

Risk Sharing and the Demand for Insurance: Theory and Experimental Evidence from Ethiopia

Erlend Berg, Michael Blake, and Karlijn Morsink*

June 2018

Abstract

Households, organisations and governments commonly engage in risk sharing. The residual risk, however, is often considerable. In response, many policy makers consider the introduction of parametric or index insurance. This raises the question of how demand for insurance depends on the extent of pre-existing risk sharing. We contribute to the literature in two ways. First, we present a simple model to analyse demand for both standard indemnity insurance and index insurance in the presence of risk sharing. Second, we conduct an artefactual field experiment with Ethiopian farmers, in which the predictions of the theoretical framework are borne out.

JEL: D14, D81, G22, O16

Key words: risk sharing, indemnity insurance, index insurance

1. INTRODUCTION

Households, organisations and governments commonly share risk through transfers and emergency loans, but the residual risk tends to be substantial. This is especially so in the presence of aggregate shocks, because actors who share idiosyncratic losses are typically simultaneously affected and therefore unable to fully insure each other (Samphantharak and Townsend, 2018). Notably, weather risk and natural disasters remain important sources of vulnerability, particularly in low-income countries. In response, there is increasing interest in insurance, for both governments and households, as a policy tool (Clarke and Dercon, 2016).

*Berg: Department of Economics, University of Bristol, Priory Road Complex, Bristol BS8 1TU, United Kingdom (e-mail erlend.berg@bristol.ac.uk, tel +441173317911); Blake: International Growth Centre, London School of Economics and Political Science, 4th Floor, 32 Lincoln's Inn Fields, Houghton Street, London WC2A 2AE, United Kingdom (e-mail: michael.blake@bsg.ox.ac.uk); Morsink: Department of Economics, University of Oxford, Manor Road Building, Oxford OX1 3UQ, United Kingdom (e-mail: karlijn.morsink@economics.ox.ac.uk). We are indebted to Tagel Gebrehiwot Gidey for research coordination. We are grateful for useful comments from Glenn Harrison, Mahreen Mahmud, Costas Meghir, Kalle Moene, Ollie Ramsey, Debraj Ray, Jim Roth, Elisabeth Sadoulet and audiences at the University of Bristol, the University of Kent, the 2017 CSAE conference, 2017 SEEDEC conference, and 2017 DIAL conference. This paper is an output of iiG, a research programme funded by the UK Department for International Development (DFID) to study institutions for pro-poor growth in Africa and South Asia. Views expressed are not necessarily those of DFID.

Standard indemnity insurance, which tends to be costly and fraught with asymmetric-information problems, has increasingly been supplemented by parametric or index insurance, where payouts are based on objectively verifiable indices such as rainfall measures.¹ Index insurance may improve welfare relative to indemnity insurance because it can lower costs and reduce market imperfections if payouts are sufficiently highly correlated with actual losses (Clarke, 2016). However, to appreciate the welfare implications of introducing insurance of either type, it is important to take account of its interaction with pre-existing risk-sharing arrangements (Arnott and Stiglitz, 1991; Cutler and Gruber, 1996; Attanasio and Ríos-Rull, 2000; Albarran and Attanasio, 2003). Welfare is likely to increase when formal insurance complements pre-existing arrangements, but may decrease when it substitutes them. In this paper we study how demand for both indemnity and index insurance depends on the extent of risk sharing.

The paper makes two main contributions. The first is to model both types of insurance in a single, simple theoretical framework. We show that an increase in the level of risk sharing decreases demand for indemnity insurance but increases demand for index insurance. Intuitively, this is because the greater the extent of risk sharing, the more agents will need to be insured against shocks that are common to the group, which are better covered by index insurance, rather than shocks specific to the individual, which are better covered by indemnity insurance. Risk sharing is a complement to index insurance because the former pools individual-specific shocks, including the possibility that index insurance does not pay out when an individual suffers a loss.

The second main contribution is to investigate these relationships empirically by exogenously varying risk sharing. We do so in an artefactual field experiment with farmers in Ethiopia and study demand for both indemnity and index insurance. The advantage of this approach is that it avoids the endogeneity concerns arising from the use of observational data and, hence, rigorously identifies a causal relationship between the extent of risk sharing and demand for both types of insurance. The experimental results bear out the predictions of the theory: an increase in risk sharing increases the number of units of index insurance purchased, while it decreases the number of units of indemnity insurance purchased. The findings suggest that interventions aiming to strengthen or encourage the formation of risk-sharing arrangements can increase demand for index insurance and overall insurance coverage. Interventions that build on informal financial institutions to improve access to formal finance have proven successful in the case of micro-credit (Ghatak, 1999; De Quidt, Fetzer and Ghatak, 2016; Ahlin, 2016).

¹Index insurance is also referred to as probabilistic insurance (Kahneman and Tversky, 1979; Segal, 1988; Wakker, Thaler and Tversky, 1997; McIntosh, Povel and Sadoulet, 2015). One prominent example of index insurance is Kilimo Salama, which covers 180,000 Kenyan farmers against drought and excess rainfall. Another is the Caribbean Catastrophe Risk Insurance Facility (CCRIF), set up by 16 Caribbean governments, which offers country-level index insurance against devastating hurricanes and earthquakes. In both cases, insurance is offered to stakeholders who typically already share risk through pre-existing arrangements.

As far as we know, this is the first study to exogenously vary the level of risk sharing in order to test its impact on insurance demand. This is clearly attractive from an identification perspective.² It is also directly relevant to empirical settings where the level of risk sharing is relatively fixed because, for example, it is explicitly or implicitly contracted and deviation is associated with social sanctions. Such contractual rigidity has been widely documented in traditional risk-sharing arrangements (Besley, Coate and Lounsbury, 1993; Dercon et al., 2006; Berg, 2018), suggesting that an exogenous level of risk sharing may be a reasonable assumption in a range of contexts, at least in the short term.³ The risk-sharing groups (*iddir*) in our sample all had written rules for members, including regulations on the level of contributions and punishment for failure to comply.

In line with Dercon et al. (2014), we abstract from asymmetric information issues, in particular moral hazard over the risk-taking behaviour of members of the risk-sharing group. We believe that the interaction of risk sharing and insurance under full information is interesting in its own right. It is also empirically relevant: many risk-sharing groups, especially in the context of small-scale agriculture, consist of members from close-knit rural communities where effort and risk taking are observable in practice. Several risk-sharing groups in our sample even report that they institute committees that visit the agricultural land of its members to advise on agricultural production strategies. When effort is observable by members of a risk-sharing arrangement, moral hazard issues are likely to be of secondary concern to welfare (Arnott and Stiglitz, 1991).

Our paper builds on a substantial literature concerning risk sharing and insurance. Early contributions showed that real-life risk-sharing arrangements generally fall short of providing full consumption smoothing (Cochrane, 1991; Mace, 1991; Townsend, 1994). Imperfect risk sharing may arise from an inability to insure aggregate shocks, limited commitment or information asymmetry (Coate and Ravallion, 1993; Attanasio and Pavoni, 2011). The role of insurance in this context is the subject of several papers including Arnott and Stiglitz (1991), Boucher and Delpierre (2014), De Janvry, Dequiedt and Sadoulet (2014) and Janssens and Kramer (2016). There is also a growing body of research on the demand for index insurance (Cole et al., 2013; Elabed et al., 2013; Clarke, 2016; Casaburi and Willis, 2017).

Theoretically, this paper relates to recent work on the complementarity or substitutability of risk sharing and insurance. Our set-up is similar in spirit to Dercon et al. (2014), and we generate similar predictions, but in a considerably simpler model. In

²The interaction between risk sharing and insurance is complex, as there are potential effects in both directions: the extent of risk sharing may influence the demand for formal insurance, but risk sharing may also be endogenous to the introduction of formal insurance. The experimental approach to studying such complex systems is to exogenously vary one part, and observe how the others respond. In this sense, our design is complementary to that of Takahashi, Barrett and Ikegami (2018), who experimentally vary the price of index insurance and study its effect on measures of risk sharing.

³This is distinct from modalities of risk sharing where transfers and gifts are not contracted (Fafchamps and Lund, 2003).

particular, in the case of actuarially fair premia, we make no assumptions on the utility function beyond non-satiation and risk aversion, and we derive intuitive, closed-form solutions. In the case with marked-up premia and CRRA utility, we show that optimal insurance is linear in risk sharing and derive simple, analytical expressions for their slopes.

Mobarak and Rosenzweig (2012) also model the demand for index insurance in the presence of risk sharing. However, they do not consider indemnity insurance. In a setup with a pair of risk-sharing agents, they model aggregate and idiosyncratic risks as independent, additive shocks. They conclude that risk sharing *may* increase the demand for insurance when basis risk is high, but they are unable to sign the effect unambiguously. By contrast, we model risk as a single shock occurring with a probability that depends on the aggregate state (the index). This simplifies the analysis and allows basis risk to arise naturally within the framework rather than representing an additional assumption. The simplicity of the framework permits analytical solutions for demand (in the case of actuarially fair premia) or the slope of the demand curve (with marked-up premia).

We also contribute to the empirical literature on the complementarity or substitutability of risk sharing and insurance. It is known that indemnity insurance can crowd out risk sharing in both full and limited-commitment risk-sharing relationships (Cutler and Gruber, 1996; Attanasio and Ríos-Rull, 2000; Albarran and Attanasio, 2003), but the relationship between index insurance and risk sharing has received less attention. Mobarak and Rosenzweig (2012) randomly vary index insurance offers and prices, but not the extent of risk sharing. They estimate the extent to which different Indian subcastes indemnify risk, but in their specification the direct contribution of risk sharing to index insurance demand is not identified when subcaste fixed effects are included.⁴ They do not study indemnity insurance.

McIntosh, Povel and Sadoulet (2015) investigate willingness to pay for index insurance when risk sharing through a cooperative is possible and find no significant impact on demand. They attribute this in large part to enforcement problems in the ex-post risk-sharing arrangement of the cooperative: because individuals are not expecting much group risk-sharing to actually occur, they do not increase index insurance demand when in a group. This is an important finding, but leaves open the question of whether individuals would purchase more insurance if they *did* expect group risk-sharing to occur. In our model and experiment, we exogenously fix the extent of risk-sharing to abstract from enforcement problems.

Dercon et al. (2014) offer index insurance products to Ethiopian risk-sharing groups.

⁴They therefore identify the effect of risk sharing on index insurance demand indirectly, through variation in the level of basis risk in the index insurance products offered. This is achieved by randomly placing weather stations in some villages rather than others. However, the extent of exogenous variation that this generates is small: out of 19 villages, they place a weather station in 12 and no weather station in 7. Weather stations do not appear to be randomly placed within the 12 villages, so subcaste-level randomisation is not achieved, and they are unable to rule out subcaste-level confounders that vary with distance to the weather stations.

In their experiment, all participants are subject to a marketing intervention on index insurance, but in some treatment arms, risk sharing as a mechanism to manage basis risk is particularly emphasised. While they find that this increases demand for index insurance, their intervention falls short of exogenously varying the extent of risk sharing. None of these papers empirically study demand for both indemnity and index insurance in the same setting.

The rest of the paper is organised as follows. Section 2 presents our theoretical framework and key predictions. Section 3 describes the experimental design. Section 4 lays out the estimation strategy. Section 5 presents the results, and Section 6 concludes.

2. THEORY

The key predictions of our theoretical framework are that index insurance is a complement to risk-sharing arrangements, whilst indemnity insurance acts as a substitute. We show this through three related models. The first two models assume insurance is actuarially fair, and serve simply to highlight the key intuition: the more agents engage in risk sharing, the more they need to be insured against common shocks, covered by index insurance, rather than individual-specific shocks, covered by indemnity insurance. The first 'skeletal' model shows this in the simplest possible context, and the second model shows this in the context of standard insurance model with binary loss states. The third model, which we see as our primary theoretical contribution, generates the key predictions in the standard insurance model with binary loss states and marked-up premia.

2.1 Skeletal model

The central intuition can be illustrated using a 'skeletal' model as follows. An infinite number of risk-averse farmers maximise expected utility. Indexed by i , they have stochastic incomes y_i . Expected income, denoted x , is the same for all farmers and itself random, reflecting a common shock. An index is able to track x perfectly and thus underpin an index insurance product against shocks to x .

The farmers implement a costless risk-sharing mechanism, whereby a proportion θ of each farmer's realised income is paid into a common kitty which is then evenly divided between them. Since the number of farmers is infinite, each receives an amount equal to the expected contribution, $E(\theta y_i) = \theta x$, from the kitty. So after risk sharing, farmer i 's income is

$$\theta x + (1 - \theta)y_i.$$

Farmers also have access to indemnity and index insurance, which compensate for variation in x and y_i , respectively.

Then if both insurance types are actuarially fairly priced, each risk-averse farmer will fully insure, as follows:

Result 1. In the skeletal model, each farmer purchases θ units of index insurance and $1 - \theta$ units of indemnity insurance.

That is, the proportion of index insurance in the optimal insurance portfolio is equal to the proportion of the farmer's realised income that is paid into the risk-sharing pool. *In other words, the greater the level of risk sharing in the community, the greater the demand for index insurance and the lower the demand for indemnity insurance.*

Intuitively, indemnity insurance and risk sharing are substitutes because both serve the purpose of smoothing consumption within the group, that is, protecting against idiosyncratic shocks. And index insurance and risk sharing are complements because the greater the extent of risk sharing, the more the residual risk relates to the aggregate shock, which is covered by index insurance.

Note that this result does not require conditions on the utility function beyond the standard assumptions of non-satiation and risk aversion. Farmers need not have the same utility function. Furthermore, the result does not depend on the distributions of x and y_i . In particular, it does not depend on whether farmers' incomes are correlated conditional on x .

2.2 Binary loss and actuarially fair insurance

The central intuition can be embedded in a more standard insurance framework with a binary loss state. This section studies the case of actuarially fair insurance and continues to assume only that the farmers are risk-averse. The next section studies the case with marked-up insurance premia for the CRRA family of utility functions.

As above, consider an infinite number of risk-averse farmers maximising expected utility. Each farmer has endowment y and faces a possible loss $L < y$. With probability q , a shock occurs that is common to all farmers, henceforth referred to as a drought. If there is a drought, each farmer incurs the loss with probability P . If there is no drought, the loss is incurred with probability $p < P$. The expression $P - p$ can be interpreted as a measure of how predictive drought occurrence is of actual losses.

For simplicity we assume that farmer losses are independent conditional on drought outcome.⁵ In other words, all factors affecting the probability of loss, aside from the common shock, are assumed to be idiosyncratic and uncorrelated. However, unconditionally

⁵The direction of our predictions would not change if we allowed for conditionally correlated losses. It would mean that the level of risk sharing of idiosyncratic losses that can be achieved is lower and thus the demand for indemnity insurance will be higher.

losses are positively correlated, since each farmer is more likely to incur the loss if there is a drought.

This loss and index structure is more in line with Clarke and Kalani (2011), Dercon et al. (2014), and Clarke (2016) than with Mobarak and Rosenzweig (2012) or Boucher and Delpierre (2014). Unlike the latter two papers, in our model there is either a loss or no loss for each participant, and the aggregate state does not affect wealth directly, but modifies the probability of loss.

However, unlike Mobarak and Rosenzweig (2012) and Dercon et al. (2014), but in common with Boucher and Delpierre (2014), we assume that the index tracks expected loss perfectly. This simplifying assumption allows us to focus on the trade-off between risk sharing and the demand for insurance. Continuing improvements in index insurance design also makes this assumption increasingly realistic.⁶

As in the skeletal framework above, the farmers engage in costless risk sharing: irrespective of individual outcomes, they all pay a proportion θ of their assets into a common kitty which is divided equally between them. Hence each farmer pays either θy or $\theta(y-L)$, depending on individual outcome, into the kitty, and receives $\theta(y-PL)$ from the kitty if there is a drought and $\theta(y-pL)$ if there is not.

There are thus four possible outcomes for each farmer, as shown in Table 1.

Each farmer has access to two forms of insurance: indemnity insurance and index insurance. Let α be the number of units of indemnity insurance purchased by a farmer. Each unit of indemnity insurance pays out L if and only if the farmer suffers an individual loss, irrespective of whether or not there is a drought. And let β be the number of units of index insurance taken out. Each unit of index insurance pays out L if and only if there is a drought, irrespective of whether the farmer suffers a loss.

Since both types of insurance are actuarially fair, premia equal expected losses. Hence the cost per unit of cover is $(qP + (1-q)p)L$ for indemnity insurance and qL for index insurance.

With insurance, the outcome for a farmer in each state of the world is as shown in Table 2. (The probabilities are as in Table 1 and omitted to save space.) The second row in the table, where the index insurance pays out but there is no loss, represents the upside basis risk associated with index insurance. The third row in the table, where the farmer incurs a loss but the index insurance does not pay out, represents the downside basis risk associated with index insurance.

The farmers' problem is to maximise expected utility over these four outcomes with respect to insurance decisions α and β . At the optimum, the risk-averse farmers choose α and β to ensure full consumption smoothing, such that in each of the four states of the

⁶Innovations in the design of index insurance products have led to indices that approximate expected income, such as average area yield or satellite indices based on Normalised Difference Vegetation Indices (NDVI).

world each farmer gets their autarchic expected income $y - (qP + (1 - q)p)L$.

Simple algebraic rearrangement shows that this full consumption smoothing can be achieved by farmers purchasing indemnity and index insurance as follows.

Result 2. In the binary-loss insurance framework, optimal indemnity and index insurance purchases are given by:

$$\begin{aligned}\alpha &= 1 - \theta && \text{and} \\ \beta &= \theta(P - p)\end{aligned}$$

Since full consumption smoothing is achieved costlessly, this is an optimal solution for all utility functions satisfying non-satiation and risk aversion.

The expressions for optimal insurance demand show how indemnity insurance and index insurance are, respectively, a substitute and a complement to risk sharing. Furthermore, the expression for β reproduces a key result from Clarke (2016): demand for index insurance is decreasing in basis risk. This is because as $P - p$ decreases, the index is less informative of losses. Indeed, when $P = 1$ and $p = 0$, so that $P - p = 1$, the index is perfectly informative of individual losses, and index insurance becomes equivalent to indemnity insurance. As the expressions for α and β show, in this case one unit of effective indemnity insurance is bought.

Three key assumptions of this setup are worth noting. First, both types of insurance are assumed to be actuarially fair. In practice, while governments often subsidise insurance, privately provided insurance is generally sold with a mark-up. This assumption is relaxed in the next section where we consider the case of marked-up premia.

Second, our theoretical model assumes an infinite risk-sharing pool, which clearly does not obtain in practice. However, real risk-sharing groups are often comprised of relatively large networks of family and friends (Fafchamps and Lund, 2003), so an infinite risk-sharing pool may in some settings be more realistic than, or at least an alternative benchmark to, the two-person risk pools of, for example, Mobarak and Rosenzweig (2012) and Dercon et al. (2014). Furthermore, since these papers, along with that of Boucher and Delpierre (2014), generate similar results without the assumption of an infinite risk-sharing group, it is clear that the assumption of an infinite pool is not necessary to obtain similar theoretical results. Rather, the assumption serves to simplify the model and thereby highlight its intuition.

Third, we assume that the extent of risk sharing is fixed exogenously. This permits the isolation of the effect of risk sharing on insurance demand and is consistent with the experimental design. Our approach is complementary to that of (Takahashi, Barrett and Ikegami, 2018), who vary the price of insurance and study its effect on measures

of risk sharing. The assumption of exogenous risk sharing is also directly relevant to settings where participation in risk-sharing networks is socially enforced, and where local risk-sharing institutions are perhaps based on tradition and resistant to change.

2.3 Binary loss and marked-up premia

The sections above assumed that both types of insurance were sold with actuarially fair premia. This may be realistic in cases where insurance is subsidised.⁷ However, in practice insurance is often not fully subsidised. Furthermore, in theory the presence of fully subsidised indemnity insurance generates the question of why farmers wouldn't simply purchase full indemnity insurance, rather than engage in a combination of risk sharing, index and indemnity insurance.

This section therefore examines the case of marked-up insurance premia. The model is otherwise exactly as above. However, only the constant relative risk aversion (CRRA) family of utility functions is considered.

Assume that indemnity insurance premia are marked up by a factor $M > 1$ above the actuarially fair rate, and that index insurance premia are marked up by a factor $N > 1$. It will be assumed that the mark-up on indemnity insurance is greater than that on index insurance, $M > N$. This assumption is based on the key conceptual advantages of index over indemnity insurance, including lower verification costs and reduced information asymmetry, and is also consistent with observed premia and with the literature. Clarke and Kalani (2011) use loadings of 60% for indemnity insurance and 20% for index insurance. Reported commercial loadings range from 70% to 430% for rainfall index insurance (Cole et al., 2013, Table 1) and from 140% to 470% for indemnity insurance (Hazell, 1992, Table 1).

The following additional quantities are defined for convenience:

$$r = qP + (1 - q)p$$

$$s = Mr$$

$$t = Nq$$

$$w = tP + (1 - t)p$$

These are interpreted as follows: r is the unconditional probability of loss occurrence, s is the marked-up premium per unit of indemnity cover, t is the marked-up premium per unit of index cover and w is defined analogously to r , but where the probability of

⁷For example, the Weather Based Crop Insurance Scheme in India covers 9.3 million farmers. For farmers who grow food crops, the cost to the farmers themselves is less than 2% of the commercial premium. In Mexico, the cost of the index-based drought insurance provided by CADENA to 2 million smallholder farmers is fully assumed by the state and federal governments (Cai, de Janvry and Sadoulet, 2016).

drought occurrence q has been multiplied by N .

We will assume $s \leq 1$. This is plausible because a value of $s > 1$ would imply that the cost of the indemnity insurance would exceed the insured loss, and no rational agent would buy it.

Consumption levels in the four states of the world are now as follows:

$$\begin{aligned} c_1 &= y - L + \theta(1 - P)L + \alpha(1 - s)L + \beta(1 - t)L \\ c_2 &= y - \theta PL - \alpha sL + \beta(1 - t)L \\ c_3 &= y - L + \theta(1 - p)L + \alpha(1 - s)L - \beta tL \\ c_4 &= y - \theta pL - \alpha sL - \beta tL \end{aligned}$$

The agent's problem is to set the insurance levels α and β so as to maximise expected utility over these four states:

$$\max_{\alpha, \beta} \quad qPu(c_1) + q(1 - P)u(c_2) + (1 - q)pu(c_3) + (1 - q)(1 - p)u(c_4)$$

Result 3. When the instantaneous utility function satisfies constant relative risk aversion, the optimal levels of insurance cover α and β vary with the extent of risk sharing as follows:

$$\begin{aligned} \frac{d\alpha}{d\theta} &= -\frac{k + 1 - \alpha}{k + \theta} \\ \frac{d\beta}{d\theta} &= \frac{(P - p)k + \beta}{k + \theta} \end{aligned}$$

Here,

$$k = \frac{y/L - s}{s - w}$$

does not depend on θ , α or β .

Proof. Introducing the short-hand notation $u'_i \equiv u'(c_i)$, the first-order conditions of the problem are given by:

$$\begin{aligned} \frac{\partial EU}{\partial \alpha} &= (1 - s)qPu'_1 - sq(1 - P)u'_2 + (1 - s)(1 - q)pu'_3 - s(1 - q)(1 - p)u'_4 = 0 \\ \frac{\partial EU}{\partial \beta} &= (1 - t)qPu'_1 + (1 - t)q(1 - P)u'_2 - t(1 - q)pu'_3 - t(1 - q)(1 - p)u'_4 = 0 \end{aligned}$$

Consider an initial value for θ . The optimal insurance levels α and β for this level of risk sharing are implicitly given by the first-order conditions. Dividing these conditions

by the marginal utility in one of the four states of the world, say u'_4 , we obtain:

$$(1-s)qP\frac{u'_1}{u'_4} - sq(1-P)\frac{u'_2}{u'_4} + (1-s)(1-q)p\frac{u'_3}{u'_4} - s(1-q)(1-p) = 0$$

$$(1-t)qP\frac{u'_1}{u'_4} + (1-t)q(1-P)\frac{u'_2}{u'_4} - t(1-q)p\frac{u'_3}{u'_4} - t(1-q)(1-p) = 0$$

Since s, t, q, P and p are constants that do not depend on θ, α or β , it is clear that, as θ is varied, the first-order conditions will remain satisfied as long as u'_1, \dots, u'_4 stay in fixed proportion to each other.

For CRRA utility with parameter ρ , we have

$$\frac{u'_i}{u'_j} = \left(\frac{c_i}{c_j}\right)^{-\rho},$$

so this implies keeping the consumption levels $c_1 \dots c_4$ in fixed proportion to each other as θ is varied. Say,

$$\frac{c_1}{c_4} = K_1, \quad \frac{c_2}{c_4} = K_2 \quad \text{and} \quad \frac{c_3}{c_4} = K_3,$$

for constants K_1, K_2 and K_3 .

Taking the natural logarithm and differentiating, we find:

$$\frac{d \ln c_1}{d\theta} = \frac{d \ln c_2}{d\theta} = \frac{d \ln c_3}{d\theta} = \frac{d \ln c_4}{d\theta}$$

Substituting in for the consumption levels and differentiating with respect to θ gives:

$$\frac{1-P + (1-s)\frac{d\alpha}{d\theta} + (1-t)\frac{d\beta}{d\theta}}{y/L - 1 + (1-P)\theta + (1-s)\alpha + (1-t)\beta} = \frac{-P - s\frac{d\alpha}{d\theta} + (1-t)\frac{d\beta}{d\theta}}{y/L - P\theta - s\alpha + (1-t)\beta}$$

$$= \frac{1-p + (1-s)\frac{d\alpha}{d\theta} - t\frac{d\beta}{d\theta}}{y/L - 1 + (1-p)\theta + (1-s)\alpha - t\beta} = \frac{-p - s\frac{d\alpha}{d\theta} - t\frac{d\beta}{d\theta}}{y/L - p\theta - s\alpha - t\beta}$$

This system contains only two linearly independent equations. Solving them for $\frac{d\alpha}{d\theta}$ and $\frac{d\beta}{d\theta}$ gives the expressions above. \square

Corollary 1. When indemnity insurance has a higher loading than index insurance, that is when $M > N$, demand for indemnity insurance decreases in risk sharing and demand for index insurance increases in risk sharing.

Proof. When $M > N$,

$$s - w = M[qP + (1-q)p] - [NqP + (1-Nq)p]$$

$$= (M - N)r + (N - 1)p$$

$$> 0.$$

Since $L < y$ and $s < 1$, the ratio k in Result 3 is positive. Then, as $\alpha \leq 1, \beta \geq 0, \theta \geq 0$

and $P > p$, it follows that $\frac{d\alpha}{d\theta} < 0$ and $\frac{d\beta}{d\theta} > 0$. \square

Corollary 2. The derivatives $\frac{d\alpha}{d\theta}$ and $\frac{d\beta}{d\theta}$ do not vary with θ , that is, demand for indemnity insurance and demand for index insurance are both linear in risk sharing.

Proof. Differentiating the equations in Result 3 with respect to θ gives $\frac{d^2\alpha}{d\theta^2} = \frac{d^2\beta}{d\theta^2} = 0$. \square

Corollary 3. The quantity

$$\frac{1 - \theta - \alpha}{(P - p)\theta - \beta}$$

remains constant as θ varies. It can be interpreted as the ratio of the gap in indemnity insurance cover (relative to the no-loading, full insurance case) to the gap in index insurance cover.

Proof. From the proof of Result 3, we know that the ratios of the different consumption levels remain constant as θ varies. In particular,

$$\frac{c_2}{c_4} = \frac{c_4 + \beta L - (P - p)\theta L}{c_4} = 1 - \frac{(P - p)\theta L - \beta L}{c_4}$$

and

$$\frac{c_3}{c_4} = \frac{c_4 - L + \alpha L + \theta L}{c_4} = 1 - \frac{L - \theta L - \alpha L}{c_4}$$

are both constant. But then

$$\frac{L - \theta L - \alpha L}{(P - p)\theta L - \beta L} = \frac{1 - \theta - \alpha}{(P - p)\theta - \beta}$$

is also constant. \square

Figure 1 illustrates optimal demand for insurance as a function of risk sharing when the utility function satisfies constant relative risk aversion (CRRA) with parameter $\rho \rightarrow 1$ (logarithmic utility), $\rho = 2$ and $\rho = 3$. The figure was generated by numerically optimising expected utility over α and β with all other parameter values as in the experiment.⁸ It illustrates that the slopes of the demand curves, negative for indemnity insurance and positive for index insurance, do not depend on θ or ρ . It also shows that optimal α can be negative for sufficiently high θ , implying that the agent would want to *sell* indemnity insurance.

2.4 Summary of predictions

All three theory sections generate the following two key predictions:

1. Demand for indemnity insurance decreases with the level of risk sharing.
2. Demand for index insurance increases with the level of risk sharing.

⁸That is, $y = 200$, $L = 100$, $q = 1/6$, $P = 5/6$, $p = 1/6$, $M = 1.44$ and $N = 1.20$.

3. CONTEXT AND EXPERIMENTAL DESIGN

3.1 Context

The theoretical predictions were tested in an artefactual field experiment with 400 farmers from 10 villages in one district of Tigray, Ethiopia. These farmers were members of 14 different *iddir*. An *iddir* is an association of, typically, 50–200 individuals who are connected by ties of family, friendship, geographical area, occupation or ethnic group (Mauri, 1967). The objectives of an *iddir* are to provide mutual aid and financial assistance in case of emergencies. Every *iddir* has written statutes, by-laws and records of contributions and payouts. The rules define membership procedures, payout schedules, contributions and also a set of fines and other measures for non-payment of contributions (Aredo, 1993; Dercon et al., 2006; LeMay-Boucher, 2009). *Iddir* differ in the extent to which insurance contributions depend on the occurrence of a shock. One type only collects contributions when a shock occurs and payments are made directly to the member who incurs the loss. A second type collects a small monthly payment in addition to contributions when shocks occur. A third type collects only regular monthly payments from their members (insurance premia), saved in a communal *iddir* fund. Insurance claims by members are then paid out of this fund.

Iddir, especially those defined by a particular rural area, are stable institutions. In our sample, subjects report having been a member of their *iddir* for on average 10 years, with a maximum of 30 years. Eighty-four per cent of subjects report that they contribute only in case of shocks, 10% make monthly contributions, and 6% do both. In our sample, the average monthly contribution is 2 birr.⁹ Average contributions in case of shocks are 13 birr, with a standard deviation of 22 birr. The *iddir* also institute committees that visit the agricultural land of their members, and advice on agricultural production strategies.

Because the experiment is concerned with the demand for both indemnity and index insurance, care was taken to select villages without prior exposure to either form of insurance. However, a number of index insurance products with a similar design to the product used in our experiment are offered to farmers in other parts of Ethiopia. Nyala Insurance and Oromia Insurance Company offer weather-based index insurance products to smallholder farmers. These products were taken up by 12,000 smallholder farmers in 2015 and 45,000 farmers in 2016. In 2016, Oromia paid out 2.6 million birr (USD 140,000) in claims. The contracts offer a fixed payout if cumulative monthly precipitation recorded at the nearest National Meteorological Agency station falls below a threshold that is calculated using data on past rainfall and yields. For 2017, 200,000–300,000 farmers have purchased index insurance products from these two companies. Payouts are made when a Normalised Difference Vegetation Index (NDVI), based on satellite data, falls below a

⁹The exchange rate at the time of the experiment was 1 USD to 60 birr.

defined threshold.

3.2 Recruitment and selection

The 400 farmers in our sample were randomly selected from ten randomly selected villages (kebele) in Hitalo Wajira district in Tigray, Ethiopia. Villages were selected subject to a maximum travel time from Mekele city of three hours and a minimum distance between the villages of ten kilometres. In advance of the study, permission was obtained from local administrative officials called Development Agents (DAs). The DAs provided lists of farmers for each village. The 40 farmers per village that participated in the study were randomly drawn from this list. The DA informed the farmers that they were invited to participate in a survey and an experiment about weather, crops and agricultural decision-making. They were told that they could earn 100–200 birr and that the session would take no more than three hours. Farmers were also told that their answers to the survey and decisions in the experiment would be recorded by enumerators, but that responses would be anonymised. They were asked to come to the Farmer Training Centre in their village on the specified day. They would first take part in the experiment and then be fielded the survey. One hundred per cent of invited farmers attended. Subjects were informed that they would be given 200 birr in an envelope at the start of the session, and that during the experiment, they would be exposed to the risk of losing 100 birr. They were also told that they would be able to purchase insurance to mitigate the risk. The payoffs were substantial in the context of rural Ethiopia, where the daily wage for unskilled labour was in the range 50–150 birr at the time.

3.3 The experiment

The experiments took place in the Farmer Training Centres. As soon as farmers entered, they were seated in private cubicles constructed from desks, chairs and cardboard dividers. The ten sessions were randomly allocated to four treatment arms, as shown in Table 3. Participants were provided with comprehensive training on the insurance game as a group and through individual instruction from the enumerators. The same eight enumerators were present at all the sessions, each enumerator in charge of training and collecting the responses for five subjects. Before the actual experiment, there were un-incentivised practice runs to aid and assess understanding. Since many of the subjects were illiterate, the experiment was explained orally with the help of visual aids, and physical randomisation devices (coloured tokens and dice). After playing the games, subjects were fielded a survey to collect background information.

3.3.1 Aggregate and idiosyncratic shocks

Subjects were introduced to the concepts of aggregate and idiosyncratic shocks. The aggregate shock in the game was framed as a possible region-wide drought, the occurrence of which depended on the colour of a token drawn from an envelope. There were 6 tokens in the envelope: 1 yellow and 5 blue. If the yellow token was drawn, the drought occurred, and if a blue token was drawn, it did not. The aggregate shock was drawn at the front of the room so that all participants could see.

Each participant then rolled a die to determine whether or not he or she would incur the loss. The type of die rolled by the individual participants depended on whether or not the drought occurred: If the drought occurred, each participant would roll a die with 5 yellow sides and 1 green side. If there was no drought, each participant would roll a die with 1 yellow side and 5 green sides. If a subject's die showed yellow, he or she suffered a loss of 100 birr, whereas a green die throw resulted in no loss.

Thus, as in the theory section, the aggregate shock (drought or not) was informative of the probability of suffering individual losses. Each participant's die throw was independent, and there was no group-level loss. But since the choice of die depended on the aggregate shock, farmer losses were unconditionally correlated. The probabilities and states of the game, before considering risk sharing and insurance, are presented in Figure 2. The unconditional expected net payoff in autarchy was $200 - 2 \cdot (5/36) \cdot 100 \approx 172$ birr.

3.4 Risk sharing

As in our model, the extent of risk sharing was not left to the participants but varied experimentally. In the treatment arms without risk sharing, participants could not share risk with each other during the session.¹⁰

In the risk-sharing treatment arms, participants were randomly and anonymously paired with another participant in the same session. The identity of the partner was not revealed to participants. They were informed that, irrespective of the aggregate shock, insurance decisions and payouts, they would be sharing losses with their partner such that if one suffered the 100 birr loss but the other did not, then 50 birr would be transferred from the latter to the former. This corresponds to $\theta = 1$ in the model.

In line with the model, losses were shared but insurance premia and payouts were not. Thus, in the absence of any insurance purchase and payouts, risk sharing participants' wealth would be the same, although differing decisions over insurance purchase could lead

¹⁰It is still possible that they shared risk after the session, even though subjects did not know with whom they were teamed up. But it is not clear that someone who had done well in the experiments would choose to reveal this. Even if they did, they might have difficulty identifying a deserving recipient, since it would be in everybody's interest to understate their earnings.

to different final outcomes. This setup ensures that participants do not have an incentive to free-ride on their partner’s insurance, which might otherwise depress demand for insurance and confound the mechanism studied here (De Janvry, Dequiedt and Sadoulet, 2014; Janssens and Kramer, 2016). This is arguably realistic in an agricultural context where yield outcomes are observable to all, but insurance contracts are probably not.

3.5 Insurance

All subjects received information about indemnity and index insurance. They were told that indemnity insurance would pay out in the case of a loss, irrespective of whether or not the drought occurred, and that index insurance would pay out in the case of drought, irrespective of whether or not they incurred a loss.

Participants could purchase 0, 1 or 2 units of insurance. Each unit of insurance was associated with a payout of 50 birr, that is, half the potential loss. Hence, 2 units would provide full insurance. The premium per unit of cover was 20 birr for indemnity insurance and 10 birr for index insurance. Table 4 presents unconditional probabilities, expected payouts and premia for the two insurance types. In line with the literature and commercial insurance pricing, indemnity insurance was priced with a higher loading (mark-up above the actuarially fair rate) than index insurance.¹¹ Lower verification costs and reduced potential for moral hazard are two main reasons why index insurance can be expected to be sustainable at a lower mark-up than indemnity insurance. Given their different loadings, we do not attempt to compare the *levels* of demand for index and indemnity insurance to comment on which is more attractive, but rather we examine how demand for each *varies* when risk sharing is introduced.

Table 5 shows net payoffs for an individual participant for each state of the world and each possible insurance decision, in the absence of risk sharing. Net payoffs are given by the endowment (200), minus any insurance premia, minus the loss (100) if incurred, plus any insurance payout. From the table it is clear that participants in the index insurance treatments were exposed to downside basis risk with a probability of 5/36.

The subjects’ understanding of the game was assessed by conducting practice rounds where subjects were asked, based on their insurance decisions and the realisations of the aggregate shock and idiosyncratic losses, if they would receive a payout (yes or no) and how much the insurance claim payment would be (0, 50, 100 birr).

Table 6 summarises the structure and timing of the game.

¹¹Clarke and Kalani (2011) use loadings of 60% for indemnity insurance and 20% for index insurance. Reported commercial loadings range from 70% to 430% for rainfall index insurance (Cole et al., 2013, Table 1) and from 140% to 470% for indemnity insurance (Hazell, 1992, Table 1).

4. ESTIMATION STRATEGY

We use three econometric specifications to estimate the effect of risk sharing on insurance demand. Because the outcome variable is discrete with three possible outcomes (each participant buys 0, 1 or 2 units of insurance), we begin with an ordered probit. Next we run probit regressions where the outcome variable is a binary indicator for buying any (1 or 2 units) insurance versus none (0 units). Finally, we estimate a double probit hurdle model, where a first probit is used to model the decision regarding whether to buy insurance or not (the hurdle), and a second probit models the decision regarding whether to buy 1 or 2 units of insurance, conditional on buying at least one.

The key difference between the hurdle model and the ordered probit is that the hurdle model allows the outcome variable to be thought of as the result of two distinct choices by the participant, whereas in the ordered probit the decision is best thought of as a single choice. We use both models because it is uncertain how the decision processes of the subjects are best represented. However, all three specifications yield qualitatively similar results.

It is helpful to define a latent variable Y_i^* that is linear in the independent variables:

$$Y_i^* = X_i\beta + \epsilon_i, \quad i = 1, \dots, N$$

The vector of independent variables X_i includes the main treatment variable (a binary indicator for risk sharing), a constant term and, depending on the specification, fixed effects and control variables. The error terms ϵ_i are standard normally distributed.

In the probit regressions, the binary dependent variable Y_i is 0 if the subject buys 0 units of insurance and 1 if the subject buys 1 or 2 units of insurance. The observed binary outcome Y_i is modelled as follows:

$$Y_i = \begin{cases} 0 & \text{if } Y_i^* \leq 0 \\ 1 & \text{if } Y_i^* > 0 \end{cases}, \quad i = 1, \dots, N$$

For the ordered probit regressions, the latent variable Y_i^* is defined as above. But now the observed outcome variable (the number of units of insurance purchased) Y_i takes on values 0 through 2 according to the scheme

$$Y_i = \begin{cases} 0 & \text{if } Y_i^* \leq \mu_1 \\ 1 & \text{if } \mu_1 < Y_i^* \leq \mu_2 \\ 2 & \text{if } Y_i^* > \mu_2 \end{cases},$$

where μ_1 and μ_2 are threshold parameters to be estimated along with the coefficients β .

The probability of each outcome is:

$$\begin{aligned}\Pr(Y_i = 0) &= \Phi(\mu_1 - X_i\beta) \\ \Pr(Y_i = 1) &= \Phi(\mu_2 - X_i\beta) - \Phi(\mu_1 - X_i\beta) \\ \Pr(Y_i = 2) &= 1 - \Phi(\mu_2 - X_i\beta)\end{aligned}$$

Here, Φ is the cumulative distribution function for the standard normal. The likelihood for the ordered probit can be written

$$L = \prod_{i=1}^N \sum_{n=0}^2 \mathbb{1}\{Y_i = n\} \Pr(Y_i = n),$$

where $\mathbb{1}\{\}$ is the indicator function.

To allow for the possibility that the data have been generated through a process whereby subjects separately decide to purchase zero or a positive amount of insurance, and then decide the number of units to purchase (Mullahy, 1986; Cameron and Trivedi, 1998; Botelho et al., 2009), we estimate the parameters of a double probit hurdle model (Cribari-Neto and Zeileis, 2009). The first probit, with outcome Y_i' , determines whether the individual buys at least 1 unit of insurance. This is the ‘hurdle’. The second probit, with outcome Y_i'' , determines whether the individual buys 1 or 2 units of insurance, conditional on buying at least 1.

Thus, in the hurdle model, the number of units of insurance purchased is given by

$$Y_i = Y_i'(1 + Y_i'').$$

Denoting by β' and β'' the parameter vectors for the first and second probit, respectively, the likelihood function for the hurdle model is

$$\begin{aligned}L &= \prod_{i \in \Omega_0} \Phi(-\beta' X_i) \cdot \prod_{i \in \Omega_1 \cup \Omega_2} \Phi(\beta' X_i) \\ &\cdot \prod_{i \in \Omega_1} \Phi(-\beta'' X_i) \cdot \prod_{i \in \Omega_2} \Phi(\beta'' X_i),\end{aligned}$$

where $\Omega_0 = \{i | Y_i = 0\}$ is the set of non-purchasers and $\Omega_1 = \{i | Y_i = 1\}$, $\Omega_2 = \{i | Y_i = 2\}$ are the subjects who choose 1 and 2 units of insurance, respectively.

5. RESULTS

Of the 400 subjects who started the session, three were unable to complete due to unanticipated family or work engagements. These three were excused and received their participation fee. Each of the four questions testing understanding in the practice rounds

were answered correctly by more than 80% of subjects. 72% of the subjects answered all four questions correctly.

Table 7 presents summary statistics for each of the four treatment arms: indemnity insurance without (column 1) and with (column 2) risk sharing, and index insurance without (column 4) and with (column 5) risk sharing. As these arms technically form two separate experiments, one for each type of insurance, we test for balance separately. Column 3 presents t -tests for equality of means across the indemnity insurance treatments, and column 6 does the same for index insurance. Out of 20 balance tests, equality is rejected at the 5% level for only one (marital status for index insurance). This rejection rate is close to what would be expected statistically (5%), but we still present regressions that control for all these observable characteristics.

The last row of Table 7 presents the main outcome variable, the number of units of insurance purchased. The mean comparisons foreshadow our main results: risk sharing decreases the mean number of units of indemnity insurance purchased by 27% (from 1.73 to 1.26) and increases the mean number of units of index insurance by 130% (from 0.60 to 1.38), and the differences are highly statistically significant.

As well as repeating the means, Table 8 shows the proportion of individuals purchasing 0, 1 or 2 units of insurance by treatment. The proportion of individuals purchasing at least one unit of indemnity insurance *decreases* from 0.99 to 0.86 when risk sharing is introduced, while the proportion of individuals purchasing at least one unit of index insurance *increases* (doubles) from 0.46 to 0.93.

Regression analysis confirms these findings. Table 9 shows ordered probit regression results for the number of units of insurance purchased (0, 1 or 2). Columns 2, 3, 5, and 6 include enumerator fixed effects. Column 3 and 6 also control for sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in Tropical Livestock Units and land owned in hectares. In all regressions, the risk-sharing coefficient is significant at the 1% level. For indemnity insurance, the overall effect is negative, as predicted by the theory. The effect is driven by an increase in the probability that subjects purchase 0 or 1 unit of insurance, and a decrease in the probability that subjects purchase 2 units of insurance. For index insurance, the overall effect is positive, implying that risk sharing increases demand for index insurance, as predicted. The effect is driven by a reduction in the probability that individuals purchase 0 units of insurance and an increase in the probability that subjects purchase 2 units of insurance.

Table 10 presents probit regressions where the dependent variable is an indicator for whether the participant bought a positive amount (1 or 2 units) of insurance. In line with the predictions, risk sharing decreases the probability of buying indemnity insurance (by approximately 20%), while it increases the probability of buying index insurance (by approximately 35%). The risk-sharing coefficient is everywhere significant at the 1% level.

Table 11 shows maximum-likelihood estimates for the hurdle model. In column 1, it is shown that the probability of buying any indemnity insurance, and also the probability of buying two units of insurance conditional on buying at least one, both decrease significantly with risk sharing. This is robust to the inclusion of enumerator fixed effects and control variables in columns 2 and 3. Furthermore, and with reference to the lower half of the table, risk sharing increases the predicted probability of buying 0 or 1 unit of indemnity insurance, but decreases the probability of buying 2 units. For indemnity insurance, the coefficients on risk sharing are everywhere significant at the 1% level.

For index insurance (columns 4–6), risk sharing is associated with a positive and significant (at the 1% level) increase in the propensity to buy at least one unit of insurance. The effect on the propensity to buy two units of insurance, conditional on buying at least one, is also positive but not significant at the 5% level. Risk sharing decreases the predicted probability of buying no insurance, and increases the predicted probabilities of buying both 1 and 2 units.

Thus the hurdle model confirms the ordered probit results, but suggests that, for index insurance, the main effect of risk sharing is to increase the probability of buying some insurance rather than to increase the number of units purchased from 1 to 2. For indemnity insurance, it appears that risk sharing decreases both the propensity to buy some insurance rather than no insurance, and also the propensity to buy 2 units conditional on buying at least 1.

Table 12 presents a robustness check where the ten sessions are left out of the regressions one by one. All specifications are ordered probit regressions and include enumerator fixed effects and control variables. In the upper panel, it is shown that the negative and significant effect of risk sharing on the demand for indemnity insurance is robust to dropping any of the six indemnity-insurance sessions from the analysis. The risk-sharing coefficient is everywhere significant at the 1% level. In the lower panel, it is shown that the positive and significant effect of risk sharing on the demand for index insurance is robust to dropping three out of four sessions (significant at the 1% level). When dropping session 8, one of the index insurance sessions without risk sharing, the coefficient on risk sharing remains positive but is no longer statistically significant.

Given that there were only six indemnity insurance and four index insurance sessions, and the consequent loss of power associated with dropping one of them, we believe that the robustness check confirms that our main results are not driven by any one outlying session.

6. CONCLUSION

We present a simple and parsimonious theoretical framework, showing that risk sharing is a substitute for indemnity insurance but a complement to index insurance. To test

these predictions we present the first experimental evidence on demand for indemnity and index insurance where the extent of risk sharing is varied exogenously. In an artefactual field experiment with low-income farmers in Ethiopia, the theoretical predictions are confirmed. Our results show that indemnity insurance is a substitute for risk sharing, while index insurance complements it. This highlights the importance of considering existing risk-sharing arrangements when introducing formal insurance contracts and, in particular, encouraging risk sharing when introducing index insurance.

References

- Ahlin, Christian.** 2016. “Group lending, matching patterns, and the mystery of micro-credit: Evidence from Thailand.” Unpublished manuscript.
- Albarran, Pedro, and Orazio Attanasio.** 2003. “Limited commitment and crowding out of private transfers: Evidence from a randomised experiment.” *Economic Journal*, 113(486): C77–C85.
- Aredo, Dejene.** 1993. “The iddir: A study of an indigenous informal financial institution in Ethiopia.” *Savings and Development*, 17(1): 77–90.
- Arnott, Richard, and Joseph E Stiglitz.** 1991. “Moral hazard and nonmarket institutions: Dysfunctional crowding out or peer monitoring?” *American Economic Review*, 81(1): 179–190.
- Attanasio, Orazio, and José-Víctor Ríos-Rull.** 2000. “Consumption smoothing in island economies: Can public insurance reduce welfare?” *European Economic Review*, 44(7): 1225–1258.
- Attanasio, Orazio, and Nicola Pavoni.** 2011. “Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption.” *Econometrica*, 79(4): 1027–1068.
- Berg, Erlend.** 2018. “Funeral insurance: An inter-generational commitment device?” *Journal of African Economies*, 27(3): 321–346.
- Besley, Timothy, Stephen Coate, and Glenn Loury.** 1993. “The economics of rotating savings and credit associations.” *American Economic Review*, 83(4): 792–810.
- Botelho, Anabela, Glenn W Harrison, Lígia M Costa Pinto, and Elisabet E Rutström.** 2009. “Testing static game theory with dynamic experiments: A case study of public goods.” *Games and Economic Behavior*, 67(1): 253–265.
- Boucher, Steve, and Matthieu Delpierre.** 2014. “The impact of index-based insurance on informal risk-sharing arrangement.” Luxembourg Institute of Socio-Economic Research (LISER) Working Paper 13.
- Cai, Jing, Alain de Janvry, and Elisabeth Sadoulet.** 2016. “Subsidy policies and insurance demand.” NBER Working Paper 22702.
- Cameron, Colin, and Pravin K Trivedi.** 1998. *Regression analysis of count data*. Cambridge University Press.

- Casaburi, Lorenzo, and Jack Willis.** 2017. “Time vs state in insurance: Experimental evidence from contract farming in Kenya.” Unpublished manuscript.
- Clarke, Daniel J.** 2016. “A theory of rational demand for index insurance.” *American Economic Journal: Microeconomics*, 8(1): 283–306.
- Clarke, Daniel J, and Gautam Kalani.** 2011. “Microinsurance decisions: Evidence from Ethiopia.” In *Insurance Design for Developing Countries*. DPhil thesis, Daniel Clarke, University of Oxford.
- Clarke, Daniel J, and Stefan Dercon.** 2016. *Dull disasters? How planning ahead will make a difference*. Oxford University Press.
- Coate, Stephen, and Martin Ravallion.** 1993. “Reciprocity without commitment: Characterization and performance of informal insurance arrangements.” *Journal of Development Economics*, 40(1): 1–24.
- Cochrane, John H.** 1991. “A simple test of consumption insurance.” *Journal of Political Economy*, 99(5): 957–976.
- Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery.** 2013. “Barriers to household risk management: Evidence from India.” *American Economic Journal: Applied Economics*, 5(1): 104–135.
- Cribari-Neto, Francisco, and Achim Zeileis.** 2009. “Beta regression in R.” Manuscript.
- Cutler, David M, and Jonathan Gruber.** 1996. “The effect of Medicaid expansions on public insurance, private insurance, and redistribution.” *American Economic Review*, 86(2): 378–383.
- De Janvry, Alain, Vianney Dequiedt, and Elisabeth Sadoulet.** 2014. “The demand for insurance against common shocks.” *Journal of Development Economics*, 106: 227–238.
- De Quidt, Jonathan, Thiemo Fetzer, and Maitreesh Ghatak.** 2016. “Group lending without joint liability.” *Journal of Development Economics*, 121: 217–236.
- Dercon, Stefan, Joachim De Weerd, Tessa Bold, and Alula Pankhurst.** 2006. “Group-based funeral insurance in Ethiopia and Tanzania.” *World Development*, 34(4): 685–703.
- Dercon, Stefan, Ruth Vargas Hill, Daniel Clarke, Ingo Outes-Leon, and Alemayehu Seyoum Taffesse.** 2014. “Offering rainfall insurance to informal insurance groups: Evidence from a field experiment in Ethiopia.” *Journal of Development Economics*, 106: 132–143.
- Elabed, Ghada, Marc F Bellemare, Michael R Carter, and Catherine Guirking.** 2013. “Managing basis risk with multiscale index insurance.” *Agricultural Economics*, 44(4-5): 419–431.

- Fafchamps, Marcel, and Susan Lund.** 2003. "Risk-sharing networks in rural Philippines." *Journal of Development Economics*, 71(2): 261–287.
- Ghatak, Maitreesh.** 1999. "Group lending, local information and peer selection." *Journal of Development Economics*, 60(1): 27–50.
- Hazell, Peter B R.** 1992. "The appropriate role of agricultural insurance in developing countries." *Journal of International Development*, 4(6): 567–581.
- Janssens, Wendy, and Berber Kramer.** 2016. "The social dilemma of microinsurance: Free-riding in a framed field experiment." *Journal of Economic Behavior & Organization*, 131(B): 47–61.
- Kahneman, Daniel, and Amos Tversky.** 1979. "Prospect theory: An analysis of decision under risk." *Econometrica*, 47(2): 263–292.
- LeMay-Boucher, Philippe.** 2009. "Beninese and Ethiopian informal insurance groups: A comparative analysis." *Development Policy Review*, 27(3): 333–347.
- Mace, Barbara.** 1991. "Consumption volatility: Borrowing constraints or full insurance." *Journal of Political Economy*, 99(5): 928–956.
- Mauri, Arnaldo.** 1967. *Il mercato del credito in Etiopia*. A. Giuffrè. 20.
- McIntosh, Craig, Felix Povel, and Elisabeth Sadoulet.** 2015. "Utility, risk, and demand for incomplete insurance: Lab experiments with Guatemalan cooperatives." Unpublished manuscript.
- Mobarak, Ahmed Mushfiq, and Mark R Rosenzweig.** 2012. "Selling formal insurance to the informally insured." Yale Economics Department Working Paper 97.
- Mullahy, John.** 1986. "Specification and testing of some modified count data models." *Journal of Econometrics*, 33(3): 341–365.
- Samphantharak, Krislert, and Robert M Townsend.** 2018. "Risk and return in village economies." *American Economic Journal: Microeconomics*, 10(1): 1–40.
- Segal, Uzi.** 1988. "Probabilistic insurance and anticipated utility." *Journal of Risk and Insurance*, 55(2): 287–297.
- Takahashi, Kazushi, Christopher B Barrett, and Munenobu Ikegami.** 2018. "Does index insurance crowd in or crowd out informal risk sharing? Evidence from rural Ethiopia." *American Journal of Agricultural Economics*, forthcoming.
- Townsend, Robert M.** 1994. "Risk and insurance in village India." *Econometrica*, 62(3): 539–591.
- Wakker, Peter, Richard Thaler, and Amos Tversky.** 1997. "Probabilistic insurance." *Journal of Risk and Uncertainty*, 15(1): 7–28.

A. EXPERIMENTAL INSTRUCTIONS

[Instructions for enumerators and session leaders in italic font.]
[Instructions for participants in regular font.]

There are four types of sessions: Indemnity insurance without risk sharing, indemnity insurance with risk sharing, index insurance without risk sharing, and index insurance with risk sharing. Each participant takes part in only one session, and all forty participants in a session are members of the same iddir. Since most subjects are illiterate, each decision problem is explained with visual aids. All sessions start with a group-level introduction, followed by a number of un-incentivised practice games and then the main, incentivised, experiment. Throughout it is important to explain to the participants that the main experiment will affect their earnings, while the practice games do not. At the end of the session, a survey is fielded to each participant.

First a simple “drought and loss” practice game is played, without risk sharing or insurance. The aim is to familiarise participants with the basic risk structure of the experiment. Then, the insurance concepts are introduced. This is followed by practice games for indemnity insurance, index insurance, and risk sharing. Finally, the main experiment is conducted. Care is taken to explain that the final experiment involves real monetary payoffs.

A.1 Introduction

Tables are numbered 1–40 and the assigned number of the participant is preceded by the session number, for example, for session 4 and participant 11: 4.11. In all sessions the session leader reads the following to the participants:

Thank you for participating. In this experiment we will ask you to make decisions about insurance against possible monetary losses that will be determined by drawing tokens and rolling dice. You will receive an initial endowment of 200 birr at the start of the experiment. In the event that you experience a loss, the loss amount will be taken from your initial endowment.

You will have the option of buying insurance. It is up to you to decide whether to buy insurance or not. There is no right or wrong answer. We are interested in learning about what you prefer.

The session will take no more than three hours. At the start you will be given 200 birr, and you may experience a loss of 100 birr. You will be offered different insurance products against the payment of an insurance premium. The insurance premium will be deducted from your endowment. Your net pay will depend on your decisions as well as the token drawn and the die rolled. This means that you should think carefully about the decisions you make.

On the table in front of you there is a white envelope containing 200 birr and two dice, one blue and one yellow. Please count the money and check that the dice are there. You will now receive instructions but later on, when we play the experiment, you may lose some of the money.

We kindly request that you do not discuss with the other participants, and that you make your decisions in private. One of the enumerators will come

to you to record your decisions. If you have a question about the experiment, you can ask your enumerator.

Before we play the main experiment, we will explain everything carefully and play some practice rounds. These practice rounds will not involve real money.

A.2 The drought and loss game

In all sessions the session leader reads the following to the participants:

I have six tokens in front of me; five are blue and one is yellow. Let's imagine that the blue tokens represent good weather — no drought — and the yellow token represents drought. I will put these tokens in a bag and draw one. If I draw the yellow token this means that there is a drought. If I draw a blue token this means that there is no drought. The outcome of this draw is common to all participants in the room.

The session leader now plays two rounds.

Even though the draw of the weather outcome through the tokens is the same for all of you, the quality of your yield may differ. In drought years, some farmers may still have a good crop, and in no-drought years, some farmers may still have a bad crop. However, in no-drought years you are more likely to have a good crop, while in drought years you are more likely to have a bad crop.

In front of you are two dice, one yellow and one blue, corresponding to the colour of the tokens. The yellow die has one green side (good crop) and five yellow sides (bad crop). If a yellow token (drought) is drawn, all participants will throw the yellow die. In drought years it is more likely that your crop will be bad (five out of six) but there is still a chance that your crop is good (one out of six).

The blue die has five green sides (good crop) and one yellow side (bad crop). If a blue token (no drought) is drawn, everybody will throw this die. In no-drought years, it is more likely that your crop will be good (five out of six), but there is still a chance that your crop is bad (one out of six).

Now you are all asked to throw the die that corresponds to the token drawn. This will determine if you experience a loss (bad crop, yellow side) or not (good crop, green side).

The participants play two practice rounds. In these practice rounds, the session leader draws a blue token once and a yellow token once, and the participants throw the corresponding die. The enumerators check that the correct die is used, and explain what the net payout to the participant would be based on the endowment and die throw.

A.3 Introducing indemnity insurance

In all sessions, the session leader reads the following to the participants:

I will now introduce you to indemnity insurance. It pays out if you incur a loss, irrespective of whether there is a drought. That is, payout is tied to your individual die throw rather than the group token. The overall probability of incurring a loss in the drought game is $5/18$. You will incur a loss of 100 birr if the die throw results in a yellow side. You will not incur a loss if your die throw results in a green side.

The overall probability of throwing yellow is $5/18$. A fair premium that covers your whole loss is then 28 birr ($5/18 \times 100$). There is a 44% loading that represents the costs for the administration of the insurance.

You will have the option of buying indemnity insurance against the loss. You can buy 0, 1, or 2 units of insurance. Each unit of insurance costs 20 birr and pays out 50 birr in case of a loss. So buying 2 units of insurance costs 40 birr and pays out 100 birr in case of a loss.

Two practice rounds of the game with insurance are played. The enumerator records the outcomes and choices of the participants.

Practice round 1: no drought

1. *Does the participant want to take out insurance? Record 0, 1 or 2 units, and let the participants pay the correct premium.*
2. *A blue session-level token is selected, implying that there is no drought.*
3. *Each participant throws a die to determine their individual outcome. The enumerator verifies that the correct die (the one with five green sides and one yellow side) is used and records the outcome, yellow or green. Let the participant pay 100 birr in case of loss (yellow die throw).*
4. *Ask the participant: Do you think you will get an insurance payout? Record yes or no.*
5. *Ask the participant: How much do you think you will get in insurance payout? Record 0 birr, 50 birr or 100 birr.*
6. *Pay out the correct amount (50 or 100 birr) if participant has insurance and the die outcome is yellow.*
7. *Remind the participant that this was a practice round and reverse payments.*

Practice round 2: drought. Repeat the steps above, but select the yellow token and ensure participants throw the die with five yellow sides and one green side.

A.4 Introducing index insurance

In all sessions, the session leader reads the following to the participants:

A second type of insurance is index insurance. It pays out if there is a drought, irrespective of whether you incur a loss. That is, the payout is tied to the group token draw rather than your individual die throw. The probability of a drought occurring in the drought game is $1/6$. An actuarially fair premium is therefore 17 birr. With 20% loading this implies a premium of 20 birr. Index

insurance can be cheaper than indemnity insurance, but there is a risk that you will not receive a payout (because there is no drought) even though you incur the loss.

You will incur a loss of 100 birr if your die throw results in a yellow side. You will not incur a loss if your die throw results in a green side. The insurance pays out if and only if there is a drought, irrespective of whether you incur a loss. So it is the token from the bag, and not the die throw (crop loss), that determines payout. It is therefore possible (if the token is blue but your die throw is yellow) to experience the 100 birr loss and still not receive a payout.

You can buy 0, 1, or 2 units of insurance. Each unit of insurance costs 10 birr and pays out 50 birr in the case of a drought. So buying 2 units of insurance costs 20 birr and pays out 100 birr in the case of a drought.

Two practice rounds of the game with index insurance are played. The enumerator records the outcomes and choices of the participants.

Practice round 1: no drought

- 1. Does the participant want to take out insurance? Record 0, 1 or 2 units, and let the participants pay the correct premium.*
- 2. A blue session-level token is selected, implying that there is no drought.*
- 3. Each participant throws a die to determine their individual outcome. The enumerator verifies that the correct die (the one with five green sides and one yellow side) is used and records the outcome, yellow or green. Let the participant pay 100 birr in case of loss (yellow die throw).*
- 4. Ask the participant: Do you think you will get an insurance payout? Record yes or no.*
- 5. Ask the participant: How much do you think you will get in insurance payout? Record 0 birr, 50 birr or 100 birr.*
- 6. There is no insurance payout because the token draw is blue.*
- 7. Remind the participant that this was a practice round and reverse payments.*

Practice round 2: drought. Repeat the steps above, but select the yellow token and ensure participants throw the die with five yellow sides and one green side. Now all participants with insurance will receive a payout as the token is yellow.

A.5 Introducing risk sharing

In all sessions, the session leader reads the following to the participants:

In the next practice game you will still be able to buy insurance, but instead of bearing losses on your own, you will now share losses with another participant in this session. You will not learn who this other person is, nor will they learn who you are. Losses will be shared irrespective of the insurance decisions made by you and the other person. So, if one of you suffers the 100 birr loss and the other doesn't, 50 birr will be transferred to the person with the loss

from the person without the loss. If neither of you incurs a loss, or if both do, then no transfer is made.

Two practice rounds of the game with risk sharing are played, one for indemnity insurance and one for index insurance. The enumerator records the outcomes and choices of the participants.

Practice round 1: Indemnity insurance with risk sharing

- 1. Does the participant want to take out insurance? Record 0, 1 or 2 units, and let the participants pay the correct premium.*
- 2. A session-level token is drawn to determine whether or not a drought occurs. The outcome of this draw is common to all participants in a session and determines the die to be thrown by the individual participants. Record the colour of the token drawn, blue or yellow.*
- 3. Each participant throws a die to determine their individual outcome. The enumerator verifies that the correct die is used (depending on the token drawn) and records the outcome, yellow or green. Let the participant pay 100 birr in case of loss (yellow die throw).*
- 4. Ask the participant to wait while the enumerator verifies from the session leader whether or not the risk-sharing partner incurred a loss. Explain the partner's outcome to the participant and make the correct risk-sharing adjustment: pay 50 birr, receive 50 birr or no transfer.*
- 5. Ask the participant: Do you think you will get an insurance payout? Record yes or no.*
- 6. Ask the participant: How much do you think you will get in insurance payout? Record 0 birr, 50 birr or 100 birr.*
- 7. Give the insurance payout to the participant if the outcome of the die throw is yellow.*
- 8. Remind the participant that this was a practice round and reverse payments.*

Practice round 2: index insurance with risk sharing. Repeat the steps above, but charge the correct premium for index insurance and condition the insurance payout on the draw of the token, not on the die roll.

A.6 The main (incentivised) experiment and survey

The session leader reads to the participants:

So far we have been playing practice games in order to increase your understanding. The final game, played with real money, will now start. Whatever is left of the 200 birr after the token draw, die roll and transfers is yours to keep and does not have to be returned to the session leader or enumerator. After the experiment is finished, the enumerator will come to you and ask some questions about you and your household, farm and iddir. Please wait until the enumerator instructs you that everything is finished and that you may leave. We are very grateful for your participation.

Depending on the session, participants are told that the final game will involve indemnity or index insurance, and that it will be played with or without risk sharing. The enumerator records the outcomes and choices of the participants.

- 1. Does the participant want to take out insurance? Record 0, 1 or 2 units, and let the participants pay the correct premium.*
- 2. A session-level token is drawn to determine whether or not a drought occurs. The outcome of this draw is common to all participants in a session and determines the die to be thrown by the individual participants. Record the colour of the token drawn, blue or yellow.*
- 3. Each participant throws a die to determine their individual outcome. The enumerator verifies that the correct die is used (depending on the token drawn) and records the outcome, yellow or green. Let the participant pay 100 birr in case of loss (yellow die throw).*
- 4. In sessions with risk sharing only: Ask the participant to wait while the enumerator verifies from the session leader whether or not the risk-sharing partner incurred a loss. Explain the partner's outcome to the participant and make the correct risk-sharing adjustment: pay 50 birr, receive 50 birr or no transfer.*
- 5. Give the correct insurance payout to participant. In indemnity-insurance sessions, the payout is made if the outcome of the die throw is yellow. In index-insurance sessions, the payout is made if the token drawn is yellow. Either way, if a payout is made, the participant receives 50 birr for each unit of insurance purchased.*

Table 1: Outcomes after risk sharing, without insurance

Drought	Loss	Probability	Outcome after risk sharing
Yes	Yes	qP	$\theta(y - PL) + (1 - \theta)(y - L)$
Yes	No	$q(1 - P)$	$\theta(y - PL) + (1 - \theta)y$
No	Yes	$(1 - q)p$	$\theta(y - pL) + (1 - \theta)(y - L)$
No	No	$(1 - q)(1 - p)$	$\theta(y - pL) + (1 - \theta)y$

Table 2: Outcomes after risk sharing and insurance

Drought	Loss	Outcome after risk sharing and insurance
Yes	Yes	$\theta(y - PL) + (1 - \theta)(y - L) - \alpha(qP + (1 - q)p)L - \beta qL + \alpha L + \beta L$
Yes	No	$\theta(y - PL) + (1 - \theta)y - \alpha(qP + (1 - q)p)L - \beta qL + \beta L$
No	Yes	$\theta(y - pL) + (1 - \theta)(y - L) - \alpha(qP + (1 - q)p)L - \beta qL + \alpha L$
No	No	$\theta(y - pL) + (1 - \theta)y - \alpha(qP + (1 - q)p)L - \beta qL$

Table 3: Number of sessions by treatment arm

	Indemnity insurance	Index insurance
Without risk sharing	3	2
With risk sharing	3	2

Note: Allocation of sessions to the four treatment arms. In all, ten sessions were held, and each participant in a session received the same treatment.

Table 4: Characteristics of insurance products offered

	Indemnity insurance	Index insurance
Payout per unit of cover	50	50
Unconditional probability of payout	5/18	1/6
Expected payout (=actuarially fair premium)	$125/9 \approx 13.9$	$25/3 \approx 8.3$
Premium charged per unit of cover	20	10
Implied loading / mark-up above fair premium	44%	20%

Note: Each participant was offered either indemnity or index insurance, but not both, and could choose between 0, 1 or 2 units of cover. The premia were set so as to make the loadings roughly in line with the literature.

Table 5: States and payoffs

Prob	State	Units	Payoff with indemnity insurance		Payoff with index insurance	
			Calculation	Net	Calculation	Net
5/36	Drought, loss	0	$200 - 100$	100	$200 - 100$	100
		1	$200 - 20 - 100 + 50$	130	$200 - 10 - 100 + 50$	140
		2	$200 - 2 \cdot 20 - 100 + 2 \cdot 50$	160	$200 - 2 \cdot 10 - 100 + 2 \cdot 50$	180
1/36	Drought, no loss	0	200	200	200	200
		1	$200 - 20$	180	$200 - 10 + 50$	240
		2	$200 - 2 \cdot 20$	160	$200 - 2 \cdot 10 + 2 \cdot 50$	280
5/36	No drought, loss	0	$200 - 100$	100	$200 - 100$	100
		1	$200 - 20 - 100 + 50$	130	$200 - 10 - 100$	90
		2	$200 - 2 \cdot 20 - 100 + 2 \cdot 50$	160	$200 - 2 \cdot 10 - 100$	80
25/36	No drought, no loss	0	200	200	200	200
		1	$200 - 20$	180	$200 - 10$	190
		2	$200 - 2 \cdot 20$	160	$200 - 2 \cdot 10$	180

Note: Payoffs in birr for each combination of drought, loss, type of insurance and number units of insurance purchased. The (no drought, loss) state is where index insurance leads to downside basis risk. Risk-sharing transfers (for participants in risk-sharing sessions only) come in addition.

Table 6: Insurance game timing

-
1. Each participant is given a 200 birr endowment.
 2. Each participant buys insurance (0, 1 or 2 units) and pays the premium.
 3. The aggregate shock is drawn at the group (session) level.
 4. Each participant rolls a die to determine whether (s)he incurs a loss.
 5. Participants incurring a loss pay 100 birr to enumerator.
 6. *Participants are informed of their partner's loss outcome.*
 7. *Any loss is shared with partner: receive 50, send 50 or no adjustment.*
 8. Insurance pays out.
-

Note: The italicised items only apply to treatment arms with risk sharing.

Table 7: Summary statistics and balance checks

	(1)	(2)	(3)	(4)	(5)	(6)
	Indemnity insurance, without risk sharing	Indemnity insurance, with risk sharing	Difference t -test (2)-(1)	Index insurance, without risk sharing	Index insurance, with risk sharing	Difference t -test (5)-(4)
Female	0.37 (0.48)	0.39 (0.49)	0.025 (0.063)	0.47 (0.50)	0.42 (0.50)	-0.045 (0.080)
Age	48.0 (13.6)	44.7 (13.5)	-3.28* (1.75)	42.3 (13.9)	38.6 (10.1)	-3.68* (1.96)
Married	0.84 (0.37)	0.82 (0.39)	-0.016 (0.050)	0.80 (0.40)	0.62 (0.49)	-0.18** (0.074)
Household head	0.88 (0.32)	0.88 (0.32)	0.00100 (0.042)	0.91 (0.28)	0.97 (0.16)	0.062* (0.037)
Literate	0.55 (0.50)	0.59 (0.49)	0.037 (0.065)	0.40 (0.49)	0.50 (0.50)	0.10 (0.080)
Household size	5.41 (1.82)	5.57 (1.71)	0.16 (0.25)	5.35 (2.22)	4.97 (1.90)	-0.39 (0.37)
Tropical livestock units	1.86 (2.42)	1.80 (1.81)	-0.061 (0.28)	1.46 (1.60)	1.86 (1.66)	0.40 (0.26)
Land owned, in hectares	0.72 (0.39)	0.66 (0.38)	-0.063 (0.050)	0.64 (0.31)	0.64 (0.39)	0.0047 (0.056)
Observations	120	120	.	80	80	.

Note: Summary statistics and balance checks. Columns 1, 2, 4 and 5 present means of variables by treatment arm, with standard deviations in parentheses. Columns 3 and 6 present difference t -tests, with standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Insurance demand by treatment

	Indemnity insurance without risk sharing	Indemnity insurance with risk sharing	Index insurance without risk sharing	Index insurance with risk sharing
Proportion who purchase...				
no insurance (0 units)	0.01	0.14	0.54	0.08
partial insurance (1 unit)	0.26	0.46	0.32	0.48
full insurance (2 units)	0.73	0.40	0.14	0.45
some insurance (1 or 2 units)	0.99	0.86	0.46	0.93
Mean units purchased	1.73	1.26	0.60	1.38

Table 9: Ordered probit regressions of units of insurance purchased

	Units of indemnity insurance purchased			Units of index insurance purchased		
	(1)	(2)	(3)	(4)	(5)	(6)
Risk sharing	-0.93*** (0.16)	-0.92*** (0.16)	-1.10*** (0.18)	1.23*** (0.19)	1.28*** (0.19)	1.03*** (0.23)
Marginal change in probability of buying...						
zero units	0.12	0.11	0.13	-0.36	-0.35	-0.24
one unit	0.21	0.21	0.23	0.0084	0.010	0.0065
two units	-0.33	-0.32	-0.36	0.35	0.34	0.24
Enumerator fixed effects	No	Yes	Yes	No	Yes	Yes
Controls	No	No	Yes	No	No	Yes

Note: Ordered probit regressions. The dependent variable is the number of units of insurance purchased. Columns 3 and 6 include control variables for sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in Tropical Livestock Units and land owned in hectares. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Probit regressions

	Indemnity insurance purchased			Index insurance purchased		
	(1)	(2)	(3)	(4)	(5)	(6)
Risk sharing	-1.32*** (0.39)	-1.35*** (0.39)	-1.77*** (0.51)	1.53*** (0.25)	1.62*** (0.25)	1.81*** (0.38)
Marginal change in probability of buying any insurance	-0.16	-0.17	-0.23	0.41	0.42	0.32
Enumerator fixed effects	No	Yes	Yes	No	Yes	Yes
Controls	No	No	Yes	No	No	Yes

Note: Probit regressions. The dependent variable is a binary indicator for whether the participants purchased any insurance (1 or 2 units) or not. Columns 3 and 6 include control variables for sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in Tropical Livestock Units and land owned in hectares. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11: Double probit hurdle model

	Units of indemnity insurance purchased			Units of index insurance purchased		
	(1)	(2)	(3)	(4)	(5)	(6)
<hr/>						
Pr($Y \in \{1, 2\}$):						
Risk sharing	-1.32*** (0.39)	-1.35*** (0.39)	-1.77*** (0.51)	1.53*** (0.25)	1.62*** (0.25)	1.81*** (0.38)
Pr($Y = 2 Y \in \{1, 2\}$):						
Risk sharing	-0.73*** (0.18)	-0.74*** (0.18)	-0.92*** (0.21)	0.50* (0.26)	0.49* (0.28)	0.16 (0.36)
<hr/>						
Marginal change in probability of buying...						
zero units	0.13	0.13	0.16	-0.46	-0.46	-0.43
one unit	0.20	0.20	0.21	0.15	0.18	0.24
two units	-0.33	-0.33	-0.37	0.31	0.28	0.18
Enumerator fixed effects	No	Yes	Yes	No	Yes	Yes
Controls	No	No	Yes	No	No	Yes


Note: Double probit hurdle model estimates. The dependent variable is the number of units of insurance purchased. Columns 3 and 6 include control variables for sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in Tropical Livestock Units and land owned in hectares. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 12: Robustness check: Leaving out sessions one by one

	Units of indemnity insurance purchased						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Risk sharing	-1.10*** (0.18)	-0.74*** (0.20)	-0.84*** (0.21)	-1.34*** (0.22)	-1.30*** (0.23)	-1.23*** (0.20)	-1.62*** (0.25)
Enumerator fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Session dropped	None	#1	#2	#5	#6	#9	#10
	Units of index insurance purchased						
	(1)	(2)	(3)	(4)	(5)		
Risk sharing	1.03*** (0.23)	10.8*** (1.38)	1.31*** (0.27)	0.84*** (0.31)	0.24 (0.28)		
Enumerator fixed effects	Yes	Yes	Yes	Yes	Yes		
Controls	Yes	Yes	Yes	Yes	Yes		
Session dropped	None	#3	#4	#7	#8		


Note: Ordered probit regressions, leaving out one session at a time. The dependent variable is the number of units of insurance purchased. All regressions include enumerator fixed effects and control variables for sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in Tropical Livestock Units and land owned in hectares. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1: Optimal insurance demand as a function of risk sharing



Note: Numerically estimated optimal demand for indemnity (α) and index (β) insurance as a function of the risk-sharing parameter (θ). All parameters are as in the experiment, and the utility function satisfies constant relative risk aversion (CRRA) with parameter values $\rho \rightarrow 1$ (logarithmic utility), $\rho = 2$ and $\rho = 3$. The slopes of the demand curves, negative for α and positive for β , do not depend on θ or ρ . Note that optimal α can be negative, implying that agents would want to sell, rather than buy, indemnity insurance when θ is sufficiently high.

Figure 2: Aggregate shock and idiosyncratic losses



Note: Extensive-form autarchy game. Probabilities are presented next to the branches of the tree. The states are presented at the nodes of the tree representing drought, no drought, loss and no loss. The drought occurs with probability $1/6$. The loss occurs with probability $5/6$ if there is a drought and with probability $1/6$ if not.