ELSEVIER

Journal of Development Economics
Vol. 67 (2002) 245–274

JOURNAL OF Development ECONOMICS

www.elsevier.com/locate/econbase

Localized and incomplete mutual insurance

Rinku Murgai [a], Paul Winters [b], Elisabeth Sadoulet [c,*], Alain de Janvry [c]

[a] The World Bank, 1818 H Street NW, Washington, DC 20433, USA
[b] School of Economic Studies, University of New England, Armidale, NSW 2351, Australia
[c] Department of Agricultural and Resource Economics, 207 Giannini Hall, University of California, Berkeley, CA 94720, USA

Received 1 May 2000; accepted 1 May 2001

Abstract

This paper proposes an explanation of the configuration of mutual insurance groups and of the quality of insurance within each group on the basis of two types of transaction costs: “association” costs in establishing links with insurance partners and “extraction” costs in using these links to implement insurance transfers. We show that optimal insurance arrangements can range from full insurance to autarky within risk pools that span a range of possibilities from cluster formation to community-wide risk sharing. We use a unique data set on canal water trading among households in Pakistan to illustrate the effects of transaction costs in localizing and limiting exchange.

JEL classification: D23; D71; 017
Keywords: Mutual insurance; Transaction costs; Water trading

1. Introduction

People who live in rural areas, particularly those who depend on agricultural incomes, must cope with pervasive uncertainty. There is an extensive literature on a variety of informal transfer mechanisms that help households cope with uncertainty by pooling risk

* Corresponding author.
E-mail addresses: rmurgai@worldbank.org. (R. Murgai), pwinters@metz.une.edu.au (P. Winters), sadoulet@are.berkeley.edu (E. Sadoulet), alain@are.berkeley.edu (A. Janvry).

PII: S0304-3878(01)00189-4

in the community.[1] Yet, empirical studies of risk sharing uniformly reject the hypothesis of Pareto-efficient risk-pooling within rural communities (e.g., Deaton, 1992; Townsend, 1994; Jalan and Ravallion, 1999; Gertler and Gruber, in press).

The failure to detect full insurance at the community level has been explained by the difficulty insurance groups face in monitoring members and in enforcing rules, and therefore in designing incentive compatible and implementable contracts that yield full insurance. These explanations, which emphasize the role of information and social capital, focus on incomplete informal insurance within fixed-size groups such as entire villages. However, relatively little work has focused on localized and incomplete insurance among clusters of individuals within these larger entities, even though the importance of information and enforcement considerations in demarcating the boundaries of exchange networks is widely noted. This paper attempts to explicitly explain the boundaries of mutual insurance groups as well as the amount of insurance within each group on the basis of different types of transaction costs.

When there are transaction costs of monitoring behavior, enforcing participation, and coordinating transfers, households must value the anticipated benefits of risk pooling against these costs when forming connections that underlie exchange. This endogenous process of insurance group formation may result in insurance among only a subset of community members for whom monitoring is easier and enforcement mechanisms (primarily social pressure) are more effective.[2] Indeed, empirical observations suggest that groups are likely to coalesce around individuals who are neighbors, kin, from the same caste, etc. (see, e.g., Scott, 1976; Posner, 1980 and citations therein; Ellsworth, 1989; Fafchamps, 1992; Grimard, 1997). Rejection of Pareto-efficient risk pooling at the community level may thus be due to selecting the village, rather than the appropriate subset of community members, for applying the test of full insurance.

The arguments are organized as follows. In Section 2, we briefly review the literature on mutual insurance and social exchange, with particular attention to different types of transaction costs that affect reciprocal exchange. We draw a distinction between two classes of transaction costs: "association" costs of establishing links with insurance partners and "extraction" costs of using these links to implement insurance transfers.

In Section 3, we lay out a simple model that examines the role of these transaction costs in determining both the *configuration* of insurance-motivated exchanges and the quality of insurance. The configuration (in terms of size) of insurance groups depends upon the value of risk pooling benefits to each household and upon the costs of group maintenance, both of which may be affected by an array of characteristics of the households. We show that

[1] Informal risk pooling arrangements range from gift giving to multipurpose rotating savings and credit associations, mutual aid societies, labor groups, and funeral clubs which sometimes incorporate implicit insurance transactions (Bardhan and Udry, 1999). Reciprocal relationships such as might occur between moneylender and borrower, landlord and sharecropper, patron and clients, or a trader and her suppliers also provide channels for sharing risk.

[2] These considerations are documented by Scott (1976), who notes that in most cases, households cannot count with as much certainty or for as much help from fellow villagers as they can from near relatives and close neighbors. The reciprocities of kinship, particularly among bilateral kindred, diminish perceptibly the more distant the bond.

the optimal insurance arrangement can range from partial insurance at the community level, to full insurance in clusters, partial insurance in clusters, and situations in which insurance does not arise at all.

In Section 4, we turn to a unique data set on canal water exchange among households along two tertiary irrigation canals in Pakistan to examine the role of transaction costs in localizing and limiting exchange. Two key advantages of this data set are that the universe of potential exchange partners is clearly defined and transaction costs between every pair of households within this universe are explicitly measured. Our empirical analysis has two parts. We first examine the evidence that transaction costs localize exchange by shaping the choice of trading partners and then assess how the size of localized clusters and intensity of exchange within clusters are limited by these costs. Conclusions are drawn in Section 5.

2. Conceptualizing mutual insurance

There are a number of studies that have tested for efficient risk pooling, or full insurance, in a variety of rural contexts (Deaton, 1992; Townsend, 1994; Grimard, 1997; Ligon, 1998; Ligon et al. 1999; Gertler and Gruber, in press). The test is based on the proposition that in a Pareto-efficient allocation of risk, household consumption comoves with average consumption in the risk pool and is unaffected by its own income. The hypothesis of full insurance is rejected in all cases. This widely observed failure of full insurance has been attributed to the lack of good mutual information and effective enforcement mechanisms that are central to supporting insurance arrangements (Kimball, 1988; Coate and Ravallion, 1993; Ligon, 1998). Information and enforcement problems within the risk pool may result in only partial insurance in some states of nature when transfers are restricted by incentive compatibility or participation constraints.

Another more recent approach to understanding incomplete insurance has been to ask whether the village is the appropriate risk pool for sustaining informal insurance. For example, Grimard (1997) who uses the same data as Deaton (1992), hypothesizes, based on anthropological evidence from Côte d'Ivoire, that the correct risk pool is not the community but rather the ethnic group. He finds that there is greater risk pooling within ethnic groups than communities. Fafchamps and Lund (2000) provide evidence that risk-sharing takes place primarily within small groups of family and friends, and that even within these groups not all shocks are perfectly insured. Goldstein (2000) finds that insurance networks in Ghana are organized along gender lines, with women insuring primarily with other women rather than the entire community.

There is considerable additional evidence that lends support to the hypothesis that insurance may be localized to small social spaces within villages, or to groups that may well transgress village boundaries. Empirical observations of rural communities suggest that transactions tend to occur between partners or groups within which there is good mutual information and access to enforcement mechanisms, primarily through social pressure (e.g., Posner, 1980; Ellsworth, 1989; Fafchamps, 1992). Households may also sort themselves into participants and nonparticipants in risk-sharing arrangements along wealth lines because of differences in the benefits obtained from participation (see Hoff, 1997, and various citations therein). Patterns of reciprocal exchange also tend to display a certain degree of persistence over time. As Kranton (1996) shows in the context of market exchange, this persistence in exchange-flows can become self-sustaining because the mode of exchange itself lowers costs of reciprocal exchange. It, thereby, progressively isolates long-term economic partners from competing sources of demand and supply, and results in localized patterns of exchange.

However, relatively little work to date has focused on the endogenous processes which determine the choice of insurance partners (and thus, how a potential risk pool such as a village may partition into insurance clusters), and the quality of insurance within the risk pool.[3] In this paper, we develop a simple model that considers both decisions jointly, as a function of different types of transaction costs. We classify the costs of reciprocal exchange into two types. The first are fixed or sunk costs of co-ordination and information processing that must be incurred for reciprocal exchange between a group of individuals to function. These *association costs* include costs such as those of searching for potential partners, establishing relationships, and coordinating activities. In the context of mutual insurance, association costs therefore depend on the number of members in an insurance cluster but not on the degree of risk pooling within the group.

Transaction costs of association have important implications for the patterns of exchange within communities. From a set of potential partners for exchange, households will prefer partners with whom they have the lowest costs of association and, as a result, insurance arrangements may be localized to entities smaller than the community.[4] Previous studies that have focused on the village as the appropriate insurance group, within which one would expect full insurance implicitly, assume zero or sufficiently low association costs among village inhabitants and infinite or extremely high association costs with households outside the village.

The second class of transaction costs, which we shall refer to as *extraction costs*, arise when exchange is plagued by moral hazard due to imperfect monitoring and enforcement. These are costs that vary with the level of transfer requested from the partner. Individuals that need to transfer resources to partners are likely to be more reluctant to comply for larger amounts or for actions which have a higher cost to them. They will expend more effort trying to avoid having to share, hiding their luck, being absent from home when expecting the request, etc. In such cases, the community has to expend more effort to induce these individuals to comply. Extraction costs can thus be expected to increase with the quality of insurance.

[3] For an exception, see Hoff (1997) who develops a model in which the composition of the insurance pool is itself determined by the degree of risk sharing. Her paper shows that in the context of imperfect information, in equilibrium, insurance arrangements may organize along wealth lines, with poor households participating and richer households opting out of the scheme.

[4] Tests of whether insurance deteriorates with 'distance' by Rashid (1990) (cited in Townsend, 1995) in Pakistan, and Townsend (1995) in Thailand, provide suggestive evidence in support. Both studies, however, test for the village (or region) as the smallest potential insurance group and aggregate up to the national level.

The particular structure of association and extraction costs that prevails among members of a community will together determine if insurance is localized to clusters or inclusive of all households, and if it is imperfect or efficient within the community. To establish this relation, we develop in the next section a model of mutual insurance where these costs endogenously determine both the configuration and performance of informal insurance arrangements.

3. Theory: cluster size and optimal risk-pooling

Suppose there are multiple risk-averse households who face intertemporally variable and independent income streams. All households are assumed to be similar ex-ante, with identical preferences defined over own income. Each household receives income y with probability π and suffers a loss L (i.e., receives income $y-L$) with probability $1-\pi$. Each household's expected utility in the absence of any kind of informal insurance is:

$$Eu = \pi u(y) + (1-\pi)u(y-L). \tag{1}$$

Since all households are risk-averse ($u' > 0$ and $u'' < 0$) and face uncertain income streams, there are potential gains from state-contingent transfers between them. Consider the informal insurance arrangement in which each household that does not suffer a loss agrees to transfer a share α, $0 \leq \alpha \leq 1$, of the loss L to every household which receives income $(y-L)$. Since all households are assumed to be identical, transfer arrangements are symmetric. No transfers are made if all households receive the same income. If there are n people in an insurance cluster, expected utility over income for a household is given by:

$$\begin{aligned} Eu(\alpha,n) = {} & \pi\left[\sum_{x=0}^{n-1} p(x)u(y-(n-1-x)\alpha L)\right] \\ & + (1-\pi)\left[\sum_{x=0}^{n-1} p(x)u(y-L+x\alpha L)\right]. \end{aligned} \tag{2}$$

x is defined as the number of *partners* in the risk pool that do not suffer a loss, and correspondingly, $p(x)$ is the probability that x households of the $(n-1)$ partners do not have a loss. Since income realizations within the risk pool are assumed to be independent of each other, $p(x)$ is distributed as a binomial with $p(x) = \binom{n-1}{x}\pi^x(1-\pi)^{n-1-x}$.

The first term in square brackets in Eq. (2) is expected utility when the household suffers no loss. According to the risk sharing arrangement, the household must transfer αL to each of its $(n-1-x)$ partners who received a low income, with a total transfer of $(\alpha L)(n-1-x)$. Similarly, the second term in square brackets is the expected utility for the household when it suffers a loss. In this case, the household receives a total transfer of

$(\alpha L)(x)$ from the x households who did not incur a loss.[5] While α does not vary across states of nature, the total transfer made (received) by a household in any state depends on the total number of households that incurred a loss (gain), and therefore on aggregate cluster income in that state.[6]

Intuitively, the household's motivation in entering a mutual insurance arrangement is to transfer income from good to bad states in order to receive steady income across states of nature, lowering the variance of income without affecting the expected value. Mean and variance of expected income under the insurance arrangement are given by:

$$\mu = y - (1-\pi)L,$$

and

$$\sigma^2 = \pi(1-\pi)L^2[1 + n(n-1)\alpha^2 - 2(n-1)\alpha] = \sigma_0^2[1 + n(n-1)\alpha^2 - 2(n-1)\alpha],$$

where ${\sigma_0}^2$ is the variance of income under autarky (no insurance). Since expected income does not depend on cluster size (n) or the degree of risk pooling (α), optimal cluster size and risk-pooling are chosen to reduce the variance of income below autarky levels while leaving mean income unchanged.

Mathematically, the optimal insurance arrangement is derived by maximizing $Eu(\alpha,n)$ with respect to α and n. Assuming that the cost of increasing cluster size beyond the community is prohibitively large, we include the constraint that cluster size cannot exceed N, the size of the community.[7] Further, we impose the implementability constraint that utility from insurance has to be at least as great as the autarky utility level:

$$\begin{aligned} &\max\nolimits_{\alpha,n}\ Eu(\alpha,n) \\ s.t.\quad &Eu(\alpha,n) \geq Eu(0,1) \\ &0 \leq \alpha \leq 1 \\ &n \leq N. \end{aligned} \tag{3}$$

[5] Note that the value of x—the number of *partners* in the risk pool who do not suffer a loss—is different from the *total* number of households that have successful outcomes. Consider a case in which 3 members of a 5-person insurance group do not have a loss. A household that did not suffer a loss must transfer $2\alpha L$ to the losers. In this case, x, which is the number of partners who are not losers, is 2. On the other hand, if the household suffers a loss, it receives $3\alpha L$ from its partners who are not losers, x in this case is 3.

[6] Transfers can also be specified as an absolute amount paid to each household that suffers a loss. What matters is that the total amount transferred or received by a household should be allowed to vary depending on the state of nature.

[7] More generally, N is the largest cluster (e.g., village, ethnic/kinship group, wealth class, region) in which insurance partnerships can be developed, with association costs prohibitively large beyond N.

3.1. Cluster size and optimal risk-pooling in the absence of transaction costs

In the absence of transaction costs, the optimal insurance arrangement would smooth consumption in such a way that each (ex-ante identical) household consumes the average level of cluster income. If z of the n households had successful outcomes, then average income in the cluster would be $y-[(n-z)(L/n)]$. Each of the z successful households would in that case be called on to transfer L/n (i.e., $\alpha=1/n$) to the $(n-z)$ households that suffered a loss. There is full insurance as total cluster losses $[(n-z)L]$ are ex-post divided equally among the households.

From the characterization of mean and variance above, we know that expected income is always $[y-(1-\pi)L]$ but the variance of income faced by each household when there is full insurance ($\alpha=1/n$) within the cluster is σ_0^2/n; this residual risk mirrors the aggregate risk faced by the cluster. Therefore, it is in the interest of each household to belong to as large an insurance cluster as possible so as to reduce the variance of average cluster income.[8]

For an infinitely large cluster, the household receives perfect insurance; i.e., no uncertainty remains and the household receives the mean income in all states. If, however, n is restricted to be no greater than a particular value N—for example, the size of the village—then the household chooses an insurance cluster of that size and the residual uncertainty is σ_0^2/N.[9] The literature on mutual insurance generally assumes, implicitly or explicitly, that transaction costs which might prevent the formation of a partnership between all members of a community are zero. This being the case, after the realization of the state of nature, transfers may flow between any two households to bring about risk pooling. If, in an absurd extreme, establishing and maintaining partnerships is indeed costless, there is no reason for a mutual insurance cluster not to be community-wide or worldwide. Real-world limits to cluster size must therefore be the result of costs relating to the formation and maintenance of partnerships.

3.2. Cluster size and optimal risk-pooling with transaction costs

If there are both association and extraction costs to insurance, whether or not there is full or partial insurance and community-level membership or cluster formation, depends on the combined costs of increasing cluster size and improving the degree of risk sharing.

[8] This intuitive result can be mathematically derived by maximizing expected utility in Eq. (2) with respect to α, for a given cluster size n. Since there are no costs to insurance, the autarky constraint will not bind and an interior solution is guaranteed for α. Solving for the first-order condition provides α equal to $1/n$. To determine the optimal cluster size, it is possible to show that the distribution of payoffs faced by a household in an insurance cluster of size n second-order stochastically dominates the distribution with a cluster of less than n households. It follows that expected utility increases as cluster size increases (Hadar and Russell, 1969).

[9] Note that we have assumed that there is no covariance between income realizations. A positive covariance of income among members of a finite sized group would diminish the effectiveness of informal insurance in reducing individual risk.

Let $c(\alpha,n)$ be the sum of the association and extraction costs paid by each group member. We assume that the overall utility obtained from the insurance scheme is separable between income and the cost incurred to enforce the contract.[10] The household expected utility function is thus written as:

$$Eu(\alpha,n) = \pi\left[\sum_{x=0}^{n-1} p(x)u(y-(n-1-x)\alpha L)\right] + (1-\pi)\left[\sum_{x=0}^{n-1} p(x)u(y-L+x\alpha L)\right] - c(\alpha,n). \tag{4}$$

The optimal transfer arrangement and cluster size are such that the marginal benefits of variance reduction are equal to the marginal cost of greater insurance or additional partners. This can best be seen by taking a second-order Taylor series approximation of utility u in each state about expected income. This yields:

$$Eu(\alpha,n) = u(\mu) + \frac{1}{2}u''(\mu)\sigma^2 - c(\alpha,n). \tag{5}$$

The first-order conditions for an interior solution are:

$$\begin{aligned} -u'(\mu)\phi(\mu)\sigma_0^2(n-1)(n\alpha-1) &= c_\alpha(\alpha,n), \\ -\frac{1}{2}u'(\mu)\phi(\mu)\sigma_0^2\alpha[\alpha(2n-1)-2] &= c_n(\alpha,n). \end{aligned} \tag{6}$$

The marginal benefit of a larger cluster or greater insurance increases with the coefficient of absolute risk aversion ϕ, the level of variance faced by the household in autarky σ_0^2, and the degree of responsiveness to variance of cluster size. The marginal utility of income, $u'(\mu)$, is a normalization factor.

To focus on the effects of association costs, consider a simplification of Eq. (5) that imposes a fixed cost A to establish a partnership with each household, and restricts extraction costs to zero. That is, $c(\alpha,n)=(n-1)A$. Since there are no costs to α, there is full income pooling within the cluster. Knowing this, the household chooses its optimal number of partners to balance the benefits of reduced income variance from increasing n with the costs of increasing cluster size. Replacing the condition that $\alpha=1/n$ in Eq. (6), optimal cluster size is given by:

$$n^2 = \frac{1}{2A}u'(\mu)\phi(\mu)\sigma_0^2. \tag{7}$$

[10] Monitoring costs can be thought of as the disutility of lost leisure time. The precise nature of these costs will vary with monitoring technology. For a similar specification, see Newbery and Stiglitz (1981).

When costs of forming partnerships are small, there is full insurance at the community level. Alternatively, if costs of establishing partnerships are high, insurance is too expensive to sustain and households are better off in autarky. In between these two extremes, for the interior solution, insurance within the community is characterized by clusters with full insurance within each cluster.

To focus on the case of extraction costs, consider the household's optimization problem when association costs are zero but there are nonzero extraction costs. Since there are no association costs, all N community members belong to the insurance cluster and only α needs to be chosen. Although the specifics of the insurance contract and implications for cluster size will vary with monitoring (or enforcement) technology, we can represent this idea with a simple linear technology where the marginal extraction cost is constant and does not depend on the state of nature or on cluster size, $c(\alpha)=E\alpha$, $E>0$. We can rewrite the first-order condition for the choice of α in Eq. (6) as follows:

$$\alpha = \frac{1}{N} - \frac{E}{(N-1)Nu'(\mu)\phi(\mu)\sigma_0^2}. \tag{8}$$

In the absence of extraction costs ($E=0$), there is full community-level insurance. As soon as cost becomes positive, α decreases below $1/N$, leading to partial insurance within the community. Utility falls from its maximum as extraction costs increase. Beyond a certain level of extraction costs, participating in the insurance arrangement leads to utility levels that are lower than those with no insurance at all and the household reverts to autarky.

Comparing Figs. 1 and 2 illustrates why both association and extraction costs may yield similar interpretations to traditional empirical tests of full insurance. In Fig. 1, as

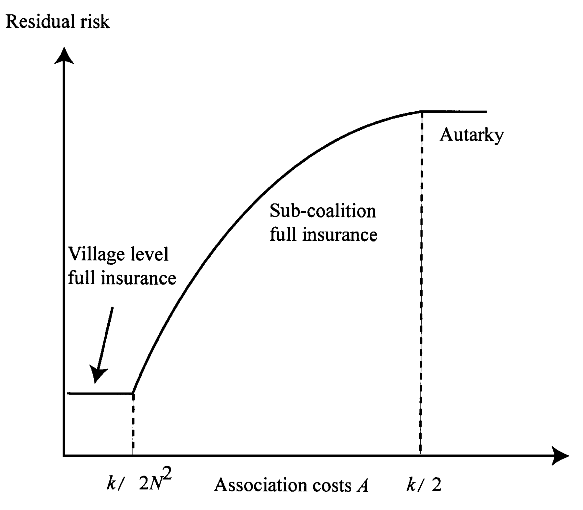

Fig. 1. Residual risk faced by a household as a function of association costs.

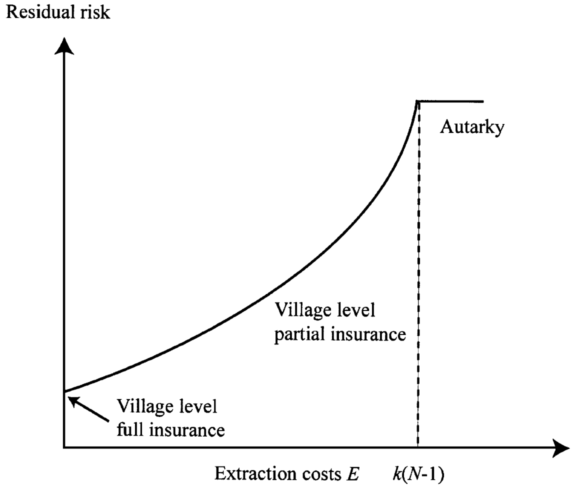

Fig. 2. Residual risk faced by a household as a function of extraction costs.

association costs increase, clusters become smaller and the residual risk faced by the household increases. Likewise, in Fig. 2, as extraction costs increase, households settle for partial insurance, and therefore face increasing degrees of residual risk. Rejection of full insurance could then be due to association costs, extraction costs, or more likely a combination of the two. The existence of clusters implies that tests of insurance at the village level might also be refuted even if there is full insurance within clusters because the entire village is presumed to be the appropriate group on which to test for efficient risk pooling.

When costs of association and extraction are positive, different insurance regimes can be summarized graphically in a two-dimensional cost space, with extraction costs on the vertical axis and association costs on the horizontal axis as shown in Fig. 3. The five regions—full community-level insurance, partial community-level insurance, full insurance within clusters, partial insurance within clusters, and autarky—are optimal responses to the combination of association and extraction costs, which satisfy the first-order conditions for an interior solution subject to the constraints that the insurance cluster cannot be larger than the community and that utility from the insurance arrangement must be at least as high as that obtained in autarky.[11]

Full insurance at the community level obtains only when the cost of extracting insurance is zero and costs of establishing partnerships are sufficiently low (region 1 on the horizontal axis such that $A < k/2N^2$ where $k = u'\ \phi\sigma_0^2$). As association costs increase ($k/2N^2 < A < k/2$), it is too costly to insure at the community level and clusters form (region 2). However, with zero extraction costs, there is full insurance within each cluster.

[11] These regimes are derived analytically in Appendix A.

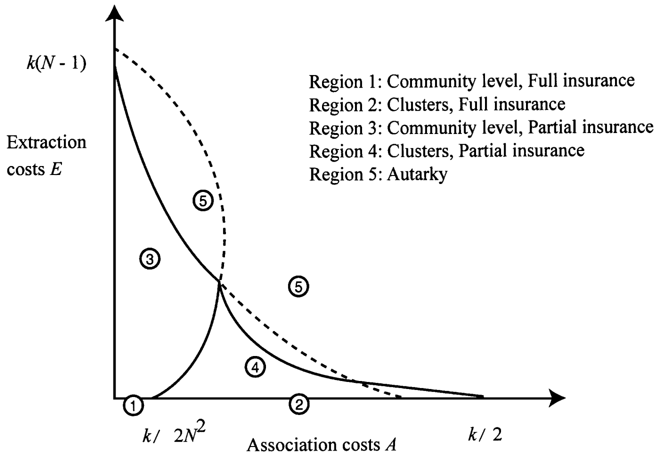

Fig. 3. Alternative insurance regimes as a function of association and extraction costs.

Notwithstanding, in this region, there is less than full insurance at the community level as households aggregate into smaller clusters.

With low costs of organizing partnerships, but positive extraction costs, the optimal insurance arrangement involves a community-wide cluster, but only partial insurance within the community (region 3). If association costs are high and extraction costs sufficiently low, the optimal response is to divide the community into smaller clusters and have partial insurance within that each of these (region 4). Finally, if either association costs are very high ($A < k/2$), or extraction costs are substantial ($E > k(N-1)$) or a combination of the two is high, insurance is too expensive and households opt for autarky (region 5).

A discussion of the slopes of the boundaries dividing insurance regimes is instructive since it illuminates the relationship between cluster size and the transfer arrangement at the optimum. The boundary between regions 3 and 4 represents the combination of costs, along which the household is indifferent between having partial insurance at the community level and partial insurance in a cluster marginally smaller than the community size ($n < N$). Along this boundary, an increase in association costs creates an incentive to reduce cluster size. A household would be willing to maintain cluster size at N only if extraction costs were higher (and, therefore, lower insurance provided within the cluster). Notice that utility decreases as costs of association and extraction increase.

The boundary between regions 3 and 5 represents the combination of costs at which the household is indifferent between partial insurance at the community level and autarky. In contrast to the previous case, along this boundary, utility is maintained at autarky levels. An increase in association costs induces a smaller cluster size. Utility is restored to autarky levels only when insurance provision increases within the cluster. This occurs as extraction costs decrease. Likewise, the downward slope of the outer curve between regions 4 and 5 arises from the households indifference between partial insurance in clusters and autarky.

3.3. Comparative statics

The benefits of insurance vary with changes in expected income μ, the exposure to risk under autarky σ_0^2, and the degree of risk aversion $\phi(\mu)$ (see Eq. (6)). As expected income increases (either as L decreases or y increases), the marginal utility of income decreases. Assuming constant or decreasing absolute risk aversion, k decreases, shifting all the boundaries and intersections closer to the origin in Fig. 3. This result is in agreement with the frequent observation that informal insurance arrangements are less likely to occur as incomes increase.[12]

Risk exposure σ_0^2 faced by the household under autarky changes with both the size of the loss L and the probability of a loss $1-\pi$. Any factor that increases the variance of income makes the household more tolerant towards costs, shifting all the boundaries and intersections away from the origin. For small losses, then, the household becomes increasingly intolerant to costs, and may prefer autarky to insurance which protects it from this small loss. The effect of reducing the probability of 'success' can be thought of as increasing the likelihood that the household will suffer a loss as well as a reduction in expected income. The optimal response, naturally, is to be willing to tolerate greater costs to establish and maintain exchange agreements. By the same mechanism, households with lower aversion to risk are less willing to tolerate the transaction costs of insurance.[13]

Finally, with the linear specification of transaction costs used above, the size of the cluster n is a decreasing function of association costs A, and the quality insurance α a decreasing function of the extraction costs E. The variation of n with respect to E is of same sign as the variation of α with A, but both can be either positive or negative.

To summarize the main results of the model, we have identified conditions under which transfers of some form take place (i.e., a risk-sharing institution exists) and under which full income pooling is achieved (i.e., the institution achieves first-best risk sharing). By making the choice of partners endogenous, conditions under which risk-sharing institutions exist in clusters of the community population are identified. Much of the previous work on limitations to mutual insurance has considered an exogenously determined risk-pool, equivalent to a village or an ethnicity. This assumption restricts insurance regimes only to changes in the degree of insurance at the community level (i.e., to changes along the vertical axis in Fig. 3), even though nothing precludes smaller insurance clusters. Comparative statics on optimal informal insurance regimes suggests that the divergence between the optimal insurance arrangement and first-best community

[12] For example, Evans-Pritchard (1940) (cited in Fafchamps, 1992) suggests that "it is scarcity not self-sufficiency that makes people generous".

[13] This implication of the model is related to the findings of Gaynor and Gertler (1995) who focus specifically on the impact of risk aversion on the degree of risk sharing and group size in medical partnerships. They show that at higher levels of risk aversion, compensation schemes are designed such that there is greater risk pooling but group size is correspondingly reduced to allow greater monitoring of physician behavior.

level risk sharing is smaller when insurance is badly needed: in situations where incomes are lower, losses larger, and the probability of bad realizations higher.

4. Empirical example: canal water trading in Pakistan

Econometric evidence on the factors that play a role in determining the configuration of exchange links between households and the consequent partitioning of communities into smaller insurance clusters is hard to come by. Previous empirical work on insurance can be broadly classified in two categories. First, there is a rich empirical literature that examines the specific mechanisms—such as gift exchange, reciprocal credit—that help pool risk in a community (e.g., Udry, 1994; Platteau, 1991). Second, a number of studies assess the overall extent of risk pooling achieved within a village/kinship/ethnicity or some other group (e.g., Deaton, 1992; Townsend, 1994).

This latter set of studies is agnostic about both, the specific mechanisms by which risk sharing takes place, as well the configuration of exchange links within a potential risk pool. The few empirical studies that reflect on the configuration of insurance arrangements use a partial approach as the focus is on insurance mechanisms from the point of view of an individual with one or more insurance partners, with no relationship drawn to the overall structure of exchanges (e.g., Rosenzweig and Stark, 1989; Fafchamps and Lund, 2000). One reason for this dearth of evidence is that few, if any, data sets provide detailed information relating every pair of households within a well-defined universe of potential insurance partners.

This paper largely avoids this problem by using data on canal water exchange from two watercourses (tertiary irrigation canals), Azim 43-L and Fordwah 14-R, in the Punjab Province of Pakistan.[14] The advantage of focusing on canal water exchange at the tertiary canal level is that the universe of potential trading partners is clearly defined since, aside from its connection to the secondary canal, the watercourse is closed to import or export of canal water. Moreover, since the data set includes every farmer that is entitled to use water from the canal (i.e., it is a census rather than a sample), both sides of every exchange can be identified, and measures of transaction costs between every pair of farmers can be constructed. Finally, as discussed below, another advantage of focusing on this system of exchange is that transaction costs of exchange along at least two dimensions can be clearly identified and measured.

The two watercourses are part of the Fordwah–Eastern Sadiqia irrigation system of southern Punjab, which lies in the cotton–wheat agroclimatic zone. Cotton and fodder are the main *kharif* (summer) season crops, with cotton by far the more important in terms of cultivated area. Sugarcane is cultivated year-round but most of its irrigation requirements are also concentrated during the summer months. Both watercourses are in the tail-end of the system, and hence have particularly unreliable canal supplies. The data used for the

[14] This example comes from field observations during a stay at the International Water Management Institute (IWMI) in Pakistan.

analysis cover every canal water exchange between every farmer in the watercourse during the *kharif* (summer) 1994 season.

In these watercourses, as in much of Pakistan and Northern India, canal water is distributed to each farmer according to a rotational water-delivery system, or *warabandi*, which provides turns to use the entire water flow in the canal to farmers at a pre-specified time each week. Each farmer who cultivates land within the watercourse command area is assigned a water turn (e.g., 20 min per acre) that is proportional to his landholding. The sequencing of water turns during the week is fixed by the location of the plot in the command area; a farmer with a plot closest to the head receives his water turn first, followed by adjacent farmers, until the tail farmer has received his turn at the end of the week, after whom the water turn sequence reverts to the head farmer for the following week.

Fig. 4 shows a diagram of the Azim 43-L watercourse, indicating the location (and hence, warabandi turn) of each of the 35 farmers who cultivate land in this command area. The Fordwah 14-R watercourse (map not included) covers a much larger area, with 94 households. The Azim 43-L watercourse is almost linear, and its command area stretches over a distance of approximately 2600 yards (38 distance units[15]). Households in the middle of the watercourse are at a distance of around 5 units to all other households, while the household located at the head is at an average distance of 24 units to the others. Lengths of water turns, corresponding to areas commanded by farmers, vary from 30 min to 21 h, averaging 4.8 h. By contrast, the Fordwah 14-R watercourse is more compact, and distance between the 94 households never exceed 43 units, with an average of 10 units. The distribution of farm size is also less unequal and water turns vary from 10 min to 7.5 h, averaging 1.8 h. In both watercourses, in response to the rigid water allocation scheme and the unreliability of actual water deliveries,[16] farmers have developed an informal system of canal water trading.[17] Data on water exchanges between farmers for the two watercourses during the kharif 1994 season show that there are significant transfers of water time within each watercourse and that no payment is provided for the transfers.[18]

[15] Unit of physical distance are the side of an acre of land, i.e., approximately 70 yards.

[16] The amount of water entering tertiary canals is highly unpredictable due to problems with operation and maintenance, overuse, and illegal diversions of water in upstream irrigation channels (Bandaragoda and ur Rehman, 1995).

[17] The theoretical model developed in Section 3 is constructed around the usual motive for insurance: smoothing consumption. As pointed out by a referee, in the context of this example, the model could just as well have been recast in terms of a cluster of risk neutral farmers arranging water transfers to maximize joint production income. The results on the role of transactions costs and on patterns of exchange would look like mutual insurance because of the curvature of the production functions (farmers within a cluster would try to equalize the marginal value product of water inputs).

[18] On the whole, trading is less frequent in Fordwah 14-R, with only 4% of the 4371 households pair exchanging full or partial turns at least once during the season, although all households except one had at least one exchange with one partner along the course. Several factors are at play. First, due to administrative preference for the Fordwah distributary, there is less variation in water availability than in Azim, thus providing lesser incentives to trade and insure water realizations. Second, participatory management of water turns is a common practice in Fordwah but not in Azim. Fordwah has a much less inequitable land distribution and a less fragmented social structure than Azim (Strosser and Kuper, 1994), which may explain the differential abilities to engage in joint management. With participatory management in Fordwah, private exchange is a less important mechanism to spread risk.

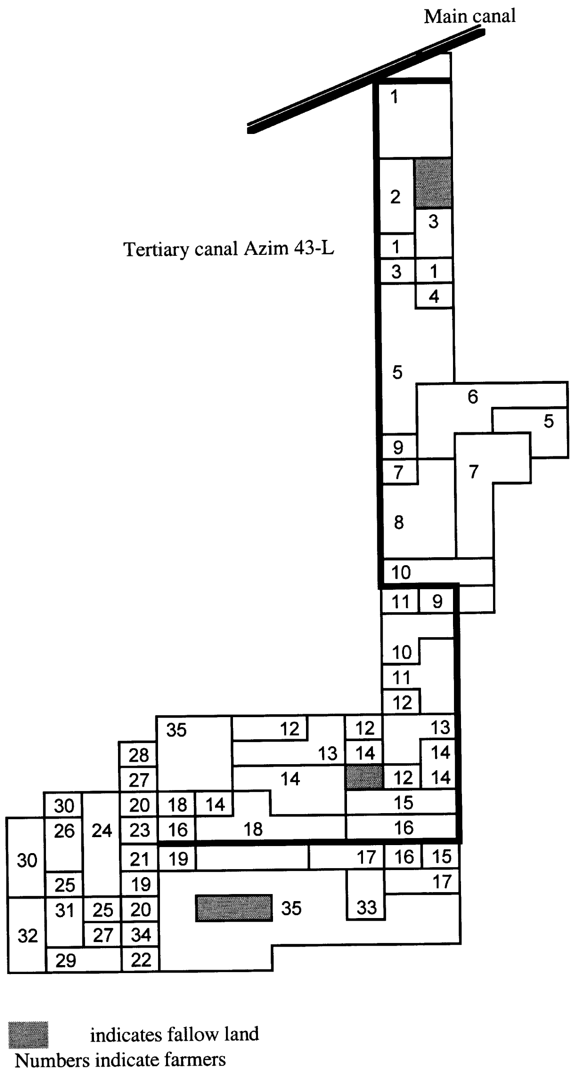

Fig. 4. Command Area Map for Watercourse Azim-43L.

Field interviews indicate that these exchanges are part of an informal arrangement in which farmers receive water in case of immediate shortfall and give water (if asked) in case of relative plenty. All but a few households have exchanged water and most of them have both given water and received water during the course of the observed season. In the Azim 43-L watercourse, households have on average received water 7 times, from 3.5 partners. In the Fordwah 14-R watercourse, households have on average received water 6 times, from 2.2 partners.

Our empirical analysis of patterns of water exchange has two parts. We first examine the evidence that transaction costs affect the formation of exchange links between farmers, and assess whether exchange is organized in clusters within the watercourse rather than as a pure network of bilateral exchanges.[19] Given the evidence we find for cluster formation within exchange networks, the second part of our empirical analysis examines how the size of clusters and quality of exchange within clusters varies with transaction costs.

4.1. Transaction costs and cluster formation in canal water exchange

There are three readily identifiable transaction costs that might affect the configuration of exchange within the watercourse. First, there are association costs due to the need of coordinating transfers with partners at different water turn distances. Water turn distance between two farmers is defined as the number of farmers who have a turn in between the two.[20] The greater the water turn distance between farmers engaged in an exchange, the greater the number of intervening farmers whose turns must be shifted to accommodate the new timing.[21] These costs of coordination rise with water turn distance, irrespective of the size of the transfer as measured by the number of minutes exchanged.

Second, there are association costs proportional to the geographical distance between farms. Very often, transfers are partial turn exchanges in which the receiving household, realizing that it needs additional water, walks over to a partner household to arrange for a turn.[22] Farmers who are close by in terms of water turn distance are not necessarily next to each other in terms of geographical distance (as measured by the minimum distance between their plots). This is most starkly apparent when one considers a water exchange between two households, one located at the head of the watercourse and the other at the tail. Even though these households are adjacent to each other in terms of water turn

[19] Since we do not have data on other types of transfers between these farmers, our analysis is limited to patterns of water exchange only. We, therefore, focus specifically on those transaction costs that are pertinent to exchanging canal water.

[20] Since a farmer can exchange his turn with farmers who have turns either before or after his won turn, we measure the water turn distance as the minimum of the number of intervening water turns when measured clockwise and counter-clockwise. For example, a farmer-pair in which one is at the head and the other at the tail end of the watercourse are considered to be adjacent to each other in terms of water turn distance, even though they are far apart when measured clockwise.

[21] For example, suppose farmer A receives a low water realization and the next farmer in line in the water turn sequence, farmer B, provides a transfer of water time equal to 10 min. Farmer A continues to water his field for an extra 10 min and farmer B delays receiving water for that amount of time. This is a straightforward transfer of water time. Suppose instead, that farmers A and C are in an informal exchange agreement and farmer C is two turns away from farmer A. If C transfers water to A for 10 min, farmer A takes the extra 10 min transferred to him, B receives water 10 min later but takes water for an extra 10 min, and delays receiving water for 10 min. This requires coordinating the activities of three farmers rather than just two as in the first case.

[22] A transfer to secure additional water, even a few minutes, can be extremely important in terms of its value to the household since these small exchanges help complete watering a field, which would otherwise have to be irrigated from scratch again.

distance, walking from the tail-end to the head-end of the watercourse to request a partial turn transfer in the midst of one's *warabandi* turn is impractical.

Finally, among extraction costs, there are costs of enforcement that may limit the effectiveness of a risk-sharing arrangement between farmers. Monitoring costs are unlikely to be important in this system of water exchange. Water deliveries and water needs are readily observable (and verifiable) by all farmers along the watercourse. We have no explicit measure of extraction costs but use kinship relations to identify partnerships within which enforcement difficulties may be more readily resolved. Since costs of establishing exchange links may also be lower among kin, in the empirical analysis that follows, the kinship variable should capture lower extraction costs, and possibly lower association costs as well.

The relationship between these transaction costs—water turn, geographical, and social distance—and the pattern of exchange between potential partners is evident in the transfer matrix for the Azim 43-L watercourse (Fig. 5). Households are numbered in terms of their

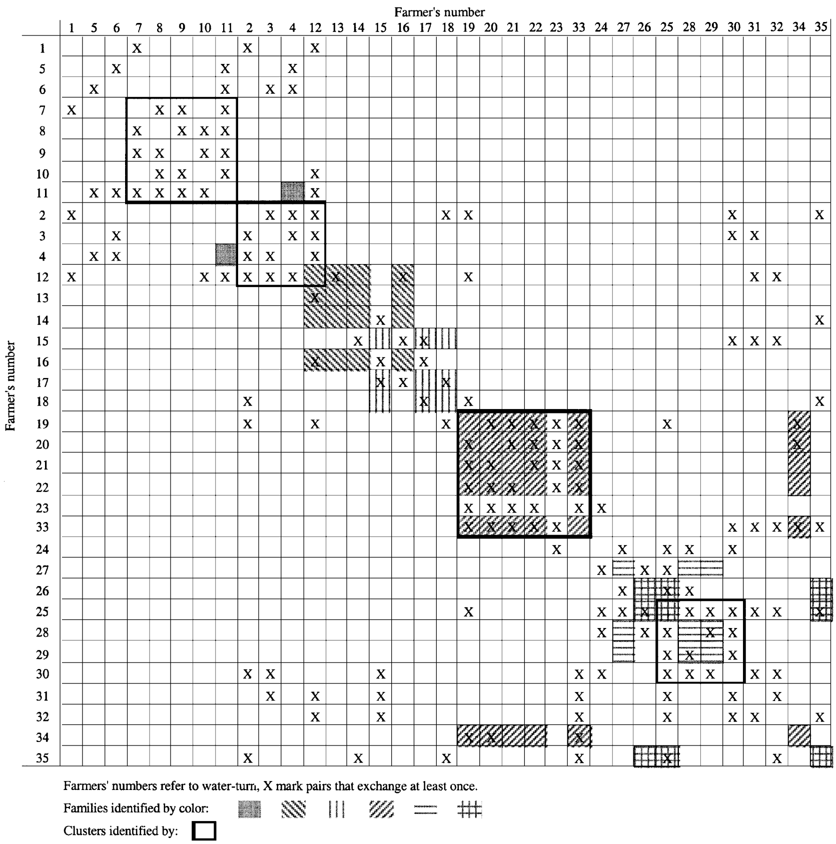

Fig. 5. Matrix of exchanges between farmers of the Azim 43-L watercourse.

water turn, with household 1 being closest to the head of the watercourse and 35 at the tail-end. Cell ij of the matrix records whether households i and j exchanged water at least once during the season. A simple inspection of the transfer matrix reveals preferential exchanges between households who are next to each other in terms of water turn distance (i.e., there is a greater propensity to trade close to the diagonal). Households that are close in terms of water turn distance but far apart in terms of geographical distance (i.e., household pairs where one is at the head-end and the other at the tail-end) do not exchange as much. When kinship relationships are examined (illustrated by the different patterns of shading in Fig. 4), not all kin exchange with each other but there is clearly preferential trading between family members. However, families also tend to be contiguous in terms of water turns, which likely enhances the propensity of kin to exchange. In addition to bilateral exchange, there is also evidence of cluster effects, i.e., with intensive exchanges among members of the group, and few exchanges outside the group, even with neighbors who are close by.[23]

We turn to an econometric analysis of the effect of these transaction costs on the configuration of exchange patterns between households in the two watercourses. In particular, we test the hypothesis that a model of exchange localized within clusters explains the pattern of exchanges better than a model in which exchanges are characterized by purely bilateral relationships. In a bilateral exchange model, the probability of exchange between any two farmers i and j is a function of some measure of transaction cost d_{ij} between the two partners:

$$p(e_{ij} = 1) = f(d_{ij}) \qquad \text{(bilateral exchange model)},$$

where e_{ij} is a binary variable representing whether the pair exchanged water at least once during the season. By contrast, in a model of localized exchange within clusters, exchange between two partners in a cluster depends on the participation of other members in the cluster, and hence is a function of characteristics of the whole cluster. Hence, a pure cluster model could be written as:

$$p(e_{ij} = 1) = f(c_{ij}, c_k) \qquad \text{(cluster model)},$$

where c_{ij} is a binary variable equal to 1 if i and j are members of the same cluster and c_k are characteristics of the cluster k to which i and j belong. The fundamental feature of the bilateral exchange framework is that whether a link is created or not is largely determined by the characteristics of the pair of partners, while in a cluster model, links are determined by the characteristics of all the members of the cluster, not just the two partners.

[23] Note that clusters do not consist solely of water turn neighbors. In the matrix, some households (e.g., 2, 3, 4, and 33) have been placed next to the households they exchange with frequently to uncover a potential cluster configuration.

The econometric model can be written with an index function that nests both models:

$$e_{ij}^{*} = \alpha d_{ij} + (\beta c_{ij} + \gamma c_k) + \varepsilon_{ij}, \quad \text{with } e_{ij} = 1 \text{ iff } e_{ij}^{*} > 0. \tag{9}$$

e_{ij}^{*} is the unobserved latent variable that captures the net benefit of exchange and ε_{ij} is an error term. To ensure that c_k does not capture some unobserved characteristic of the pair ij (and estimation of γ is therefore unbiased), we choose the average distance (water turn, geographical, or social) d among all pairs in the cluster, excluding the pair ij, as our measure of cluster characteristics. A model of pure bilateral exchange is rejected in favor of a model of cluster formation superimposed on the web of bilateral exchanges if the coefficient γ is significantly different from zero.

One potential concern with the estimation of Eq. (9) is that there is no exogenous variable that determines clusters. It is difficult to set cluster boundaries using priors based on the water turn and geographical distance measures since there is no clear technical constraint to trading water beyond a certain distance, and once we allow for overlapping clusters, the number of possible cluster configurations we could potentially test for becomes very large. The only 'prior' that could be used to define cluster boundaries is kinship relations, but as we show later, given the evident importance of water turn and geographical distance in the system of exchange, kinship-based clusters are likely to lack explanatory power.

Because observed cluster characteristics (c_k) are endogenous, we cannot test the pure bilateral exchange model (H_0) against the bilateral-and-observed-cluster model (H_1). However, while the choice of a single configuration of clusters C_0 is endogenous, the set $\{C_s, s \in S\}$ of all potential configurations of clusters is exogenous. We therefore test the bilateral exchange model against the hypothesis (H_2) that there exists *at least one* cluster configuration in addition to the bilateral exchange model. Hence, the model of exchanges under (H_2) comprises bilateral exchanges plus the set of all possible cluster configurations: $\{$(bilateral exchanges and configuration of cluster C_s), $s \in S\}$, i.e., (bilateral exchanges and C_0) or (bilateral exchanges and C_1), etc. Since H_0 is embedded in H_1, itself embedded in H_2, rejecting H_0 against H_1 is sufficient to reject H_0 against H_2. Hence, if we reject the hypothesis that γ is equal to 0 with at least one cluster configuration C_0, then we will have shown that there exists at least one cluster configuration that, in combination with bilateral exchanges, better represents the observed exchanges than would a pure bilateral exchange model.

A potential concern with this test is that since there are a huge number of possible combinations of clusters, even if the true model was a pure bilateral exchange model, the variance of the error term could generate configurations of exchanges that would generate cluster effects. We address this problem as well as the protocol defined to determine the cluster immediately after reporting the results.

Table 1 reports maximum likelihood probit estimates of the probability of exchange between farmer-pairs for the Azim 43-L watercourse. The chosen cluster configuration C_0 consists of the four clusters indicated in Fig. 5. As discussed earlier, costs of exchange between a pair of households are represented by social distance (kinship) and by physical distance (water turn and geographical distance). To identify the role of these characteristics, we need to control for the other determinants of exchange, particularly the differential benefits of insurance according to where the households are located in the

Table 1
Configuration of water exchanges on Azim 43-L watercourse

	Mean value of variables	Bilateral exchange model		Bilateral and cluster exchange model	
		Marginal effect[a]	z[b]	Marginal effect[a]	z[b]
Bilateral exchange model					
Physical distance (d_{ij})					
Water turn distance	9.0	− 0.017	− 5.9	− 0.012	− 3.9
Not geographical neighbors	0.84	− 0.178	− 5.9	− 0.153	− 4.3
Social proximity (d_{ij})					
Same family	0.05	0.078	1.3	− 0.059	− 1.1
Same family[a] × water turn distance	0.28	0.025	3.5	0.018	2.1
Cluster model					
Same cluster (c_{ij})	0.06	–	–	0.823	9.5
Percent of other pairs in cluster that are family members (c_k)		–	–	0.121	2.1
Control variables for the benefits of insurance					
Average length of water turn	4.8	0.003	1.0	0.009	2.8
Location on canal[c]	12.0	0.001	1.2	0.003	2.1
Number of observations	595	595		595	
Pseudo-R^2		0.25		0.36	
Log likelihood		− 192.5		− 153.2	
Test against bilateral and cluster exchange model		chi^2 (2) = 78.7 Prob > chi^2 = 0.0000			

Dependent variable: probability that the household pair traded at least once (probit estimation).

[a] Marginal effects, multiplied by 100, computed at the means of the explanatory variables. Computed as discrete changes for dummy variables (not geographical neighbors, same family, and same cluster).

[b] z statistic of the corresponding parameter. Standard errors obtained by bootstrapping with 1000 repetitions.

[c] Location of the household closest to the head of the canal = continuous variable from 1 to 35, from head to tailenders.

watercourse and the length of their water turns.[24] To identify the impact of cluster characteristics on the probability of exchange between any two partners, we use the percentage of pairs in the cluster (excluding ij) that are kin.[25]

[24] Households located at the tail-end are more likely to benefit from insurance, but at the same time, have a lower capacity to repay with water since their expected water endowment is lower due to seepage losses (no allowance is made in the warabandi schedule to account for greater seepage losses incurred by tail-end households). It is possible that these households repay through some other means. In addition, it is possible that the location variable is capturing the net effect of greater insurance demand on the part of a tail-end households but a lower probability of finding a partner given his lower ability to repay. Households with longer water turns have more flexibility in their water management strategies and are therefore, less likely to demand insurance.

[25] The regressions for both watercourses were estimated using all three – water turn, geographical, social – average distance variables. However, since the variables are highly collinear, the most significant variable was retained for the final specification.

Estimates in Table 1 show that pattern of exchange is affected by water turn and geographical distance, both of which have negative effects on the probability of exchange between households.[26] Kinship is also a powerful explanatory factor. In the bilateral exchange model, being from the same family offsets the negative effect of water turn distance. These results confirm the role of distance, and therefore transaction costs, in limiting exchange among community members.[27] When cluster specific effects are introduced, the intensity of family links in the cluster significantly increases the probability of exchange between any two households within the cluster. The bilateral exchange model is rejected against the model that includes clusters. Several other cluster configurations also result in rejection of the pure bilateral exchange model. Therefore, our test does not prove nor rely on the fact that the particular configuration that we retain is composed of the actual clusters that farmers have formed. The analysis simply shows that there exists at least one configuration of clusters that outperforms the pure exchange model.[28]

As mentioned above, there is a potential concern that any particular cluster configuration could have been generated by random errors in a pure bilateral exchange model. In order to assess this hypothesis, we generated 150 matrices of exchanges with the bilateral term and a random error of mean and variance equal to the mean and variance of $\beta c_{ij} + \gamma c_k + \varepsilon_{ij}$. We then submitted those 150 matrices as well as the observed matrix of exchanges to a common protocol for detection of clusters.[29] The two models of bilateral exchange and bilateral exchange-and-cluster were then estimated for each of the 150 matrices and for four replications of the observed exchange matrix. The 150 values for the Chi^2 statistics of the test of the bilateral exchange model range from 1 to 71, with a mean and a median of 24.8, and the highest 95% at 45.6. So indeed random errors can generate some combination of clusters. Yet, in comparison to these statistics, the Chi^2 for the four replication matrices of the observed exchanges were 57.5, 76.2, 78.4, and 78.7. Hence, three of these values are far above any value for the simulated matrices, and the lowest value is the fourth highest in the group of 150. Evidently, the matrix of observed exchanges has a structure of clusters significantly more pronounced than could have been generated by random errors.

[26] Since there are multiple observations for each household in the data set (i.e., with each potential partner), the disturbances are likely to be correlated across observations. Therefore, standard errors are obtained by bootstrapping the sample with 1000 replications.

[27] Apart from entailing lower costs of risk-sharing, kinship may also support insurance arrangements because of altruistic motives that increase the benefits of participation in the insurance arrangement (Foster and Rosenzweig, 2001).

[28] As mentioned earlier, the only alternative for forming clusters based on exogenous characteristics in kinship. We tested for kinship-based cluster effects but these were not significant.

[29] Specifically, the task was assigned to two students who were asked to detect clusters in the matrices. A user friendly interface allowed them to permute rows and columns and place clusters in the matrix, and to view them on a screen. They had a maximum of 10 min for each matrix to select the best combination of clusters. We included the observed matrix of exchanges twice for cluster indentification, interspersed between other matrices, without telling the students. The students did not seem to recognize matrices that they had previously treated, suggesting that these two trials can be considered to be independent.

Table 2
Configuration of water exchanges on F14-R watercourse

	Mean value of variables	Bilateral exchange model		Bilateral and cluster exchange model	
		Marginal effect[a]	z[b]	Marginal effect[a]	z[b]
Bilateral exchange model					
Physical distance (d_{ij})					
Water turn distance	23.8	− 0.09	− 7.8	− 0.04	− 3.7
Not geographical neighbors	0.91	− 1.18	− 7.0	− 0.46	− 1.7
Social proximity (d_{ij})					
Same family	0.04	1.28	3.1	0.41	1.2
Same family[a] × water turn distance	0.61	0.04	2.3	0.02	0.9
Cluster model					
Same cluster (c_{ij})	0.09	–	–	9.56	11.2
Average distance among other pairs in cluster (c_k)		–	–	− 0.04	− 3.4
Control variables for the benefits of insurance					
Average length of water turn	1.8	− 0.001	0.0	0.03	0.8
Location on canal[c]	31.7	0.003	1.8	0.0003	0.2
Number of observations		4371		4371	
Pseudo-R^2		0.32		0.43	
Log likelihood		− 477.3		− 388.7	
Test against bilateral and cluster exchange model		chi^2 (2) = 177.3 Prob > chi^2 = 0.0000			

Dependent variable: probability that the household pair traded at least once (probit estimation).

[a] Marginal effects, multiplied by 100, computed at the means of the explanatory variables. Computer as discrete changes for dummy variables (not geographical neighbors, same family, and same cluster).

[b] z statistic of the corresponding parameter. Standard errors obtained by bootstrapping with 1000 repetitions.

[c] Location of the household closest to the head of the canal = continuous variable from 1 to 94, from head to tailenders.

The model of bilateral exchange with clusters explains water exchanges better than the pure bilateral exchange model in the Fordwah 14-R watercourse as well (Table 2).[30] These results confirm the finding that clusters play an important role in sustaining mutual insurance. It should be emphasized that the fact that households exchange in small groups is not, by itself, evidence that risk is shared inefficiently within the watercourse (see also Fafchamps and Lund, 2000). As long as groups overlap so that no households are left out, water transfers could flow across clusters to pool risk across the entire watercourse.[31]

[30] In Fordwah, the same method suggests the existence of seven clusters. The distance variable retained for characterizing the clusters is the average physical distance among the pairs in the cluster.

[31] Note that the econometric specification does not preclude the existence of overlapping clusters.

Since we do not have data on water shocks and endowments, we cannot measure the degree to which cluster-based exchange reduces risk-pooling at the watercourse-level, or the quality of insurance within the clusters. The objective of the empirical work in the next section is therefore more modest: to assess whether cluster size and the intensity of exchange respond to transaction costs as suggested by the theory.

4.2. Cluster size and intensity of exchange

Given the evidence of cluster formation within exchange networks, we examine how the size of clusters and the intensity of exchange (as measured by the percentage of pairs that exchange at least once) within clusters vary with transaction costs. For the two watercourses, we use the 11 clusters specified earlier in Section 3 (four in the smaller Azim 43-L and seven in the larger Fordwah14-R). Overall, 49.6% of the farmers belong in an insurance cluster, with a slightly greater degree of participation in Azim (54.2%) than in Fordwah (47.8%) (Table 3); 44.2% of all pairs of farmers that exchange at least once belong in clusters, and 46.8% of all water transfers are done within a cluster. Clusters vary in size from 4 to 11 partners, with an average of 5.9; 20.9% of the pairs of households in clusters are kin, while overall in these two communities, only 3.8% of the pairs are kin. Within clusters, more than two-thirds of the pairs exchanged water at least once in the season.

The theory developed above predicts that cluster size decreases with association costs and that the quality of insurance decreases with extraction costs. Both the cluster size and the quality of insurance are predicted to increase with exposure to risk. These predictions cannot be tested for the selected clusters (for example, by examining the impact of member characteristics on cluster size) since the cluster itself and the selection of its members are endogenous. Therefore, we approach this choice from an individual's

Table 3
Characteristics of the exchange clusters

Participation in exchange clusters		Watercourse		Overall
		Azim 43-L	Fordwah 14-R	
Percentage that belong in clusters:				
Farmers		54.2	47.8	49.6
Pairs of farmers that exchange at least once		38.7	47.3	44.2
Exchanges		41.6	49.0	46.8
Characteristics of the 11 clusters	**Average**	**Standard deviation**	**Minimum**	**Maximum**
Number of members	5.9	2.0	4	11
Percent of pairs that are kin	20.9	23.5	0.0	66.7
Average water turn distance	5.1	4.6	1.7	16.9
Average geographical distance	2.3	2.4	0.2	8.2
Total number of exchanges	33.5	35.6	14	138
Percent of pairs that exchange at least once	72.9	22.5	40.0	100
Average number of exchanges per pair	1.5	0.4	0.7	2.2

point of view.[32] Each farmer is characterized by indicators of distance to potential partners and of exposure to risk. Geographic distance, assumed to be positively correlated to association costs A, is characterized by the average geographical distance to the closest four water turn neighbors. Social distance, assumed to be correlated with both association costs A and extraction costs E, is characterized by the number of family members that each farmer has along the canal. Risk exposure is captured positively by location along the canal and negatively by the length of water turn. Since potential partners are themselves heterogeneous, in reality each farmer will choose not only the size of his cluster but also his specific cluster partners, as a function of his costs of extraction, costs of association, and exposure to risk. The matching that is actually obtained obviously has to be a compromise between the ideal cluster that each one would want for himself and the difficulty of finding partners. Nevertheless, we expect to find that each individual is in a group that corresponds to his desired cluster in terms of size and quality of insurance.

Table 4 reports regression results on the determinants of cluster size and intensity of exchange at the individual level, i.e., for the 129 farmers that are on either one of the two watercourse command areas. Cluster size and intensity of exchange are defined as follows:

$$n_i = N_c \text{ if } i \in \text{ cluster } c, \quad 0 \text{ if } i \text{ does not belong to any cluster,}$$

$$I_i = M_c/[N_c(N_c - 1)/2] \text{ if } i \in \text{ cluster } c, \quad 0 \text{ if } i \text{ does not belong to any cluster,}$$

where N_c is the number of members of cluster c and M_c is the number of pairs that exchange at least once in cluster c. Approximately 50% of the households do not belong to any cluster, and therefore, have zero values for both variables. Because of the preponderance of zero values and the discrete nature of cluster size, the cluster size regression is estimated with a negative binomial regression model (the Poisson model, which assumes overdispersion equal to 0, was rejected). The intensity of exchange regression is estimated with Tobit maximum likelihood to account for the censoring.

In the negative binomial regression, the conditional mean and variance of the number of members are related to the exogenous variables x, the parameters β, and the parameter of overdispersion α, as follows:

$$E(n_i \,|\, x_i) = e^{\beta' x_i} \quad \text{and} \quad \mathrm{var}(n_i \mid x_i) = e^{\beta' x_i}(1 + \alpha e^{\beta' x_i}).$$

The regression for cluster size suggests that it is an increasing function of the number of family members and decreases with geographical distance. Predicted cluster size for an average farmer in Azim 43-L is 2.4 members. This size ranges between 2.0 and 3.1 for a farmer with no kin and with 5 kin (the maximum value observed), respectively. Estimates predict that the most isolated farmer would not belong to a cluster, while the farmer with closest neighbors (geographically) in this watercourse would belong to a cluster of nearly 3 members. With the shortest water-turn, the average farmer is predicted to belong in a

[32] See Sadoulet and Carpenter (1999) for a similar approach for analyzing the formation of credit groups.

Table 4
Determinants of the size and intensity of exchange in clusters

	Mean value of variable on watercourses		Number of members in cluster (Negative binomial regression)		Percent of pairs that exchange (Double bound Tobit)	
	A43	F14	Coefficient	z^a	Coefficient	z^a
Transactions costs						
Number of family members along the canal (-A, -E)	1.8	3.3	0.09	2.3	0.065	2.2
Average distance to the closest four neighbors (A)	15.5	4.1	− 0.29	− 2.0	− 0.12	− 1.9
Risk exposure						
Location on Azim 43-L watercourse	18	–	− 0.03	− 1.3	− 0.02	− 1.1
Location on Fordwah 14-R watercourse	–	47.5	0.004	0.8	− 0.002	− 0.7
Own water turn length	4.8	1.8	− 0.169	− 3.0	− 0.112	− 2.6
Fordwah 14-R watercourse effect	–	1	− 0.91	− 2.0	− 1.27	− 2.8
Intercept	1	1	2.15	4.9	1.38	3.1
Parameter of overdispersion			2.34	4.6		
Number of observations	35	94	129		129	
Number of left-censored observations (at 0)					65	
Number of right-censored observations (at 1)					14	
Wald chi^2 (6)			20.9		21.6	

[a] z statistics based on robust standard errors.

cluster of 4.9 members while with the longest water-turn (corresponding to 50 acres) he would not participate in a cluster.[33]

Empirical results on the intensity of exchange also confirm some of the predictions of the theory.[34] For the average farmer in Azim 43-L, the intensity of exchange within the cluster is predicted to be 48%. This ranges between 36% and 69% for farmers who have no kin and 5 kin, respectively. It increases to 96% for the cluster of the farmer with the shortest water turn. For the Fordwah 14-R watercourse, where farmers have more kin and live closer to each other but face lower uncertainty in water delivery, the predicted cluster size is larger but the intensity of exchange is predicted to be lower.

[33] Location on the canal is not significant. As mentioned earlier, the location variable may be capturing the net effect of greater insurance demand on the part of a tail-end household but a lower probability of finding a partner given his lower ability to repay. As a result, the coefficient on location can be of either sign.

[34] Using the mean of the latent variable, $\gamma' x_i$, where x represents the exogenous variables, and γ the vector of parameters.

5. Concluding remarks

This paper has drawn attention to the role of transaction costs in shaping both the configuration of insurance-motivated exchanges and the quality of mutual insurance. We have formalized the determination of the number of partners in an informal insurance agreement and the degree of risk pooling as a function of different types of transaction costs: costs of establishing associations that are fixed per partner, and extraction costs to implement income transfers that are a function of the level of insurance. Which of these costs are important in a particular context depends on the specific purpose of the insurance system, the community characteristics, etc. The specificity of these costs influences the size of the cluster as well as the optimal set of state-contingent transfers. In particular, high association costs combined with low extraction costs will lead to clusters with full insurance, while low association costs with high extraction costs will lead to community-level partial insurance. The first case in particular implies that tests of risk pooling at the village level, or other reference group, may be refuted because the village is not necessarily the logical reference group.

Data from water transfers along two watercourses in Pakistan, where water delivery is subject to idiosyncratic random shocks, show that households cope with variability by exchanging water bilaterally with neighbors and family members, and also with members of tightly knit clusters. This pattern of exchange suggests that transaction costs are important in localizing exchange by shaping the choice of insurance partners. We also find that the size of the mutual insurance cluster decreases with association costs while the intensity of exchanges decreases with extraction costs.

The results of this paper raise several considerations for policy. First, they suggest that to improve the quality of informal insurance, it is imperative to focus policy levers on reducing association and extraction costs, which lead to lower insurance in the group of interest: lower association costs allow higher quality insurance by broadening insurance clusters to the full community; lower extraction costs allow full instead of partial insurance in the chosen insurance group. These costs can be reduced by improving the flow of information among community members. This can be done by centralizing information in simple clearinghouse mechanisms and/or by facilitating decentralized two-way communications through for instance cellular phones. Second, evidence of kinship as an organizing principle for self-selected clusters points to the relevance of social capital in promoting informal risk sharing. Risk sharing takes place preferentially in social arenas that facilitate rapid information flows, impose norms of fairness and reciprocity, and apply social sanctions on defaulting parties. These functions of local institutions are nontrivial because they solve inherent problems of coordination, asymmetric information, and contract enforcement that can be prohibitively costly for outsiders to solve. If effective mutual insurance is to transcend clusters of kin, social capital needs to be created. The role of outside agents to help build community social capital beyond local descent groups has been explored by the World Bank (1997), Durston (1999), Arnold (2001), and Abraham and Platteau (2001). While there is no single recipe, and the time dimension can be substantial, key elements of this undertaking include the rediscovery of ancestral norms of reciprocity in the community, the training of community members in the basic principles of governance of local organizations, and the provision of new opportunities to derive short-term material benefits from the use of social capital.

Acknowledgements

For their comments and suggestions, we are indebted to Jean-Louis Arcand, Pranab Bardhan, Peter Berck, Pierre Dubois, Martin Fournier, Hanan Jacoby, Ethan Ligon, Pedro Olinto, Martin Ravallion, Loïc Sadoulet, Paul Seabright, Brian Wright, and seminar participants at CERDI in Clermont-Ferrand, the World Bank, ARQADE in Toulouse, and the University of California at Berkeley. We thank Frederico Finan for assistance in setting the system for the visual analysis of matrices. We are grateful to Pierre Strosser and to the IWMI staff in Lahore and at the field station in Hasilpur for their assistance and for the data used in this analysis. The views expressed in this paper are those of the authors alone; they do not necessarily represent the view of the World Bank, its Executive Directors, or the countries they represent.

Appendix A. Derivation of insurance regimes

The insurance regimes in Fig. 3 are derived from the optimization problem:

$$\begin{aligned}Eu(\alpha,n) = {} & u(\mu) + \frac{1}{2}u''(u) \\ & \times \left\{\pi(1-\pi)L^2\left[1+n(n-1)\alpha^2 - 2(n-1)\alpha\right]\right\} \\ & - E\alpha - A(n-1),\end{aligned}$$

such that $0 \leq \alpha \leq 1$, $1 \leq n \leq N$, and $Eu(\alpha,n) \geq Eu(0,1)$.

The first-order conditions for the choice of α and n for an interior solution are given by:

$$k(n-1)(1-n\alpha) = E \quad \text{and} \quad \frac{k}{2}\alpha(\alpha+2-2n\alpha) = A,$$

$$\text{where} \quad k = u'\phi\sigma_0^2 = -u''\pi(1-\pi)L^2.$$

(1) Boundary between regions 1 and 2 (on the horizontal axis) is the case when the household is indifferent between full insurance at the community level and full insurance in a cluster marginally smaller than the community. It is characterized by the first-order conditions for α and n when extraction costs (E) are zero and $n = N$. That is, $\alpha = 1/N$, and $\{k/2N\}\{\alpha + 2 - 2N\alpha\} = A$. This implies that $A = k/2N^2$.

(2) Boundary between regions 2 and 5 (on the horizontal axis) is the case when the household is indifferent between full insurance in a cluster and autarky. It is characterized by the first-order conditions for α and n when extraction costs are zero and $n = 1$. That is, $\alpha = 1/n$, $\{k/2\}\alpha\{\alpha + 2 - 2n\alpha\} = A$, and $n = 1$. This implies that $A = k/2$.

(3) Boundaries between regions 3 and 4, and between regions 3 and 5 are cases when the household is indifferent between partial community-level insurance and partial insurance in clusters marginally smaller than the community (boundary between 3 and 4), and when it is

indifferent between partial community-level insurance and autarky (boundary between 3 and 4).

Indifference between partial community-level insurance and partial insurance in clusters marginally smaller than the community is derived from the first-order conditions for α and n when $n=N$. These are $k(N-1)(1-N\alpha)=E$ and $(k/2)\alpha(\alpha+2-2N\alpha)=A$. Together, these conditions imply a curve between E and A when a household that faces positive extraction costs and association costs chooses $n=N$ and an optimal α. The curve is characterized by $A=\{k/2\}\{(k(N-1)-E)/(k(N-1)N)^2\}\{2EN+k(N-1)-E\}$, with A an increasing function of E in the range $0 \leq E \leq \{k(N-1)^2/2N-1\}$ and a decreasing function beyond this range. When association costs are zero, $E=k(N-1)$, and when extraction costs are zero, $A=k/2N^2$, as derived in case 1. In the range where A is an increasing function of E, the cluster size and transfer arrangement are complements. The complementarity arises from the structure of variance. In this region, the marginal benefit of risk reduction from changing cluster size is positive with respect to α. That is, $-(d/d\alpha)(d\sigma^2/dn)>0$. Beyond this region, as E increases to $k(N-1)$, A decreases as the cross partial of variance with respect to α and n changes sign. Therefore α and n are substitutes in this part of the boundary.

In addition to this relationship between E and A, the optimal solutions are characterized by a participation constraint which requires that utility from insurance exceeds the utility a household retains in autarky. The constraint, when combined with the first-order conditions for α and n, and the condition that $n=N$, implies that $\{E-k(N-1)\}^2-2Ak(N-1)^2N=0$. When association costs are zero, $E=k(N-1)$ in the positive vertical axis, and when extraction costs are zero, $A=k/2N$. In between these points, the curve is downward sloped and convex. For $E>k(N-1)$, the relationship between E and A is upward sloped. However, for these values of E, the optimal α is negative, which violates the constraint that α lies between 0 and 1.

Between regions 3 and 5, the household is indifferent between partial community-level insurance and autarky. Therefore, the boundary is defined by the first-order condition for α as well as the participation constraint when $n=N$. Along this boundary, utility is maintained at autarky levels. An increase in association costs induces a smaller cluster size. Utility is restored to autarky levels only when insurance provision increases within the cluster. This occurs as extraction costs decrease.

In contrast, between regions 3 and 4, the household is indifferent between partial community-level insurance and partial insurance in a cluster; the boundary is determined only by the two first-order conditions when $n=N$. Along this boundary, an increase in association costs creates an incentive to reduce cluster size. A household would be willing to maintain cluster size at N only if extraction costs were higher (and, therefore, lower insurance provided within the cluster).

(4) Boundary between regions 4 and 5 is the case when the household is indifferent between partial insurance in clusters and autarky. This boundary is obtained when the first-order conditions for α and n, and the participation constraint are met. That is, $k(n-1)(1-n\alpha)=E$, $(k/2)\alpha(\alpha+2-2n\alpha)=A$, and $Eu(\alpha,n)=Eu(0,1)$. With some algebra, the three conditions can be reduced to $n\alpha^2=2A^*$; $n\alpha=(1/3)(\alpha+2)$; and $(2/3)(1-\alpha)-\alpha E^*-2(n-1)A^*=0$ where $E^*=E/k$ and $A^*=A/k$. Total differentiation of these three conditions with respect to α, n, E^*, and A^* can be used to show that E and A are negatively

related along this boundary, suggesting that cluster size and transfer arrangements are substitutes in the region.

As a final point, note that the boundary between regions 4 and 5 intersects the curves for case 3 at the same point. Here, conditions that satisfy all three curves are met: $n = N$, $Eu(\alpha,n) = Eu(0,1)$, $\alpha = (k(n-1) - E)/(k(n-1)n)$. For case 3, this point is dictated by the participation constraint and the constraint that cluster size not exceed the community size. On the other hand, for case 4, this point is a result of the community-size and the corresponding α being chosen as the optimal *interior* solution to the maximization. For this case, as extraction costs increase further, the optimal interior choice of n is greater than N, making that range infeasible.

References

Abraham, A., Platteau, J.-P., 2001. Participatory development in the presence of endogeneous community imperfections. Department of Economics and Centre de Recherche en Economie du Développement, University of Namur.

Arnold, J.E.M., 2001. Devolution of control of common pool resources to local communities: experiences in forestry. In: de Janvry, A., et al., (Eds.), Access to Land, Rural Poverty, and Public Action. Oxford Univ. Press, Oxford, pp. 163–195.

Bandaragoda, T., ur Rehman, S., 1995. Warabandi in Pakistan's canal irrigation systems: widening gap between theory and practice. International Water Management Institute Country Paper, No. 7.

Bardhan, P., Udry, C., et al., 1999. Development Microeconomics. Oxford Univ. Press, Oxford.

Coate, S., Ravallion, M., 1993. Reciprocity without commitment: characterization and performance of informal insurance arrangements. Journal of Development Economics 40, 1–24.

Deaton, A., 1992. Household savings in LDCs: credit markets, insurance, and welfare. Scandinavian Journal of Economics 94, 253–273.

Durston, J., 1999. Building community social capital. CEPAL Review 69 (December), 103–118.

Ellsworth, L., 1989. Mutual insurance and non-market transactions among farmers in Burkina Faso. PhD Thesis, University of Wisconsin.

Evans-Prichard, E.E., 1940. The Nuer. Clarendon Press, Oxford.

Fafchamps, M., 1992. Solidarity networks in preindustrial societies: rational peasants with a moral economy. Economic Development and Cultural Change 41, 147–174.

Fafchamps, M., Lund, S., 2000. Risk sharing networks in rural Philippines. Mimeo, Department of Economics, University of Oxford.

Foster, A., Rosenzweig, M., 2001. Imperfect commitment, altruism, and the family: evidence from transfer behavior in low-income rural areas. Review of Economics and Statistics 83 (3), 389–407.

Gaynor, M., Gertler, P., 1995. Moral hazard and risk spreading in partnerships. RAND Journal of Economics 26 (4), 591–613.

Gertler, P., Gruber, J., 2001. Insuring consumption against illness. American Economic Review, in press.

Goldstein, M., 2000. Chop time, no friends: intrahousehold and individual insurance mechanisms in Southern Ghana. Mimeo, Development Studies Institute, London School of Economics.

Grimard, F., 1997. Household consumption smoothing through ethnic ties: evidence from Côte d'Ivoire. Journal of Development Economics 53 (2), 391–422.

Hadar, J., Russell, W.M., 1969. Rules of ordering uncertain prospects. American Economic Review 6 (2), 65–73.

Hoff, K., 1997. Informal insurance arrangements: an equilibrium analysis. Presented at the 1997 Meetings of the Econometric Society, New Orleans.

Kimball, M., 1988. Farmers' cooperatives as behaviors towards risk. American Economic Review 78, 224–232.

Kranton, R., 1996. Reciprocal exchange: a self-sustaining system. American Economic Review 86 (4), 830–851.

Jalan, J., Ravallion, M., 1999. Are the poor less well-insured? Evidence on vulnerability to income risk in rural China. Journal of Development Economics 58, 61–81.

Ligon, E., 1998. Risk-sharing under varying information regimes: theory and evidence in village economies. Review of Economic Studies 65 (4), 847–864.

Ligon, E, Thomas, J.P., Worrall, T., 1999. Mutual Insurance and Limited Commitment: Theory and Evidence in Village Economies. Working Paper, Department of Agricultural and Resource Economics, University of California at Berkeley.

Newbery, D., Stiglitz, J., 1981. The Theory of Commodity Price Stabilization: A Study in the Economics of Risk. Clarendon Press, Oxford.

Platteau, J.P., 1991. Traditional systems of social security and hunger insurance: past achievements and modern challenges. In: Ahmad, E., Dreze, J., Hills, J., Sen, A. (Eds.), Social Security in Developing Countries. Clarendon Press, Oxford.

Posner, R., 1980. A theory of primitive society with special reference to law. Journal of Law and Economics 23, 1–53.

Rashid, M., 1990. Rural Consumption: the Evidence from Pakistan. Dissertation. University of Chicago.

Rosenzweig, M.R., Stark, O., 1989. Consumption smoothing, migration, and marriage: evidence from rural India. Journal of Political Economy 97 (4), 905–926.

Sadoulet, L., Carpenter, S., 1999. Risk-matching in credit groups: evidence from Guatemala. Mimeo, Free University of Brussels.

Scott, J.C., 1976. The Moral Economy of the Peasant: Rebellion and Subsistence in Southeast Asia. Yale University Press, New Haven.

Strosser, P., Kuper, M., 1994. Water Markets in the Fordwah/Eastern Sadiqia Area. Working Paper No. 30, International Water Management Institute, Colombo.

Townsend, R., 1994. Risk and insurance in village India. Econometrica 62 (3), 539–591.

Townsend, R., 1995. Consumption insurance: an evaluation of risk-bearing systems in low-income economies. Journal of Economic Perspectives 9 (3), 83–102.

Udry, C., 1994. Risk and insurance in a rural credit market: an empirical investigation in Northern Nigeria. Review of Economic Studies 63, 495–526.

World Bank, The. 1997. The World Bank Participation Sourcebook. Washington D.C.