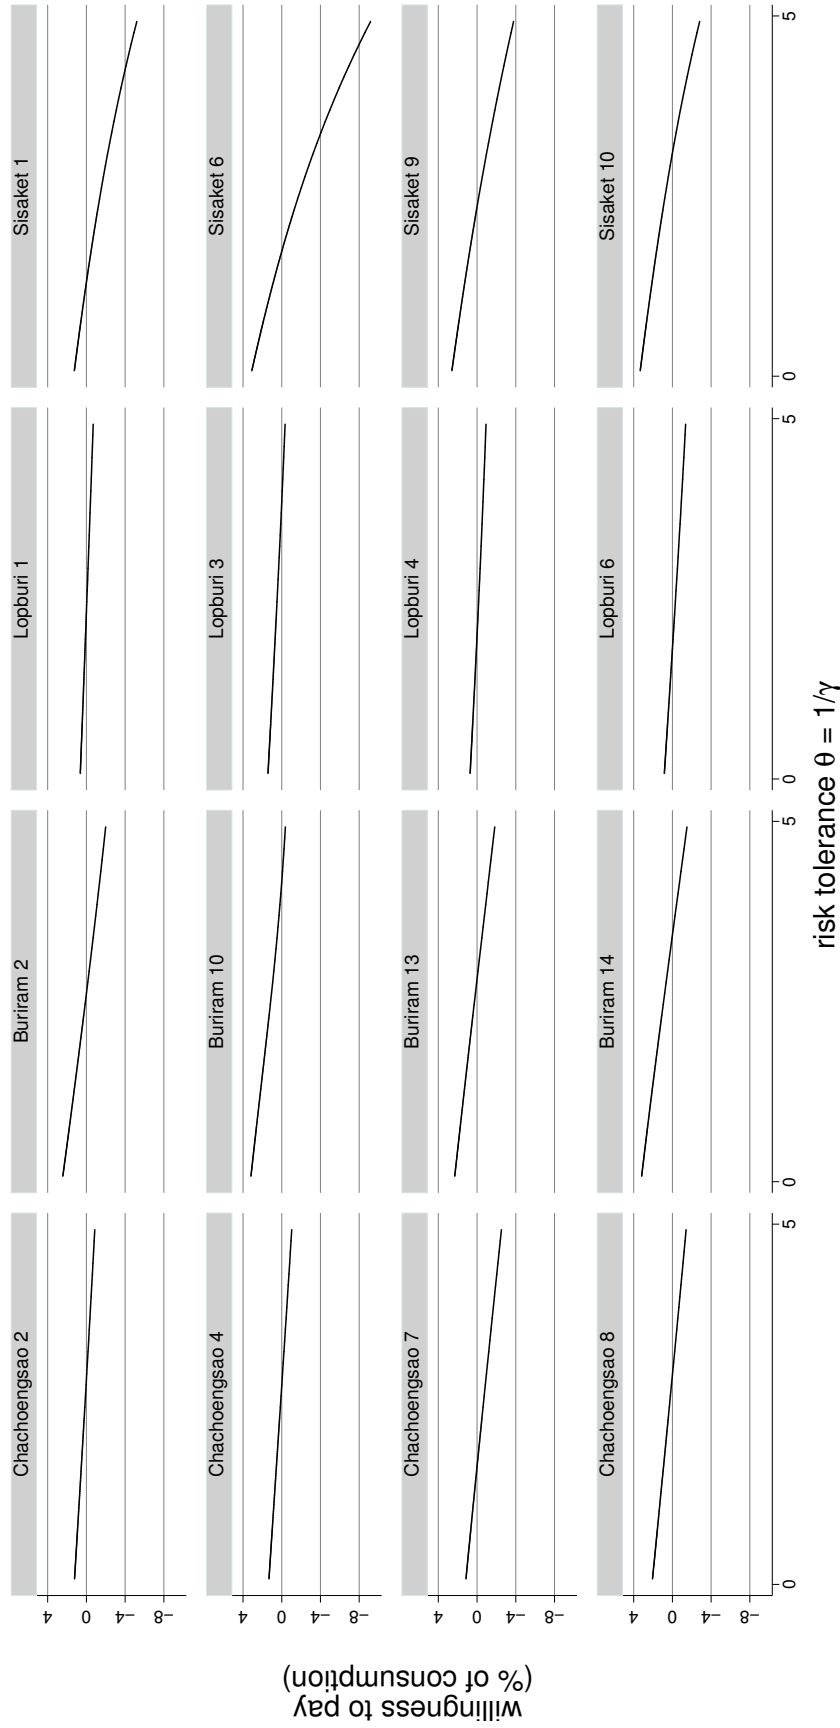


Figure 2: Welfare costs of aggregate risk.

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 loss from aggregate risk and is willing to pay to eliminate risk; negative numbers mean the household has a welfare gain from aggregate risk. Graphs assume the mean risk tolerance in each village is 1.