Risk-Matching and Diversification in Joint Liability Borrowing Groups

PRELIMINARY AND INCOMPLETE
Draft Prepared for Applied Dynamics Workshop

Christian Ahlin*

May 23, 2006

Abstract

In a lending environment of limited liability with unobserved risk types, individual lending contracts involve safe borrowers cross-subsidizing risky ones, and thus may drive safe borrowers from the market. Ghatak (1999, 2000) shows that joint liability contracts can improve efficiency because they embed an effective discount for safe borrowers, given that borrowers match homogeneously by risk-type. We extend borrower heterogeneity to a second dimension, correlated risk. We show that borrowers will anti-diversify risk within groups, in order to lower chances of facing liability for fellow group members. We directly test homogeneous sorting by risk and intra-group diversification of risk using data on Thai joint liability borrowing groups. Univariate tests involve calculating the between-group contributions of borrower differences in a village as well as rank correlations between borrower characteristics and group identity. These numbers are compared to those arising from all possible combinations of village borrowers into groups of the observed sizes. Non-parametric tests show evidence of both homogeneous sorting by risk and risk anti-diversification within groups. Multivariate analysis is carried out using Fox's (2006) maximum score estimator, which chooses payoff function parameters to maximize the frequency with which payoffs of observed groupings dominate payoffs of feasible alternatives. These results support the univariate results by suggesting that homogeneous matching along both dimensions is profitable. They also allow estimation of a key model parameter and hence calculation of the average reduction in effective joint liability due to correlated risk within groups.

*I thank Yanqin Fan, James Foster, Jeremy Fox, Xavi Gine, Alex Karaivanov, Ethan Ligon, Andrew Newman, Moto Shintani, Rob Townsend, Jonathan Zinman, and seminar participants at Vanderbilt and NEUDC 2005 for valuable input. Jisong Wu provided excellent research assistance and computational algorithms. All errors are mine.

1 Introduction

Extending financial markets to the world's poorest has been a top agenda item in economic development in recent years. Specifically, starting with lenders such as the Grameen Bank of Bangladesh in the 1970's, millions of households around the developing world have received loans intended to facilitate small-scale enterprise. The small size of these loans and the fact that they are often not backed by collateral have given rise to experimental lending techniques, including group lending. Several economists have rationalized the popularity of group lending as a remedy for informational problems in a non-secured lending context.

Consider one such context in which borrowers' distributions of returns have identical means but vary in riskiness. A lender that cannot observe risk will offer all borrowers the same terms, but effectively will charge less to risky borrowers, who fail more often, than to safe borrowers. Thus there is cross-subsidization from safe to risky borrowers, and this may cause safe borrowers to exit the market. This type of adverse selection in credit markets was analyzed by Stiglitz and Weiss (1981).

Add to this context a communal tightness, that is, that borrowers know each other well, including each other's riskiness. Ghatak (1999, 2000) shows that group lending contracts can be used in this context to overcome the bank's lack of information, partially or fully. Here a group lending contract means that the lender requires borrowers to sort into pairs and makes each borrower liable for his partner's loan. The foundational result is that groups match homogeneously in risk. Every borrower prefers a safer partner, to reduce expected liability; but safe borrowers prefer them more, since they succeed more often and liability is only operative when one succeeds.

Given homogeneous risk-matching, group lending mitigates the adverse selection problem. The lender now can screen borrowers, since safe borrowers have safe partners and so are more willing to accept joint liability (Ghatak, 2000).[1] The lender can also pool borrowers, simply offering all groups the same standard contract with joint liability (Ghatak, 1999). Even though contract terms are uniform, there is effectively a built-in discount for safe borrowers: their partners are safer and thus the joint liability clause is less costly for them in expectation. This discount can draw into the market some or all safe borrowers who would have been excluded under individual loans.

The pooling result is appealing in practical terms. It implies that even a very passive lender that offers only a standard joint liability contract is giving implicit discounts to safe borrowers, relative to the case of using individual liability contracts. This may help explain the popularity of group lending in micro-credit.[2]

However, the lynchpin in this analysis is the homogeneous risk-matching of the borrowing groups. Without it, there is no discount for safe borrowers operating through joint liability. A main contribution of this paper is to test directly for homogeneous risk-matching among borrowing groups in Thailand.[3]

A second main point of this paper points to a potentially more ambiguous side of volun-

[1] Gangopadhyay et al. (2005) refine the earlier results of Ghatak (2000).

[2] Direct and indirect evidence for adverse selection in these lending contexts can be found in Ahlin and Townsend (2006, forthcoming).

[3] This is the first question on the microfinance empirical research agenda Morduch (1999, p. 1586) lays out: "Is there evidence of assortative matching through group lending as postulated by Ghatak (1999)?".

tary group formation. In particular, suppose there is a second dimension of heterogeneity among borrowers, *correlated* risk. For example, assume there are two aggregate shocks in a village – for example, rainfall and currency risk – and households vary in their degrees of exposure to each. In this context, do groups form homogeneously with respect to the shock they are affiliated with; that is, do they diversify or anti-diversify?

The main theoretical result is that groups sort homogeneously in both dimensions: they match with similar risk-types, and among those, with partners exposed to the same risk. The rationale is straightforward. Since execution of joint liability happens only when borrower returns are dissimilar, higher correlation of group risk lessens the incidence of liability. Groups anti-diversify in order to minimize liability for their partners.

Correlated risk restricts the amount of joint liability that the lender can offer. In the model the *effective* joint liability rate turns out to equal the stated joint liability rate discounted by within-group correlation. Consider the extreme case where all groups have perfectly correlated returns. Joint liability is then never executed and group lending reduces to individual lending. An analysis of contracting in this environment of correlated risk is beyond the scope of this paper, but the analysis here does suggest that some aspects of voluntary group formation can work against efficiency.

We test empirically whether groups seem to be diversified or not, in terms of risk. The data for these and the risk-matching tests come from the Townsend Thai dataset, which includes information on borrowing groups from the Bank for Agriculture and Agricultural Cooperatives (BAAC). The BAAC is the predominant rural lender in Thailand and offers group lending contracts to borrowers with little or no collateral. Borrowers form groups autonomously, though the BAAC may veto one or more proposed members.

One feature of the two datasets to be used is that where possible, and indeed often, two groups from the same village are surveyed. Considering the village as the relevant universe for group formation, we can assess whether groups appear homogeneous by comparing across two groups in the same village. Specifically, for each village and ordered variable, we calculate rank correlations and inequality (e.g. variance) decompositions into between- and within-group components. For each village and non-ordered variable, we calculate fractionalization decompositions into between- and within-group components. We put these calculations in perspective by repeating them for all possible combinations of the village borrowers into groups of the size observed. The results from the actually observed groups can then be mapped into a *sorting quantile* reflecting how homogeneous or heterogeneous group formation is relative to all possibilities (holding group size and borrowing pool fixed).

For any given variable, villages can be found at both ends of the spectrum – homogeneous sorting (high sorting quantile) and heterogeneous sorting (low sorting quantile). Sorting quantile means and medians (across villages) suggest predominant tendencies. Using the fact that if matching is completely random over a certain variable, then village sorting quantiles are drawn from a uniform distribution, we use Kolmogorov-Smirnov (KS) tests of the sample cdf of village quantiles against the uniform distribution.

We find direct evidence for homogeneous matching by risk, statistically significant but quantitatively moderate. The KS tests reject heterogeneous sorting at the 5% level. Median sorting quantiles range from 59% to 62%,[4] suggesting that risk homogeneity is an important,

[4]That is, the median village has groups that are more homogeneously sorted than about 60% of all

but not overwhelming, consideration in group formation. It is therefore plausible that group lending gives a moderate effective discount to safe borrowers.

We also find evidence for anti-diversification within groups. This is seen in three ways: occupational homogeneity within groups, greater coincidence of bad years for income within groups than across groups, and clustering of similar income shocks within groups. Of these, the occupation homogeneity is most striking, at least in the approximately half of villages that have any occupational diversity along which groups can sort. The median village has a sorting quantile of 84-91%. For coincidence of worst years for income, the median sorting quantiles are 64-66%; for clustering of income shocks, 66%. Statistical significance in the KS tests holds for each measure but worst-year coincidence. Overall, evidence suggests that groups are responding to incentives not to diversify within groups.

We next turn to Fox's (2006) maximum score estimator, which allows for direct estimation of certain payoff function parameters (signs and/or magnitudes), including whether given borrower characteristics are complements or substitutes in the payoff function. It chooses payoff function parameter values that maximize the frequency with which observed groupings yield higher payoffs than feasible, unobserved alternatives.

We first directly test complementarity vs. substitutability of risk-types in the original Ghatak models (i.e. without correlated risk). The estimator finds for complementarity, supporting homogeneous matching by risk-type. Similar tests confirm the univariate anti-diversification results.

Next we turn to multivariate analysis, which can alleviate concerns that sorting appears homogeneous along one dimension merely due to correlation between this and another more important dimension. In a reduced form setup, coefficients support homogeneous matching along both risk-type and correlated risk dimensions; however, subsampling does not allow rejection of the opposite hypothesis in either case (though both are relatively close).

We also take a more structural approach that uses the model's exact payoff function and allows estimation of a key model parameter. In this case, again homogeneous matching along both dimensions is supported, but only in the correlated risk dimension can the opposite hypothesis be rejected, at the 10% level. The model parameter estimates also allow calculation of the average percent reduction in effective joint liability due to correlated risk within groups. We perform this calculation, but little can be said with certainty due to wide confidence intervals around the estimates.

The paper is organized as follows. The model setup and results on homogeneous sorting are in section 2. Data are described and key variables defined in section 3. Section 4 presents the methodology behind the nonparametric univariate tests, as well as the results. Section 5 presents several estimations using Fox's maximum score estimator.

2 Theory: Sorting when returns may be correlated

2.1 Model Setup

There is a measure one continuum of potential borrowers, each endowed with a project requiring one unit of capital. Every project has expected value E. However, agents and

possible combinations.

their projects differ in risk. Agent i's project yields gross returns of Y_i (i.e. "succeeds") with probability p_i and yields zero gross returns (i.e. "fails") with probability $1 - p_i$. The higher p_i, the lower agent i's risk, measured by second-order stochastic dominance. Population risk-types are distributed over $(\underline{p}, \overline{p})$, where $0 < \underline{p} < \overline{p} < 1$, according to density function $g(p)$ and cdf $G(p)$.

A lender requires potential borrowers to form groups of size two, each member of which is jointly liable for the other. Specifically, contracts are assumed to take the following form (as in Ghatak, 1999, 2000). A borrower who fails pays the lender nothing, due to limited liability. A borrower who succeeds pays the lender gross interest rate r. A borrower who succeeds and whose partner fails pays an additional joint liability payment q. Thus, a borrower of type p_i who matches with a borrower of type p_j has expected payoff

$$\Pi_{ij} = E - p_i r - p_i(1 - p_j)q, \tag{1}$$

assuming the borrowers' returns are uncorrelated. Note that

$$\frac{\partial^2 \Pi_{ij}}{\partial p_i \partial p_j} = q > 0.$$

That is, risk-types are complements in the joint payoff function and homogeneous matching by risk is the equilibrium outcome, as Ghatak has shown.

To this environment we add the potential for correlated returns. Assume the success of each project is determined by one or more unobserved latent variables. There is a continuum of such latent variables, one Z_i for every potential borrower. The Z_i's are independently and uniformly distributed on $[0, 1]$. In the zero-correlation case, agent i's success is determined solely by his individual shock Z_i. Specifically, agent i is assumed to succeed if and only if $Z_i \leq p_i$. The probability of success is thus $F(p_i)$, where F is the uniform distribution function for the unit interval; and $F(p_i) = p_i$.

In addition to the Z_i's, we assume there are two village-aggregate shocks, A and B, both uniformly distributed on $[0, 1]$. A and B are assumed to be distributed independently of each other and the Z_i's.[5]

Each villager is assumed to be exposed to risk from shock A, shock B, or neither. A fraction $\alpha \in (0, 1)$ of villagers are exposed to risk from shock A: with probability $\kappa \in (0, 1)$, their success is determined by A and otherwise their success is determined by their own Z_i. That is, given risk-type p_i, with probability κ success occurs if $A \leq p_i$ and with probability $1 - \kappa$ success occurs if $Z_i \leq p_i$. Note that since A and Z_i are independently and uniformly distributed on $[0, 1]$, the unconditional probability of success is p_i: $\kappa F(p_i) + (1-\kappa)F(p_i) = p_i$. A fraction $\beta \in (0, \alpha]$ of villagers are exposed to risk from shock B: with probability κ, their success is determined by B and otherwise their success is determined by their own Z_i. The remaining $1 - \alpha - \beta$ (which could be zero) are exposed to neither aggregate shock. One can think of α and β as capturing the extensiveness, and κ the intensiveness, of aggregate shock exposure. Aggregate shock affiliation is assumed uncorrelated with unconditional probabilities of success, the p_i's.

[5] The homogeneous sorting results of this paper hold for any amount of correlation ρ_{AB} between A and B, with $\rho_{AB} = 1$ interpreted as having one shock only.

How does this uncertainty structure affect payoffs? As argued in previous work (see Ahlin and Townsend, forthcoming and 2006), there is a unique, one-parameter class of joint output distributions given unconditional probabilities of success (p_i, p_j):

	j Succeeds (p_j)	j Fails $(1 - p_j)$
i Succeeds (p_i)	$p_i\, p_j + \epsilon$	$p_i(1 - p_j) - \epsilon$
i Fails $(1 - p_i)$	$(1 - p_i)p_j - \epsilon$	$(1 - p_i)(1 - p_j) + \epsilon$

The case of $\epsilon = 0$ is the case of independent returns considered by Ghatak. A positive (negative) ϵ induces positive (negative) correlation between borrower returns. Borrower i's payoff to borrowing with borrower j under this generalized joint distribution, given ϵ, is

$$\Pi_{ij} = E - p_i r - [p_i(1 - p_j) - \epsilon]q = E - p_i r - p_i(1 - p_j)q + \epsilon q. \tag{2}$$

This is similar to payoff 1. The one addition is that higher correlation (higher ϵ) raises payoffs by directly lowering the probability of asymmetric outcomes and, thus, of paying q.

Groups fit into one of two categories. The first category contains borrowers both exposed to the same shock, i.e. both A-risk or both B-risk. The second category includes all other combinations. It is clear that for groups in the second category, $\epsilon = 0$. This is because the shocks each borrower is exposed to – idiosyncratic and perhaps also aggregate – are independent. Returns are thus uncorrelated.

One can show that for groups in the first category (both A-risk or both B-risk),[6]

$$\epsilon = \kappa^2 \cdot \min\{p_i(1 - p_j), p_j(1 - p_i)\}. \tag{3}$$

One can also calculate the correlation coefficient between borrower returns, ρ_{ij}, to be

$$\rho_{ij} = \kappa^2 \cdot \overline{\rho}(p_i, p_j), \tag{4}$$

where

$$\overline{\rho}(p_i, p_j) \equiv \min\left\{\sqrt{\frac{p_i(1 - p_j)}{p_j(1 - p_i)}}, \sqrt{\frac{p_j(1 - p_i)}{p_i(1 - p_j)}}\right\}$$

is the maximum correlation possible between projects with unconditional success probabilities (p_i, p_j). The expression for $\overline{\rho}(p_i, p_j)$ makes clear that only if $p_i = p_j$ can projects i and j be perfectly correlated, i.e. can $\overline{\rho}(p_i, p_j) = 1$. If $p_i \neq p_j$, with positive probability one project succeeds and the other fails, so perfect correlation is impossible, i.e. $\overline{\rho}(p_i, p_j) < 1$.

Thus the parameter κ^2 has a clear interpretation as the correlation between two similarly exposed borrowers, expressed as a fraction of the maximum feasible correlation between them. Under homogeneous matching in risk-type p, κ^2 is exactly the correlation between similarly exposed borrowers.

Summarizing the above, the payoff of borrower i matching with borrower j can be written

$$\Pi_{ij} = E - rp_i - qp_i(1 - p_j) + q\kappa^2 \cdot \min\{p_i(1 - p_j), p_j(1 - p_i)\}\, C_{ij}\,, \tag{5}$$

[6] To see this, assume both borrowers are A-risk. With probability κ^2, both borrowers' success is determined by shock A. Thus, both succeed iff $A < p_i$ and $A < p_j$, i.e. with probability $\min\{p_i, p_j\}$. With probability $1 - \kappa^2$, borrowers' outcomes are independent since at least one's success is determined by his idiosyncratic shock, so both succeed with probability $p_i p_j$. Adding $\kappa^2 \min\{p_i, p_j\}$ with $(1 - \kappa^2)p_i p_j$ gives $p_i p_j + \epsilon$.

where C_{ij} equals one if borrowers i and j are exposed to the same risk and equals zero otherwise. Similarly, the joint surplus of match (i, j), that is the sum of payoffs of both borrowers, equals

$$JS_{ij} = 2E - (r+q)(p_i + p_j) + 2q\, p_i p_j + 2q\kappa^2 \cdot \min\{p_i(1-p_j), p_j(1-p_i)\}\, C_{ij}\,. \tag{6}$$

2.2 Equilibrium Matching

Our interest is in whether groups form so as to diversify risk or to amplify risk within the group. For example, will A-risks match with B-risks or with fellow A-risks? This question need not directly bear on risk-sharing: households may share risk with other households regardless of whether they are in the same joint liability borrowing group. The issue instead is how groups form in response to joint liability contracts.

Recall that all groups must fit into one of two categories. Groups in the first category have similarly exposed borrowers (both A-risk or both B-risk) while groups in the second category have diversified borrowers. We first show homogeneous risk-matching within each of the two categories of groups.

Since the second category of group involves $\epsilon = 0$, Ghatak's results hold directly and homogeneous risk-matching is the only equilibrium within this category. For the first category,

$$\frac{\partial^2 \Pi_{ij}}{\partial p_i \partial p_j} = q(1-\kappa^2) > 0$$

if $p_i \neq p_j$ and is $+\infty$ if $p_i = p_j$. Thus, risk-types are complements in the payoff function, so homogeneous sorting by risk-type is the equilibrium outcome.[7]

We next show equilibrium anti-diversification within groups. Given homogeneous risk-matching and using payoff 2 and expression 3, the payoff for a group of risk-type p in the first category (both A-risk or both B-risk) is

$$E - pr - p(1-p)q(1-\kappa^2). \tag{7}$$

The payoff of a borrowing group of type p in the second category is given by expression 1, since there is no correlation. Inspection makes clear that payoffs in the second category are less than those in the first category, for a given risk-type.

The second category will thus have zero measure of both A-risk and B-risk borrowers in equilibrium. If there were a positive measure of A-risk borrowers in this category, nearly all of them could leave their partners and match together, increasing their payoffs while leaving their partners no worse off (their payoffs remain at 1). The same argument rules out B-risk borrowers in the second category.

Thus, in equilibrium all A-risk borrowers match together, all B-risk borrowers match together, and all borrowers not exposed to either A or B match together. The intuition is simple: borrowers anti-diversify within groups to lower their chances of facing liability for their partner.

[7] An extension of the more direct argument of Ghatak (1999, Lemma 1 and Proposition 1) can alternatively be applied here to show homogeneous sorting by risk-type.

Note that correlated borrowing groups (both A-risk or both B-risk) face an *effective* joint liability rate of $q' \equiv q(1-\kappa^2)$; compare payoffs 1 and 7. Effective joint liability is thus lowered in direct proportion to the degree of correlation within the group. In a sense, correlation works against the lender by limiting the degree to which it can effectively use joint liability in contracts. Since a natural upper bound for q is r,[8] effective joint liability q' is thus capped at $r(1-\kappa^2)$, which approaches zero as correlation approaches one.

3 Data and Variable Descriptions

Data description and environment. The data come from the Townsend Thai data base, in particular from a large cross section of l92 villages, conducted in May l997, and a smaller re-survey conducted in late 1999/early 2000. The larger, "baseline" survey includes two provinces each from two contrasting regions of Thailand – central and northeast. Within each province, twelve subcounties, or tambons, were chosen. Within each tambon, a cluster of four villages was selected. In each village as many borrowing groups of the Bank for Agriculture and Agricultural Cooperatives (BAAC) as possible were interviewed, up to two. In all we have data on 262 groups, 200 of which are one of two groups representing their village. Each group designates an official leader, and the leader responded to questions on behalf of the group.

The "resurvey" data were collected from a random subset of the same villages. Included are data on 87 groups, 14 of which are the only groups in their village, 70 of which are one of two groups interviewed from the same village, and 3 of which are one of three groups interviewed from the same village. Besides sample size, there are two key differences between the resurvey data and the baseline survey data. First, in the resurvey individual group members respond to questions on their own behalf, up to five per group and on average 4.5; in the baseline survey, the group leader responds for all members. Second, several resurvey questions were designed explicitly to measure income risk and correlatedness, while the baseline survey captures mainly general information on the individual-member level.

One limitation imposed by the data is that there is little overlap across surveys in the groups interviewed, so merging the two datasets is fruitless. Consequently, multivariate analysis (discussed in section 5) must be survey-specific. The main specifications will include prob-high, shock, and worst-year, but will exclude occupation since it comes from a different survey.

The BAAC is a government-operated development bank in Thailand. It was established in l966 and is the primary formal financial institution serving rural households. The BAAC allows smaller loans to be backed only with social collateral in the form of joint liability. To borrow in this way, a borrower must belong to an official BAAC borrowing group that faces explicit liability for the loan. That is, in the event of a group member's default on a loan, the BAAC may opt to follow up with the delinquent borrower *or* other group members in search of repayment. Thus loans underwritten by a BAAC group do not in principle require land or other physical collateral, only the promise that individual members be jointly liable.

[8] If $q > r$, the group might pretend both to have succeeded when one failed, to limit group liability to $2r$ rather than $r + q$. See Gangopadhyay et al. (2005).

Groups typically have between five and fifteen members; about 15% are larger. Group formation is mostly at the discretion of the borrowers themselves. Typically, groups are born when borrowers propose a list of members to the BAAC, and the BAAC then approves some or all members. The BAAC seems to use its veto power sparingly: only about 12% of groups in the baseline survey report that the BAAC struck members from the list.[9] We know of no case where the BAAC adds members to a list or forms a group unilaterally. Thus, while the BAAC has some say about group formation, it is largely left to the borrowers themselves.

Variable descriptions. Our basic empirical strategy will involve comparing across groups within villages to determine whether homogeneity is greater within groups than across groups. To do so, measures of risk and of correlatedness are necessary. Our measure of risk takes the model quite literally, and has been used in other tests of the adverse selection environment (see Ahlin and Townsend, forthcoming). Group members were asked in the resurvey what their income would be in the coming year if it were a good year (Hi), what their income would be if it were a bad year (Lo), and what they expected their income to be (Ex). Assuming that income can take only one of two values, Hi and Lo, the probability of success, **prob** or prob-high, works out to be

$$prob = \frac{Ex - Lo}{Hi - Lo},$$

using the fact that $prob * Hi + (1 - prob) * Lo = Ex$. Another measure of risk, less directly related to the model, is the **coefficient of variation** of income.[10] Based on the same projected income distribution, this works out to be

$$\sigma/Ex = \sqrt{Hi/Ex - 1}\sqrt{1 - Lo/Ex}.$$

Since data on explicit risk exposure are lacking, correlatedness is measured in three indirect ways. First, we use information on **occupation** from the baseline survey, since exposure to aggregate shocks is likely to be relatively similar within occupations. Occupation is coded in one of twenty categories, including "rice farmer", "corn farmer", "construction worker in village", "mechanic", etc. Since the BAAC targets agricultural workers exclusively (at the time of the survey), there is less variation in occupation than would be ideal: only 46 out of 100 villages with two groups have *any* occupational heterogeneity among the two groups' borrowers. Nonetheless, we can analyze tendencies to diversify among these 46 villages.

Second, we use timing of bad income years, **worst-year**. Specifically, the resurvey asks borrowers which year of the past two was worse for household income: "one year ago", "two years ago", or "neither". If borrowers are exposed to the same aggregate shocks, bad income years are more likely to coincide; thus coincidence of bad years can proxy anti-diversification. One drawback of this measure is its coarseness, as it maps past performance into one of three categories.

Third, we calculate a direct measure of this year's household income **shock**, from the resurvey. One ingredient of this measure is next year's expected income, Ex, mentioned

[9]This is in response to a free-form question about how original members were determined when the group was founded.

[10]The coefficient of variation equals the standard deviation normalized by the mean.

above. The second is a very detailed compilation of realized business and farm income for the just-completed year, Inc. The income shock is then measured as $Shock = (Inc - Ex)/Ex$, i.e. the percent deviation of income from next year's expected value. This measure is most clearly interpretable as an income shock if we imagine households draw i.i.d. income realizations from a fixed distribution over time. Then Ex is exactly mean income, and $Shock$ is this year's realized random component of income (as a percent of mean income). If incomes are not stagnant over time but growing at the same expected rate across households, the measure would merely have to be adjusted by a constant to continue to capture the income shock. However, if household incomes are growing at different rates, then $Shock$ captures not only the income shock but differential growth rates of income. We view this as the main drawback of this measure: within-group homogeneity may imply sorting aimed at anti-diversifying risk, or it may imply sorting based on income growth rates.

Flaws notwithstanding, the three measures of correlatedness should combine to give a suggestive picture of group propensities to diversify or anti-diversify.

4 Univariate Methodology and Results

4.1 Univariate Methodology

Decompositions and rank correlations. Homogeneous sorting can be detected by comparing across groups within villages. Two general approaches are used. One is a decomposition of differences between village borrowers into within-group and between-group components. A dominant between-group component suggests homogeneous sorting. The second approach is to calculate rank correlations between a focal variable and an arbitrary group index. A high absolute-value rank correlation suggests homogeneous sorting.

More specifically, consider data on ordered variable x from two groups in village v, M and N, of respective sizes m and n: $M = (x_1, ..., x_m)$ and $N = (x_{m+1}, ..., x_{m+n})$. We calculate a variance decomposition of $X = (x_1, ..., x_{m+n})$ into between-group and within-group components. The same can be done for a number of decomposable inequality measures – we also do so with Theil's second measure of inequality, that is the mean log deviation.[11] The disadvantage of this measure is that it fails to work with non-positive values. To illustrate the approach, consider a village with group data $M = (2, 5, 6, 8)$ and $N = (1, 4, 7, 9)$. Both variance and mean log deviation would register zero-percent between-group inequality, since the group means are the same. However, $M' = (1, 2, 5, 6)$ and $N' = (4, 7, 8, 9)$ would attribute 44% of overall variance to between-group differences and 30% of the overall mean log deviation to between-group differences, reflecting the fact that groups appear to sort somewhat homogeneously by variable x.

One can also calculate rank correlations between the data X and a group index variable, y, where for example $y_1 = ... = y_m = 1$ and $y_{m+1} = ... = y_{m+n} = 2$. Rank correlations are useful in the context of sorting because they use only the ordinality (rankings) of the data and not the cardinality. That is, they are invariant with respect to any data transformation that leaves the relative rankings unchanged.

[11] The mean log deviation equals the log of (the arithmetic mean divided by the geometric mean). For a decomposition, see Foster and Shneyerov (2000).

Two rank correlation measures are used, Spearman's rho and Kendall's tau$_b$.[12] Since the group index is arbitrary, and the higher value could be assigned to either group M or group N, we take the absolute value of the rank correlation. This restricts the value to be in $[0, 1]$, as is the case in the inequality decompositions. For example, consider groups $M = (2, 5, 6, 8)$ and $N = (1, 4, 7, 9)$. Both rank-correlation measures give zero correlation between $(2, 5, 6, 8, 1, 4, 7, 9)$ and group index $(1, 1, 1, 1, 2, 2, 2, 2)$, indicating no evidence of homogeneous sorting. However, groups $M' = (1, 2, 5, 6)$ and $N' = (4, 7, 8, 9)$ register correlations of 57% under Kendall's tau$_b$ and 65% under Spearman's rho, reflecting somewhat homogeneous sorting. The correlations would be the same but negative if group M' were indexed by 2 and N' by 1; however, the absolute value of the correlation is used, ensuring indexing choice makes no difference.

A separate technique is needed for *non-ordered* categorical variables, such as occupation. We employ a decomposition approach, where the measure to be decomposed is analogous to an inequality measure in non-ordered space: fractionalization.[13] It is the probability that two randomly selected members of a population differ with respect to the (categorical) variable in question.[14] If the share of a population that fits into category j is denoted s_j, fractionalization F can be calculated as

$$F = 1 - \sum_j s_j^2.$$

For example, if in village v group M has occupations (Corn,Rice,Corn,Corn) and group N has occupations (Rice,Rice,Rice,Corn), then occupational fractionalization among borrowers village v is 50%.

Ahlin (2005) provides a decomposition of fractionalization into between-group and within-group components. Specifically, he shows that total fractionalization is a convex combination of within-group and between-group fractionalization, where between-group fractionalization is defined as the probability two randomly selected members of a population differ, conditional on being from different groups, and within-group fractionalization is the probability two randomly selected members of a population differ, conditional on being from the same group. The combination weights have to do with the number and sizes of groups. In the example of $M =$ (Corn,Rice,Corn,Corn) and $N =$ (Rice,Rice,Rice,Corn), 62.5% of total fractionalization is attributable to the between-group component. If instead $M' =$ (Corn,Rice,Rice,Corn) and $N' =$ (Corn,Rice,Rice,Corn), this number would be 50%.

Sorting quantiles. To move toward a statistical test for homogeneous sorting, the results from each technique are next placed in context. As above, consider data $X = (x_1, ..., x_{m+n})$ from two groups in village v, M and N, of respective sizes m and n. We form all possible combinations of the $m + n$ borrowers into two groups of respective sizes m and n and perform the same calculation – inequality decomposition, rank correlation, or fractionalization decomposition – on each one. The observed village grouping can then be assigned a

[12] Formulas can be found in Gibbons and Chakraborti (2003, pp. 419-420, 422-423). Kendall's tau$_b$ allows for ties; for Spearman's rho, we use the same rank for tied observations, the midrank.

[13] Fractionalization along ethno-linguistic and other dimensions has been used by a number of authors in analyses of economic growth. See Mauro (1995), Easterly and Levine (1997), and Alesina et al. (2003).

[14] Thus it is closely related to the gini coefficient, which is (half) the expected difference in incomes (relative to mean income) between two randomly selected members of a population. In both cases, the sampling is with replacement.

"sorting quantile" (or sorting quantile range, if there are ties) based on where its calculated value falls relative to this universe of possibilities. In this way, every village, variable, and technique is assigned a value (or range) in $[0, 1]$, with higher numbers representing greater homogeneity in sorting and lower numbers representing more heterogeneous sorting.

To illustrate, consider again a village with groups $M = (2, 5, 6, 8)$ and $N = (1, 4, 7, 9)$. There are $\binom{8}{4} = 70$ ways to sort eight borrowers into two groups of size four. Of these seventy combinations, sixty four register higher between-group inequality while six (including the observed combination) register exactly the same, i.e. zero between-group inequality, using either variance or mean log deviation. Thus this village is somewhere between the 0th and 8.6th percentiles in terms of group homogeneity; its sorting quantile range is $[0, 8.6]$. The somewhat wide range reflects the fact that there are ties and that the sample is relatively small. Consider groups $M' = (1, 2, 5, 6)$ and $N' = (4, 7, 8, 9)$. There are four combinations of borrowers higher, sixty two combinations lower, and four combinations tied, in terms of between-group variance. The village's sorting quantile range is thus $[88.6, 94.3]$. The same numbers apply for mean log deviation. Recall that 30% of mean log deviation across all borrowers in this village is attributable to between-group differences while 44% of the overall variance is. Translating these numbers in quantile ranges shows both that this may be considered a large amount of between-group differences, relative to the distribution of characteristics, though the decomposition numbers may seem more moderate; and that the variance and mean log deviation paint a more comparable picture than the raw decompositions might have indicated.[15]

Similar analysis produces sorting quantiles based on the rank correlation measures. The village with $M = (2, 5, 6, 8)$ and $N = (1, 4, 7, 9)$ gets a sorting quantile range of $[0, 11.4]$, based on either rank correlation measure. The village with $M' = (1, 2, 5, 6)$ and $N' = (4, 7, 8, 9)$ gets a sorting quantile range of $[88.6, 94.3]$, just as with the inequality decompositions. In general, however, the quantile ranges are often wider with rank correlations than with variance decompositions since they use less information and consequently result in ties more frequently.

The same approach can be used with fractionalization decomposition of non-ordered, categorical variables. If $M =$ (Corn,Rice,Corn,Corn) and $N =$ (Rice,Rice,Rice,Corn), thirty six combinations have less, two combinations have greater, and thirty two combinations have the same between-group contribution to overall fractionalization. Thus the village's sorting quantile range would be $[51.4, 97.1]$. A village with $M' =$ (Corn,Rice,Rice,Corn) and $N' =$ (Corn,Rice,Rice,Corn) would have thirty six combinations greater and thirty four combinations tied in terms of between-group fractionalization. Its quantile range would be $[0, 51.4]$. These quantile ranges are admittedly wide, and heterogeneous matching is likely to occur even by chance. More data and a greater number of categories would tend to shrink the ranges.

Thus for a given technique and variable, each village is assigned a sorting quantile or quantile range. Technically, every village will have a quantile range, though the range may be narrow; and the quantile may be thought of as the midpoint of this range. A higher sorting quantile reflects more homogeneous sorting (as measured by the given technique), while a lower sorting quantile reflects more heterogeneous sorting. One can then interpret villages

[15] However, the measures do not always give rise to the same sorting quantiles.

with quantiles above the 95th as exhibiting homogeneous sorting at the 5% confidence level, for example.

A nonparametric test. Rather than test sorting village by village, however, we combine villages in a single test (per variable and technique) of the overall tendency to sort homogeneously. Each village's sorting quantile is treated as a draw from the same distribution, and this distribution is compared using the Kolmogorov-Smirnov (KS) test to a benchmark distribution. An advantage of this approach is that it is non-parametric and requires no distributional assumptions.

The benchmark comparison distribution is the one that would obtain if sorting with respect to the given variable were completely random in all villages: the uniform distribution on $[0, 1]$. To see that the uniform is the appropriate comparison, consider the case of a large number of borrowers in a village, no two combinations of which (into two groups of the given sizes) result in a tie using the given technique. If each of the N, say, possible combinations is equally likely, then each $1/N$-quantile is equally likely to be observed. That is, the sorting quantile is drawn from the uniform distribution (approximately, with the difference getting arbitrarily small as N increases).

With smaller numbers of borrowers and, especially, with ties, the uniform distribution becomes less appropriate. This is because the appropriate distribution is a step function approximation of the uniform, and the steps are larger the lower is the number of combinations N (a number which varies by village) and the more ties there are. For robustness to this, we focus on the quantile ranges rather than the quantiles themselves. We assign the village not a single quantile (e.g. the midpoint of the range), but an equal chance of equaling any quantile in the quantile range, i.e. a random variable distributed uniformly over the quantile range. For example, if a village's sorting quantile is between the 92nd and 96th percentiles, we consider it not be the 94th percentile but uniformly distributed between the 92nd and 96th percentiles.

With this interpretation, the correct benchmark distribution is *exactly* the uniform regardless of the number of combinations N and the number of ties. To see this, let there be $K \leq N$ different values arising when the given technique is performed on the N combinations, with values $v_1 < v_2 < ... < v_K$. (Ties involve $K < N$.) Also, let n_i be the number of combinations that give rise to value v_i and N_i be the number of combinations that give rise to any value $v \leq v_i$, with $N_0 \equiv 0$; then $N_i = \sum_{k=1}^{i} n_i$ and $N_K = N$. If sorting is completely random, than each of the N combinations of borrowers is equally likely to obtain. With probability $p_i \equiv n_i/N$ the realized combination will result in value v_i, with quantile distributed uniformly over range $[\frac{N_{i-1}}{N}, \frac{N_i}{N}]$. The distribution of quantiles z is then

$$f(z) = \sum_{i=1}^{K} p_i \int_{\frac{N_{i-1}}{N}}^{\frac{N_i}{N}} \frac{1}{\frac{N_i - N_{i-1}}{N}} \, dz = \sum_{i=1}^{K} \frac{n_i}{N} \int_{\frac{N_{i-1}}{N}}^{\frac{N_i}{N}} \frac{N}{n_i} \, dz = \sum_{i=1}^{K} \int_{\frac{N_{i-1}}{N}}^{\frac{N_i}{N}} dz = \int_{0}^{1} dz.$$

Thus, as long as the distribution function across villages is constructed not using exact quantiles, but rather uniform distributions over quantile ranges, the correct comparison regardless of sample size and number of ties is the uniform. This is the approach we take. Weighting each village equally, we construct the sample cdf from the quantile ranges. For example, consider two villages, with respective sorting quantile ranges $[25, 30]$ and $[80, 90]$. The sample cdf using quantile ranges would be flat at zero until $(0.25, 0)$, then travel linearly upward from $(0.25, 0)$ to $(0.3, 0.5)$, then be flat at one half until $(0.8, 0.5)$, then travel linearly upward from $(0.8, 0.5)$ to $(0.9, 1)$, and finally remain flat at one.[16]

One final issue has to do with village weighting in the KS tests. Note that if all borrowers in a village have the same value for a given variable, then there is no inequality (or fractionalization), so the between-group component is not defined and the village must be dropped. This village borrower homogeneity happens not infrequently for categorical variables, such as occupation. Thus in our baseline approach, all villages with any differences are given equal and positive weights, while villages with no differences are implicitly given zero weight. A more continuous approach would be to weight each village in proportion to the amount of overall heterogeneity that exists within it. Hence, we also report results for categorical variables from weighting each village's observation in proportion to the total fractionalization in the village.[17]

4.2 Univariate Results

Sorting by risk-type. The probability of achieving the high income realization, **prob**, is a close analog to the risk-type variable in the theory and is thus the focus of our empirical tests for homogeneous risk-matching.[18] The sample cdfs of village sorting quantile ranges for *prob* based on inequality decompositions, constructed as described in the previous section, are graphed in Figure 1. Based on these quantile ranges and the variance measure, the mean (median) village is more homogeneously sorted than 57% (62%) of all possible combinations of borrowers into groups of the observed sizes. The random-matching benchmark, the uniform, is graphed as a dotted line. Using a one-sided KS test, we reject at the 5% level the hypothesis of heterogeneous sorting, that is, that the true distribution is first-order stochastically dominated by the uniform.

The distribution, including means and medians, looks similar when using mean log deviation instead of variance decompositions. However, the KS test p-value is significantly larger. The reason is due to the much smaller sample size, 15 instead of 32. This comes from the breakdown of the mean log deviation measure when there are zeros, a feature of the data from at least one respondent in 17 of the villages.

The sample cdf of village sorting quantile ranges for *prob* based on rank correlations is shown in Figure 2. Both Spearman's rho and Kendall's tau_b give mean (median) sorting percentiles of 58% (59%), and reject heterogeneous sorting at a 5% confidence level based on the KS test. Overall, the data point to statistically significant but not quantitatively overwhelming homogeneous risk-matching.

[16]The sample cdf using the midpoints of the quantile ranges would be flat at zero until a spike from $(0.275, 0)$ to $(0.275, 0.5)$, then be flat at one half before spiking from $(0.85, 0.5)$ to $(0.85, 1)$, and finally remain flat at one. This sample cdf has greater deviations on both sides of the uniform than the one we construct.

[17]For example, consider again two villages with respective sorting quantile ranges $[25, 30]$ and $[80, 90]$. Further imagine the two villages register fractionalization numbers of 0.75 and 0.25, respectively. Then the *weighted* sample cdf using quantile ranges would be flat at zero until $(0.25, 0)$, then travel linearly upward from $(0.25, 0)$ to $(0.3, 0.75)$, then be flat at one half until $(0.8, 0.75)$, then travel linearly upward from $(0.8, 0.75)$ to $(0.9, 1)$, and finally remain flat at one. That is, the first village would receive three times as much weight as the second since its overall fractionalization is three times as high.

[18]Ahlin and Townsend (forthcoming) find direct evidence for adverse selection using this measure.

A second measure of risk, though in a way not as closely related to the theory, is the **coefficient of variation** of projected income, described in section 3. Sorting tendencies based on this variable and inequality decompositions are graphed in Figure 3. According to variance decomposition, the mean (median) village is more homogeneously sorted than 63% (72%) of all possible combinations. The KS test rejects heterogeneous sorting at the 5% level. Mean log deviation decompositions give similar results, though the small sample size (due to reasons discussed above) lowers statistical significance of the KS test.

However, when judged by the rank correlations there is less evidence for homogeneous sorting by coefficient of variation; see Figure 4. The means and medians drop to 59% and 55-56%, respectively, and the KS tests come somewhat close but fail to reject heterogeneous sorting. While the coefficient of variation measure gives weaker results, we view it as auxiliary to the *prob* measure, and somewhat supportive. Overall, the data give solid evidence for a degree of homogeneous risk-matching.

Sorting by correlated risk. We next examine diversification within groups. **Occupation**al diversification, based on decomposition of occupational fractionalization, is graphed in Figure 5. The results overwhelmingly reject occupational diversification. The mean (median) village is more homogeneously sorted by occupation than 75% (84%) of all group formations. These numbers increase to 79% and 91% when villages are weighted by the overall amount of occupational heterogeneity (measured by fractionalization) in the village. Admittedly, only 46 out of the 100 villages with two groups have *any* occupational heterogeneity; but where heterogeneity exists, it clearly does not appear to be used to diversify groups. Using weighted and unweighted cdfs, the KS tests reject occupational diversification at the 1% level.

Next, consider the **worst-year** measure. This is a categorical variable which can be treated either as non-ordered or as ordered, with both "one year ago" and "two years ago" considered closer to "neither" than they are to each other. Fractionalization decompositions from the non-ordered case are reported in Figure 6 and rank correlation results from the ordered case in Figure 7.[19] The approaches give similar results, yielding unweighted and weighted means (medians) of 59-60% (65-66%) in the case of the decompositions and slightly lower means (medians) of 58% (64%) in the case of the rank correlations. Both come fairly close but fail to reject diversification reflected in non-coincidence of bad income years.

Finally, we consider coincidence of income **shock**s, measured by the percent deviation of this year's realized income from next year's expected income. Rank correlations are presented in Figure 8.[20] The rank correlation measures yield means (medians) of 59-60% (66%) and reject diversification at the 10% level. That is, income shocks are relatively coincident within groups, consistent with the view that groups anti-diversify so as to increase correlated risk.

Overall, the univariate tests paint a picture of homogeneous group composition along dimensions of unconditional risk and of correlated risk exposure, just as the theory predicts.

[19] Results from Spearman's rho are nearly identical to those from Kendall's tau_b, so are not reported.

[20] Variance decomposition yielded similar results: mean 58%, median 65%, KS p-value 0.067. Mean log deviation decomposition is omitted because this measure is undefined if there is one zero or negative number and nearly all villages have at least one such response.

5 Multivariate Methodology and Results

Univariate analysis may suggest one variable to be a key sorting dimension when in reality this variable is merely correlated with a second variable that is the true key sorting dimension. For example, one could interpret the evidence for homogeneous risk matching as driven by 1) borrowers' desires to match based on correlated risk, not unconditional risk per se, and 2) a correlation between unconditional risk and correlated risk. The latter could arise if some aggregate shocks are more volatile than others. In this case, matching based on similar shock exposure would induce homogeneous matching based on unconditional risk also.

A multivariate approach that isolated the partial effect of each variable would help alleviate these concerns. For this purpose, we turn to the maximum score estimator of Fox (2006). This estimator identifies some parameters in a joint surplus function (sum of group payoffs) up to scale. It does so by contemplating feasible, unobserved combinations of agents and choosing parameters that most frequently give these unobserved combinations less joint surplus than observed combinations. Thus the estimator exploits the idea that observed groupings should be optimal relative to feasible unobserved groupings; this is exactly the theoretical approach taken by Ghatak (1999, 2000) and in this paper.[21]

Specifically, theory predicts that given observed borrowing groups (i, j) and (i', j') within the same pool of borrowers, and joint surplus function JS_{xy}, in equilibrium

$$JS_{ij} + JS_{i'j'} \geq JS_{i'j} + JS_{ij'}. \tag{8}$$

If this condition did not hold, the groups would have sorted differently and been better off. The estimator works by choosing coefficients to maximize the score, the number of inequalities of the form 8 that are true. In our context, the inequalities come from all one-for-one borrower swaps across groups in the same village.[22] For example, if we have data on five borrowers in each of two groups in the same village, there are $5 \times 5 = 25$ one-for-one swaps to consider. Theory would also justify other kinds of inequalities, for example those arising from a one-borrower transfer or from two-for-two swaps. We do not use transfers because they change group size, an issue beyond the scope of this paper. At any rate, Fox (2006) shows consistency of the estimator when using a only subset of implied inequalities.

A preliminary structural test. We use the estimator in several ways. First, consider the original setting of Ghatak in which there is no correlation between borrower returns. Joint surplus from borrowing group (i, j) is then, from equation 6,

$$JS_{ij} = 2E - (r + q)(p_i + p_j) + 2q\, p_i p_j\,. \tag{9}$$

Using this in the key inequality condition 8 gives

$$\beta\,(p_i p_j + p_{i'} p_{j'}) \geq \beta\,(p_i p_{j'} - p_{i'} p_j)\,, \tag{10}$$

where $\beta = 2q$. Fox's maximum score estimator can be used to estimate the coefficient β, but only up to scale, i.e. the sign, since the inequality holds for any positive β or for any

[21] That is, there are no search frictions and utility is transferable.

[22] An implicit assumption in the choice of these inequalities is that it is feasible (if not desirable) to match with any other household in the village, a relatively good assumption for these areas of Thailand.

negative β. Thus the estimator can be used to tell whether risk-types are complements or substitutes in the payoff function (since the payoff function cross-partial is proportional to q), which is exactly the key to homogeneous sorting by risk-type. Note that in this example and generally, all terms in the joint surplus function (e.g. 9) that do not involve interactions between borrower characteristics drop out of condition 8; thus coefficients on those terms cannot be estimated.

Risk-type p is measured by *prob*, discussed in section 3. Since groups in the data are rarely, if ever, of size two, equation 10 must be modified to allow for this. Accordingly, if borrower i is from group M and i' from group N, then p_j is set to the average p in group M excluding borrower i and $p_{j'}$ is set to the average p in group N excluding borrower i'. Denote these averages as $p_{M\setminus i}$ and $p_{N\setminus i'}$, respectively.

If there are V villages indexed by v, and each village v has two groups, M_v and N_v, the estimate of the sign of β is then

$$max_{\beta\in\{-1,1\}} \sum_{v=1}^{V} \sum_{i\in M_v} \sum_{i'\in N_v} \beta \cdot 1\{p_i\, p_{M_v\setminus i} + p_{i'}\, p_{N_v\setminus i'} > p_i\, p_{N_v\setminus i'} + p_{i'}\, p_{M_v\setminus i}\}.$$

The estimate[23] turns out to be positive, supporting risk-type complementarity in the payoff function as theory predicts. In fact, of all inequalities in the summation that are correctly predicted for either $\beta = 1$ or $\beta = -1$, 57% are predicted by $\beta = 1$, a percentage similar to the mean sorting quantile of the univariate tests of the previous section.

Univariate tests. Stepping away from the explicit structure of the theory, we can perform the same univariate analysis for the correlated risk variables. For example, let borrower i be in group M and $P_{occ}(i, M \setminus i)$ be the percent of group M, excluding i, sharing borrower i's occupation. Then, one can estimate whether having the same occupation is a positive or negative contributor to group payoffs:

$$max_{\beta_{occ}\in\{-1,1\}} \sum_{v=1}^{V} \sum_{i\in M_v} \sum_{i'\in N_v} \beta_{occ}\cdot 1\{P_{occ}(i, M_v\setminus i)+P_{occ}(i', N_v\setminus i') > P_{occ}(i, N_v\setminus i')+P_{occ}(i, M_v\setminus i)\}.$$

The sign of the payoff contribution of anti-diversification as measured by coincidence of worst years for income can be estimated exactly analogously to the above equation. $P_{wst}(i, M \setminus i)$ substitutes for $P_{occ}(i, M \setminus i)$ and gives the percent of group M excluding borrower i that agrees with borrower i on which year was worst for income. Both coefficients β_{occ} and β_{wst} are estimated to be positive, indicating a tendency for anti-diversification. For occupation, a full 90% of all inequalities that can be predicted by $\beta = 1$ or $\beta = -1$ are predicted by $\beta = 1$; for worst-year, the percentage is 63%. Again, the numbers are similar to the respective mean sorting quantiles of the previous section, and suggest anti-diversification.

For income shocks (variable described in section 3), we use only the ordinality of the data. All unique values for income shocks are ranked, and each borrower's shock is translated into a rank relative to all borrowers in the dataset. The interactive term analogous to P_{occ} and P_{wst} is then the average absolute rank difference between borrower i and rest of group

[23] The estimator uses a strict inequality though theory requires only a weak one. Given a continuous distribution of match-specific error terms introduced to support the estimator, equalities can be ignored with probability one.

$M \setminus i$, normalized by the range of ranks in the dataset and then subtracted from one.[24] This measure, call it RS_{shk}, thus ranges from zero to one and captures average similarity in income shock ranks between one borrower and the rest of the group. The coefficient β_{shk} is estimated as positive, again supporting anti-diversification, though the percent correct attributable to $\beta = 1$ is just above 51%.

Multivariate reduced-form estimation. The multivariate version of this non-structural approach sets

$$JS_{ij} = \beta_{prob}\, p_i p_{M\setminus i} + \beta_{wst}\, P_{wst}(i, M \setminus i) + \beta_{shk}\, RS_{shk}(i, M \setminus i) \tag{11}$$

and chooses the β's to maximize

$$max_{\beta_{prob}\in\{-1,1\},\beta_{wst},\beta_{shk}} \sum_{v=1}^{V} \sum_{i\in M_v} \sum_{i'\in N_v} 1\{JS_{ij} + JS_{i'j'} > JS_{ij'} + JS_{i'j}\}.$$

Since the maximum score estimator can only estimate the β's up to a scale parameter, β_{prob} is normalized to be in $\{-1, 1\}$. Occupational data are not included here since they come from a different dataset that cannot be merged, as discussed in section 3.

Estimation is carried out using the genetic algorithm routine in Matlab. Results from this estimation give $\hat{\beta}_{prob} = 1$, $\hat{\beta}_{shk} = 0.17$ and $\hat{\beta}_{wst} = 0.38$. That is, all three measures appear as complementary in the joint payoff functions, supporting the theory of homogeneous matching by risk-type and by risk exposure.

As suggested by Fox (2006), hypothesis tests are carried out by subsampling. That is, we create 100 datasets each containing 24 villages, by randomly sampling (without replacement) from the 32 villages in the dataset. The coefficients are estimated for each dataset. The hypothesis that a given coefficient is negative can be rejected at the 5% level if 5 or less estimates of the coefficient are negative, and similarly for a 10% level test. We find that none of the coefficients can be rejected as negative at the 10% level, though each one is not far from significance (all can reject negativity at the 20% level). Thus the support is not overwhelming, but consistent with the idea that both dimensions of heterogeneity create their own incentives for homogeneity.

Multivariate structural estimation. Our final approach with the maximum score estimator uses the full relevant structure of the model. In particular, we estimate the exact joint surplus function 6 arising from the model, reproduced here with slight notational modification:

$$JS_{ij} = 2E - (r + q)(p_i + p_j) + \beta_1\, p_i p_j + \beta_2 \cdot \min\{p_i(1 - p_j), p_j(1 - p_i)\}\, C_{ij}\,, \tag{12}$$

where $\beta_1 = 2q$ and $\beta_2 = 2q\kappa^2$. The first two terms in the surplus involve no interactions, so will drop out of key condition 8. As above, we normalize β_1 to be in $\{-1, 1\}$. Since

[24] That is, if R_i is the income shock rank of borrower i, m the number of members in group M, and $Range$ the difference between the maximum and minimum income shock rank in the dataset, the measure is

$$1 - \sum_{j\in M\setminus i} [|R_i - R_j|/Range] \,/\, (m - 1).$$

$\beta_2/\beta_1 = \kappa^2$, parameter κ is identified. Recall that κ is the probability that the success of a borrower exposed to aggregate shock A or B will be determined by A or B rather than his own idiosyncratic shock. Also, κ^2 is the correlation of returns between two borrowers exposed to the same risk. Thus the intensity of correlated risk can be estimated. Finally, recall that effective joint liability q' for correlated borrowers in homogeneous risk-matching equilibrium equals $q(1-\kappa^2)$; thus κ^2 is the percent reduction in joint liability for homogeneous borrowers (and closely related to it for non-homogeneous borrowers, as we will show).

The extensiveness of correlated risk is captured in C_{ij}, which must be measured. In the theory, C_{ij} is just an indicator variable equaling one if the two borrowers are exposed to the same risk, and zero otherwise. This is proxied here with a convex combination of the worst-year and income-shock measures used above. For borrower i in group M,

$$C_{ij} = \beta_3 \, P_{wst}(i, M \setminus i) + (1 - \beta_3) \, RS_{shk}(i, M \setminus i), \tag{13}$$

with β_3 restricted to lie in $[0, 1]$. Thus the more agreement between borrower i and the rest of group M on which year was worst (higher $P_{wst}(i, M \setminus i)$) and the greater the average similarity in income shock rank between borrower i and the rest of group M (higher $RS_{shk}(i, M \setminus i)$), the more likely we take it be that borrower i is exposed to the same risk factors as the rest of group M (higher C_{ij}). We allow the data to assign relative weights (β_3) to the two measures of correlatedness.

The empirical joint surplus function analog to 12 is then, excluding non-interactive terms,

$$JS_{ij} = \beta_1 \, p_i p_{M \setminus i} + \beta_2 \cdot \min\{p_i(1-p_{M \setminus i}), p_{M \setminus i}(1-p_i)\} \, [\beta_3 \, P_{wst}(i, M \setminus i) + (1-\beta_3) \, RS_{shk}(i, M \setminus i)] \, . \tag{14}$$

A key difference between this joint surplus function and the reduced form one (expression 11) is the somewhat complex interaction between the probabilities of success and the amount of correlation. As discussed in section 2, the combination of success probabilities regulates how much correlation is possible.

The estimator is then

$$max_{\beta_1 \in \{-1,1\}, \beta_2, \beta_3 \in [0,1]} \sum_{v=1}^{V} \sum_{i \in M_v} \sum_{i' \in N_v} 1\{JS_{ij} + JS_{i'j'} > JS_{ij'} + JS_{i'j}\}.$$

One could additionally impose $\beta_2/\beta_1 \in [0, 1]$, since $\beta_2/\beta_1 = \kappa^2$; not doing so provides an extra check against the theory. Again using the genetic algorithm routine, results are that $\hat{\beta}_1 = 1$, $\hat{\beta}_2 = 0.14$ and $\hat{\beta}_3 = 0.89$.

Again risk-type is complementary in the payoff function ($\hat{\beta}_1 = 1$). However, subsampling results, again from 100 datasets of 24 villages, do not allow us to reject the hypothesis that β_1 is negative. Also, the correlation between similarly exposed borrowers κ^2 is estimated at 0.14, an admissible result if seemingly a bit low. The probability of being affected by the aggregate shock is then $\kappa = 0.37$. This positive estimate for $\hat{\beta}_2$ confirms homogeneous sorting by risk exposure; and subsampling results allow us to reject the hypothesis that β_2 is negative, at the 10% confidence level.

The estimates obtained allow for calculation of the average reduction in joint liability due to undiversified risk within groups. Rearranging equation 5 gives the payoff to borrower

i matching with borrower j as $E - p_i r - p_i(1-p_j)q'$, where

$$q' = q\left[1 - \kappa^2 \frac{\min\{p_i(1-p_j), p_j(1-p_i)\}}{p_i(1-p_j)} C_{ij}\right]. \tag{15}$$

Of course, under homogeneous risk-matching and for borrowers with correlated risk ($C_{ij} = 1$), q' simply equals $q(1-\kappa^2)$, as discussed in section 2. The subtracted quantity inside the brackets equals the percent reduction in effective joint liability due to undiversified risk. This quantity can be calculated for each borrower in the sample, as follows. The estimate of κ^2 used is $\hat{\beta}_2 = 0.14$. C_{ij} comes from equation 13 with $\hat{\beta}_3 = 0.89$ substituted in for β_3. Finally, p_i is measured directly and p_j is measured by $p_{M\backslash i}$, as discussed above.

The result is that joint liability is reduced by 5.4% due to undiversified risk. This comes from correlated risk intensity of 14% (κ^2), correlated risk extensiveness of 59% (the average C_{ij}), and reductions in correlation due to departures from homogeneous risk-matching (the term in 15 involving p_i and p_j). Given the wide confidence intervals around both $\hat{\beta}_2$ and $\hat{\beta}_3$, however, this quantity cannot be pinned down with confidence.

A next step may be to increase efficiency by including more inequalities in the maximum score objective function, e.g. k-for-k borrower swaps rather than just one-for-one.

6 Conclusion

[INCOMPLETE]

References

[1] Christian Ahlin. A note on fractionalization decomposition. Working Paper, October 2005.

[2] Christian Ahlin and Robert M. Townsend. Selection into and across contracts: Theory and field research. *Journal of Econometrics*, forthcoming.

[3] Christian Ahlin and Robert M. Townsend. Using repayment data to test across theories of joint liability lending. Vanderbilt University Department of Economics Working Paper 02-W27, December 2002.

[4] Alberto Alesina, Arnaud Devleeschauwer, William Easterly, Sergio Kurlat, and Romain Wacziarg. Fractionalization. *Journal of Economic Growth*, 8(2):155–194, June 2003.

[5] William Easterly and Ross Levine. Africa's growth tragedy: Policies and ethnic divisions. *Quarterly Journal of Economics*, 112(4):1203–1250, November 1997.

[6] James E. Foster and Artyom A. Shneyerov. Path independent inequality measures. *Journal of Economic Theory*, 91(2):199–222, April 2000.

[7] Jeremy T. Fox. Estimating matching games with transfers. Working Paper, January 9, 2006.

[8] Shubhashis Gangopadhyay, Maitreesh Ghatak, and Robert Lensink. On joint liability lending and the peer selection effect. *Economic Journal*, 115(506):1005–1015, October 2005.

[9] Maitreesh Ghatak. Group lending, local information and peer selection. *Journal of Development Economics*, 60(1):27–50, October 1999.

[10] Maitreesh Ghatak. Screening by the company you keep: Joint liability lending and the peer selection effect. *Economic Journal*, 110(465):601–631, July 2000.

[11] Jean Dickinson Gibbons and Subhabrata Chakraborti. *Nonparametric Statistical Inference*. Marcel Dekker, Inc., New York, 4th edition, 2003.

[12] Paolo Mauro. Corruption and growth. *Quarterly Journal of Economics*, 110(3):681–712, August 1995.

[13] Jonathan Morduch. The microfinance promise. *Journal of Economic Literature*, 37(4):1569–1614, December 1999.

[14] Joseph E. Stiglitz and Andrew Weiss. Credit rationing in markets with imperfect information. *American Economic Review*, 71(3):393–410, June 1981.

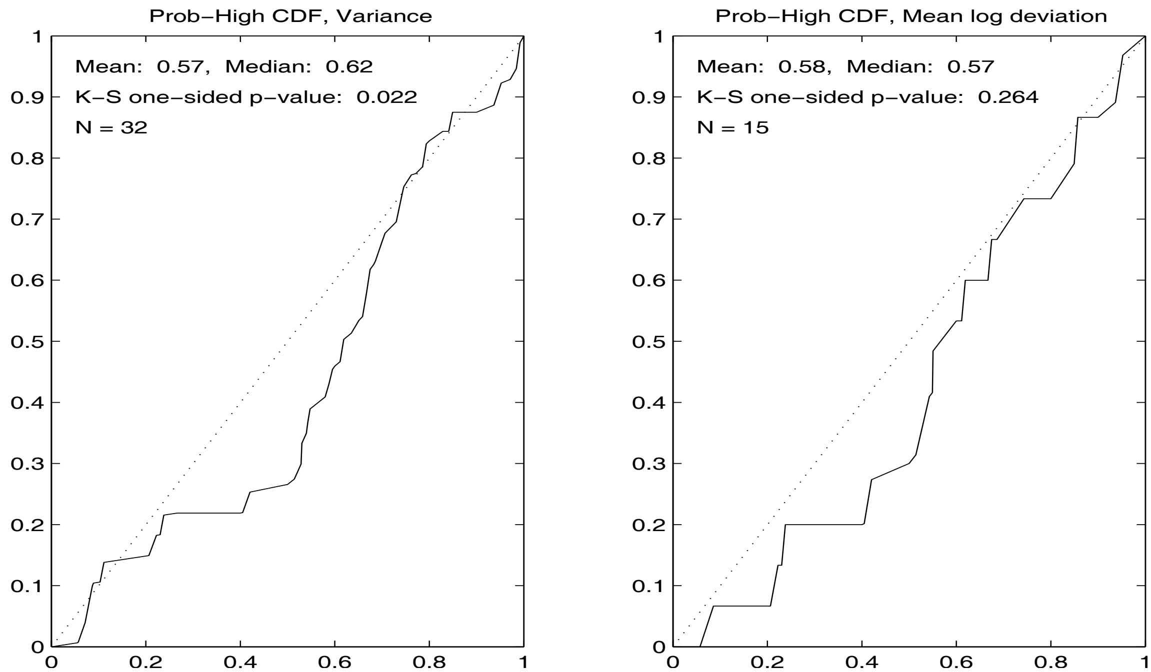

Figure 1: Sample cdf of villages based on decomposition of variance (left panel) and of mean log deviation (right panel) of *prob*, the probability of realizing high income.

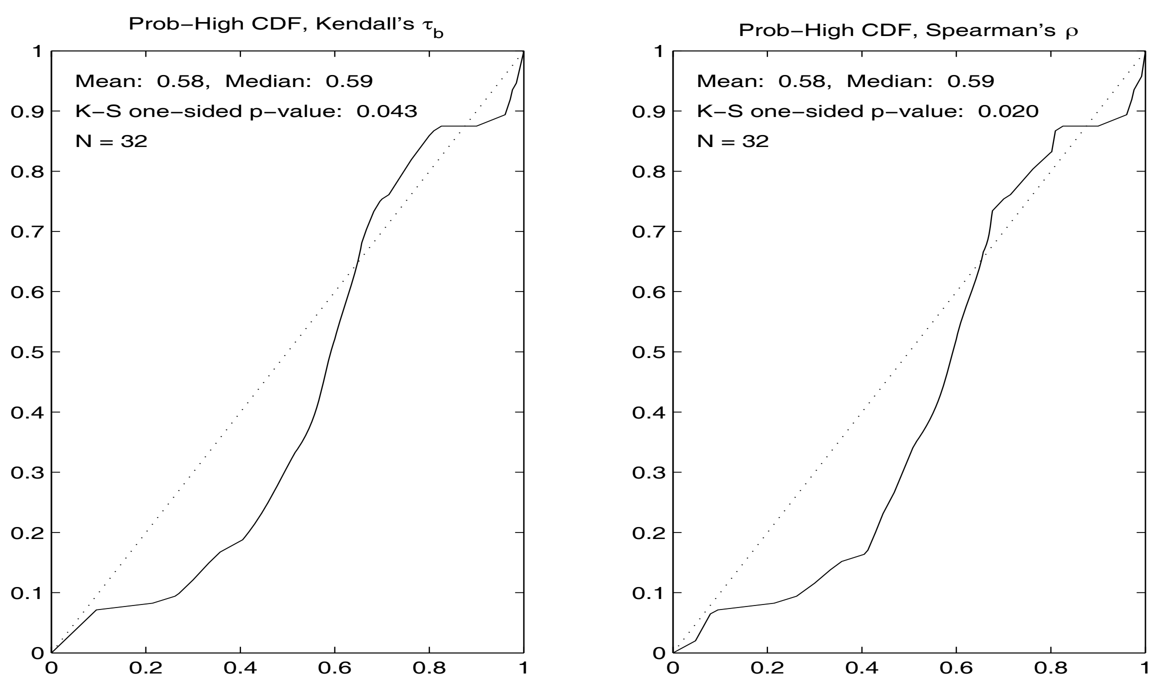

Figure 2: Sample cdf of villages based on rank correlations of *prob* and group identity.

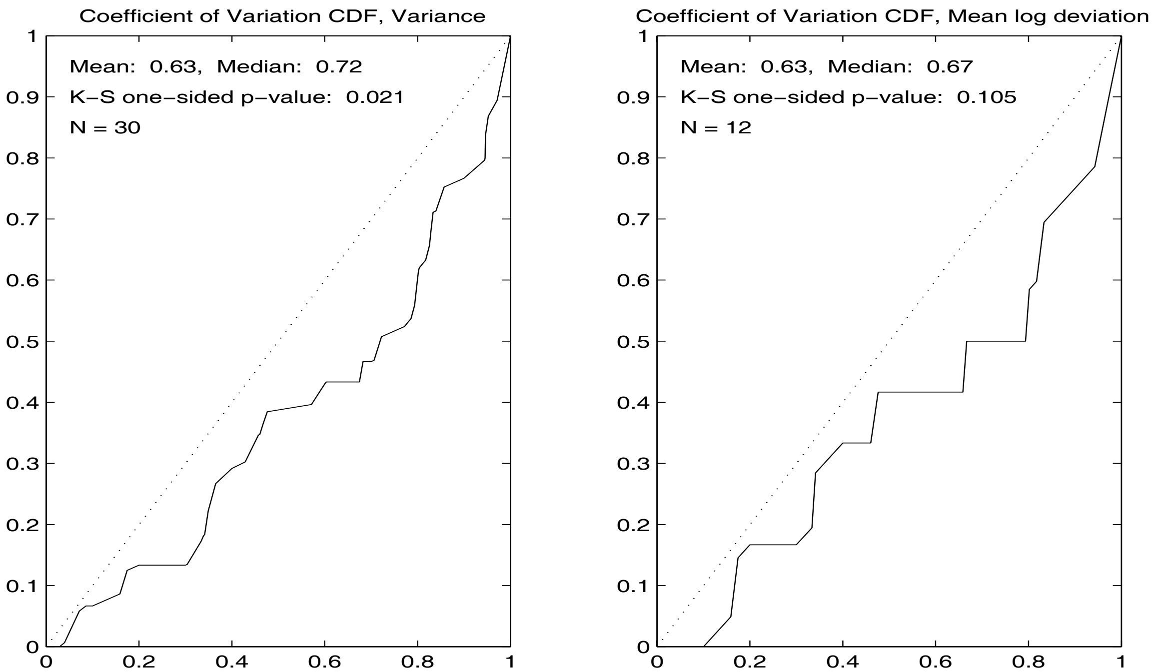

Figure 3: Sample cdf of villages based on decompositions of variance (left panel) and mean log deviation (right panel) of coefficient of variation of future income realizations.

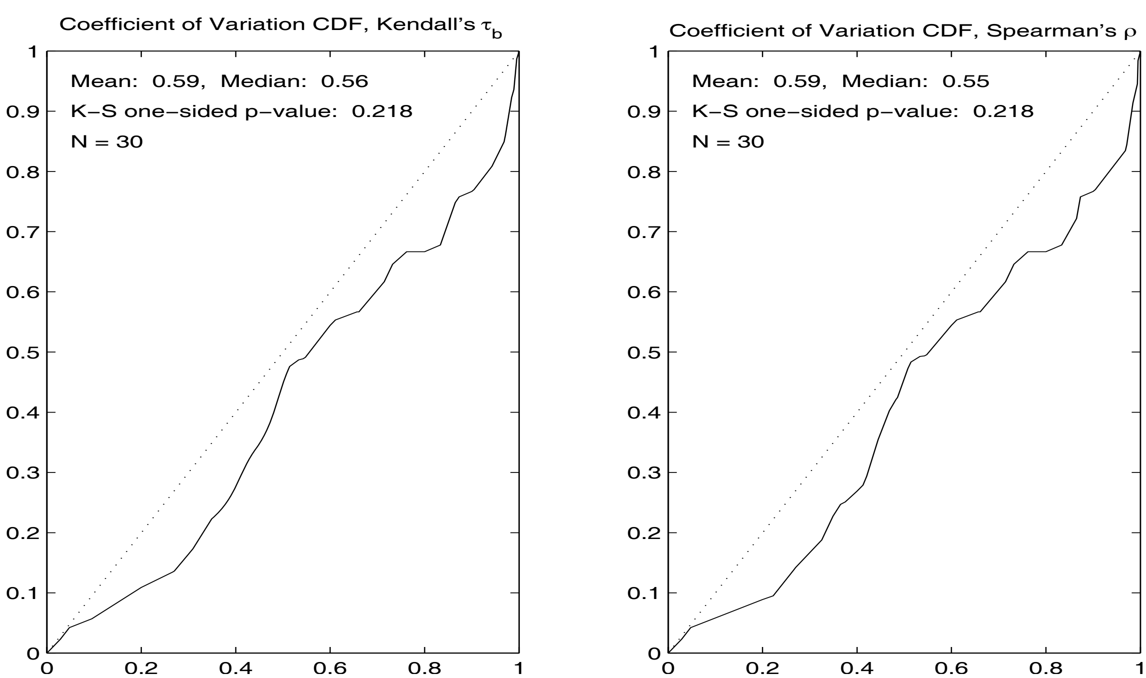

Figure 4: Sample cdf of villages based on rank correlations of future income coefficient of variation and group identity.

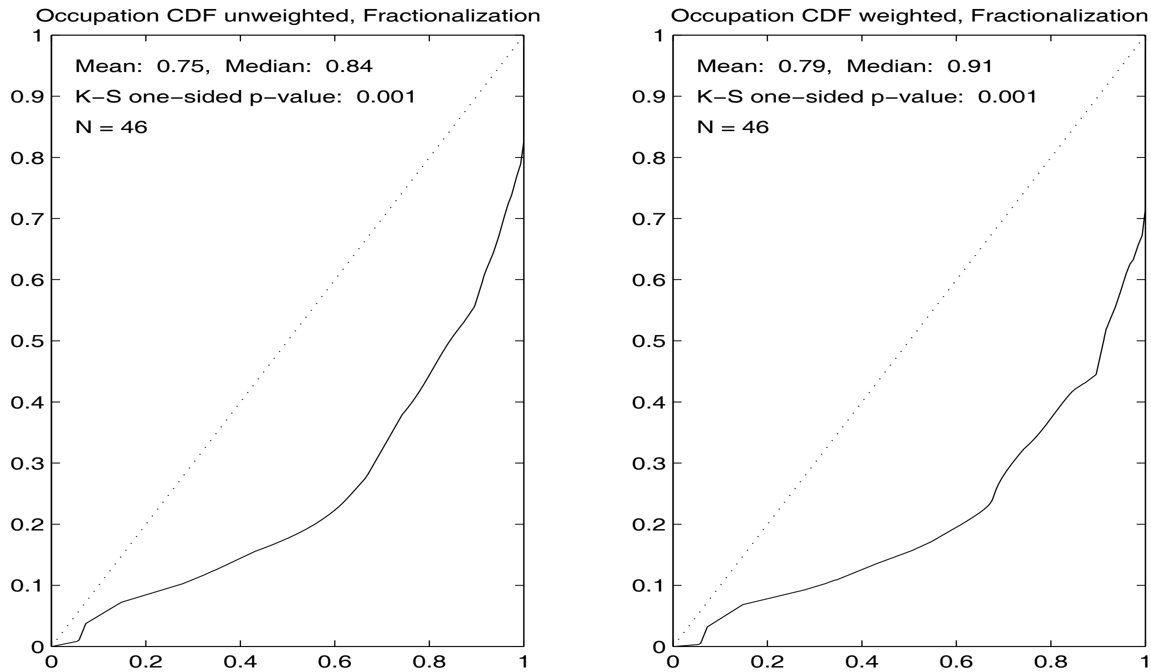

Figure 5: Sample cdf of villages based on decomposition of occupational fractionalization, unweighted (left panel) and weighted by fractionalization (right panel).

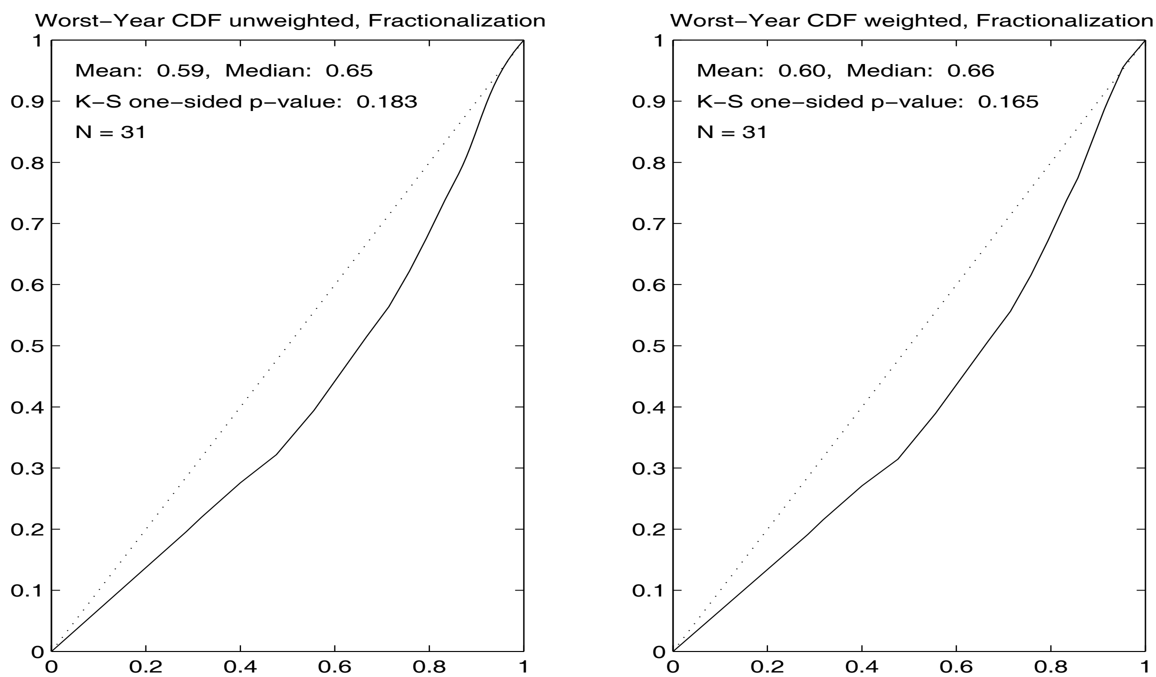

Figure 6: Sample cdf of villages based on decomposition of worst-year fractionalization, unweighted (left panel) and weighted by fractionalization (right panel).

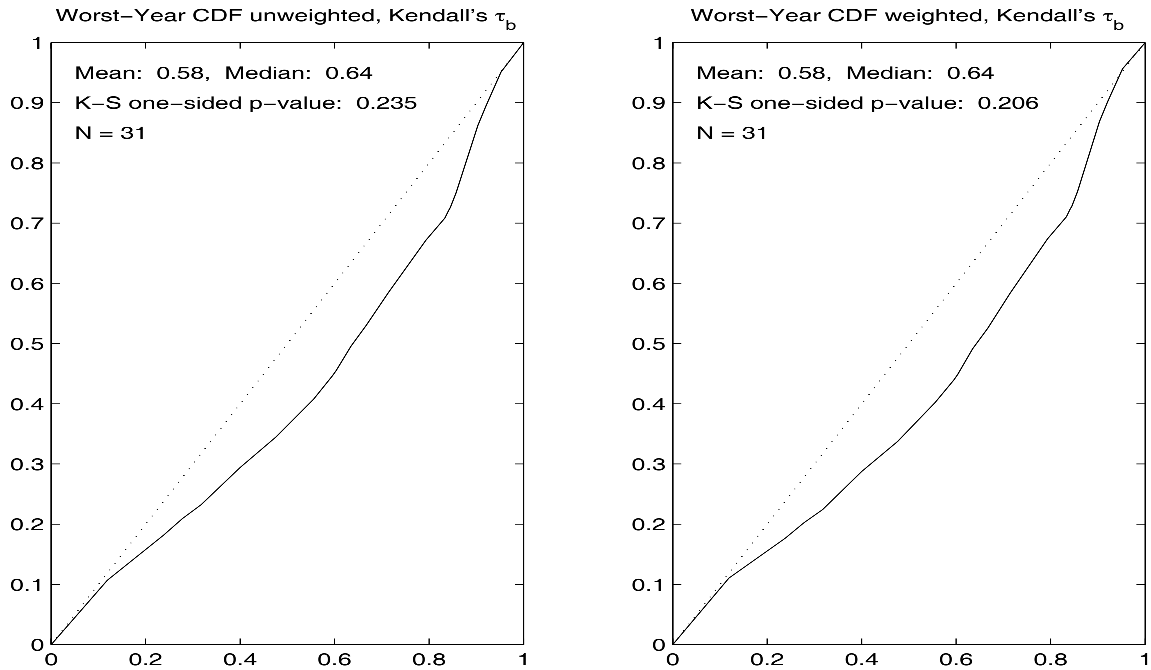

Figure 7: Sample cdf of villages based on Kendall's tau$_b$ rank correlations of worst-year and group identity.

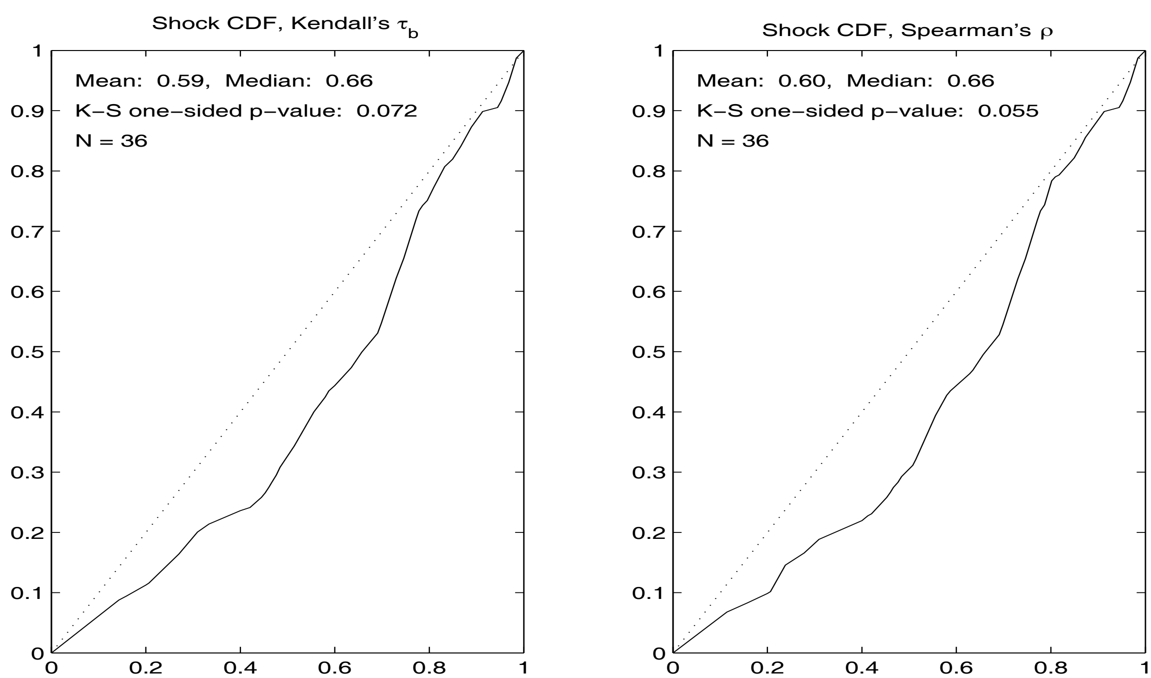

Figure 8: Sample cdf of villages based on rank correlations of income shock and group identity.