

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

Erlend Berg, Michael Blake and Karlijn Morsink*

December 2021

Abstract

Households in developing countries commonly engage in risk sharing to cope with shocks. Despite this, the residual risk they remain exposed to — often due to aggregate events such as droughts and floods — is considerable. To mitigate these risks, governments, NGOs and multilateral organizations have introduced index insurance. To appreciate its welfare implications, however, it is necessary to assess how insurance interacts with pre-existing risk sharing. We ask to what extent the demand for index insurance — as compared to standard indemnity insurance — depends on the level of pre-existing risk sharing. We contribute by developing a simple theoretical framework which shows that, relative to a state of autarky, risk sharing between agents increases demand for index insurance and decreases demand for indemnity insurance. In an artefactual field experiment with Ethiopian farmers who share risk in real life, we test and confirm these predictions.

JEL: D14, D81, G22, O16

Key words: risk sharing, indemnity insurance, index insurance, artefactual field experiment

1. INTRODUCTION

Households in developing countries commonly share risk with family, friends and community members through informal transfers and emergency loans (Fafchamps and Lund, 2003). Despite this, they remain exposed to aggregate shocks such as droughts and floods (Townsend, 1994). The risk of aggregate shocks reduces the ability of individuals to insure each other. In recent years there has been a surge in initiatives introducing both indemnity and index insurance to mitigate these aggregate risks. This is especially so in

*Berg (corresponding author): University of Bristol, School of Economics, Priory Road Complex, Bristol BS8 1TU, United Kingdom, erlend.berg@bristol.ac.uk. Blake: Pricoa Private Capital. Morsink: Utrecht University.

We are indebted to Tagel Gebrehiwot Gidey for research coordination. We are grateful for useful comments from Daniel Clarke, Stefan Dercon, 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 Improving Institutions for Pro-Poor Growth (iiG), a research programme funded by the UK Department for International Development. Views expressed are not necessarily those of the funder. Declarations of interest: none.

Asia and Africa, where large investments have been made by countries, donor agencies and multilateral organizations (Clarke and Dercon, 2016).

In the context of providing protection against aggregate shocks, standard indemnity insurance, which covers personal losses resulting from a combination of idiosyncratic and aggregate risk, tends to be costly and fraught with asymmetric-information problems. Even though innovations in information and communication technology have recently reduced some of these concerns,¹ indemnity insurance has been increasingly supplemented by index insurance because the latter can reduce moral hazard and adverse selection and minimize costly claims verification.² Index insurance provides claim payments conditional on observable and publicly verifiable indices that are proxies for aggregate shocks, such as weather indices or indices of average yields.³ This implies that the index is positively correlated with the personal losses of individuals, but imperfectly so, as idiosyncratic components of losses are not covered. As a result of this feature of index insurance, the risk of suffering a personal loss is compounded with the risk of the index not matching the personal loss. This imperfect correlation gives rise to basis risk, an idiosyncratic component of losses that is not present in standard indemnity insurance. Basis risk implies that the insured person may pay an insurance premium, experience a personal loss, but not receive a claim payment because the index is not triggered.

The fact that index insurance provides protection against aggregate shocks – which previous literature suggests are poorly insured through pre-existing informal risk sharing arrangements (Townsend, 1994) – suggests that perhaps index insurance can complement protection provided by informal risk-sharing arrangements, which are best suited to provide protection against idiosyncratic risks. This would provide an interesting contrast to the way that indemnity insurance and informal risk-sharing interact - they both provide protection against the same risks, and may therefore substitute each other (Arnott and Stiglitz, 1991; Attanasio and Ríos-Rull, 2000). In this paper we therefore study how demand for index insurance, as well as indemnity insurance, depends on the extent of pre-existing risk sharing.

We make two main contributions. The theoretical contribution is to model how demand for both types of insurance depends on risk sharing in a single, simple, framework.

¹A range of technological innovations, predominantly through the use of smart phones, has made it possible to offer indemnity crop insurance while avoiding costly loss assessments by sending insurance company representatives to visit the plots of farmers. See for example Ceballos and Kramer (2019); Misra et al. (2020); Afshar et al. (2021).

²Moral hazard is reduced because there is little that the insured can do to affect the index outcome and hence the chance of a claim payment. Adverse selection is reduced because the contract does not differentiate between the correlation of the idiosyncratic risk and the index outcome.

³Index insurance is also referred to as parametric, or probabilistic, insurance (Kahneman and Tversky, 1979; Segal, 1988; Wakker et al., 1997; McIntosh et al., 2019). 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 hurricanes and earthquakes. In both cases, insurance is offered to stakeholders who typically already share risk through pre-existing arrangements.

We show that risk sharing *increases* demand for index insurance, while confirming the familiar result that it *decreases* demand for indemnity insurance. Intuitively, this is because the greater the extent of risk sharing, which insures the idiosyncratic component of losses, the more the residual risk of the agents relates to aggregate shocks, which are insured by index insurance. Risk sharing is therefore a substitute for indemnity insurance, while it acts as a complement to index insurance. This complementarity is strengthened by the fact that risk sharing provides an opportunity for individuals to insure basis risk, the idiosyncratic component of personal losses that is introduced by index insurance. Relative to the existing literature in this area, the simplicity of our model enables us to derive closed-form solutions for both indemnity insurance and index insurance in a single framework. It also allows us to study additional questions such as the impact of insurance pricing and preferences on the demand for both index and indemnity insurance.

Our second main contribution is to investigate the relationship between risk sharing and demand for index insurance — as well as indemnity insurance — in an artefactual field experiment. The participants are 400 Ethiopian farmers from 10 villages, who in their daily lives are exposed to aggregate weather shocks and share risk through village-level informal risk-sharing arrangements. In the experiment, farmers receive an endowment and are exposed to a shock, with aggregate and idiosyncratic components, which determine potential personal losses. At the level of villages, farmers were randomly assigned to four treatment arms: indemnity insurance with risk sharing, indemnity insurance without risk sharing, index insurance with risk sharing and index insurance without risk sharing. In the experiment, risk sharing is implemented by matching each farmer, anonymously and privately, with another farmer from their group and by fixing the nature and amount of sharing. Depending on their treatment assignment, before risks are realized but knowing whether they will or will not share risk, farmers get an opportunity to protect their personal income through either index insurance or through indemnity insurance.

As far as we know, this is the first study to control risk sharing exogenously and study its effect on demand for both types of insurance. Our approach avoids endogeneity concerns arising from the use of observational data or even experimental data with endogenously formed risk sharing groups, thus permitting identification of a causal relationship between risk sharing and insurance demand.

The experimental results bear out the predictions of the theory: risk sharing increases the number of units of index insurance purchased, while it decreases the number of units of indemnity insurance purchased.

Varying the level of risk sharing experimentally in order to test its impact on insurance demand is clearly attractive from an identification perspective. The 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 approach is complementary to those of Anderberg and Morsink (2020) and Lenel and Steiner (2020), who vary the availability of insurance and study its effect on voluntary and non-reciprocal transfers.

Varying the level of risk sharing to study its impact on insurance demand is also directly relevant to empirical settings where the level of risk sharing is relatively fixed because 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 et al., 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.

The participants in our artefactual field experiment are Ethiopian farmers who are members of *iddir*. Iddir are associations of individuals connected by ties of family, friendship, geographical area, occupation, or ethnic group (Mauri, 1967). The objective of an iddir is to provide mutual transfers and financial assistance in emergencies (Dercon et al., 2006; Pankhurst, 2008; Hoddinott et al., 2009; Aredo, 2010). The risk-sharing groups in our sample all have written rules for members, including regulations on the level of contributions and punishment for failure to comply. This context is especially suitable for the current study because iddir members are very familiar with informal risk-sharing transfers.

Our paper builds on a substantial literature on risk sharing and insurance (e.g., Townsend (1994); Arnott and Stiglitz (1991); Albarran and Attanasio (2003); De Janvry et al. (2014); Janssens and Kramer (2016)). We contribute to recent work on the complementarity or substitutability of risk sharing and insurance.⁴ Our theoretical set-up is similar in spirit to Dercon et al. (2014), and we generate similar predictions, but in a considerably simpler model. In the case of actuarially fair premia, our model is more general as we make no assumptions on the utility function beyond non-satiation and risk aversion, whereas Dercon et al. (2014) require additional assumptions on the third and fourth derivatives. Furthermore, 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) model the demand for index insurance (but not

⁴We are aware of a literature that specifically studies 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 of the risk-sharing groups in our sample form committees that visit the plots of its members to advise on agricultural 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).

indemnity insurance) in the presence of risk sharing. In a set-up 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 individual 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 represent 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).

In terms of empirical evidence it is known that indemnity insurance can crowd out risk sharing in both full and limited-commitment risk-sharing arrangements (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) study the effect of risk sharing on demand for index insurance by offering insurance to Indian subcastes that vary in the degree to which they share risk. While the insurance offers are random, the level of risk sharing in these groups is not, so that the impact of risk sharing on index insurance is potentially confounded with other subcaste-level time-invariant factors.⁵

Dercon et al. (2014) also study the impact of risk sharing on demand for index insurance empirically, by varying the framing of an insurance marketing intervention to Ethiopian farmers. In one treatment arm, the marketing intervention emphasises how risk sharing can be used as a mechanism to manage basis risk. They find that this intervention, relative to an intervention that does not focus on risk sharing, increases demand for index insurance. However, it is not evident that the marketing intervention affects actual risk sharing, and it remains possible that the intervention influences demand for index insurance through other mechanisms.

McIntosh et al. (2019) investigate willingness to pay for index insurance when farmers have the option to share risk with peers in a cooperative. They find no significant impact of the risk-sharing option on demand for index insurance. They attribute this in large part to enforcement problems in the ex-post risk-sharing arrangement of the cooperative: because individuals are not expecting substantial risk sharing to actually occur, demand for index insurance does not increase. This is an important finding, but it leaves open the question of whether individuals would purchase more insurance if risk sharing was expected. In our model and experiment, we exogenously fix the extent of risk sharing to abstract from enforcement problems.

Three recent articles complement the present paper by studying the opposite direction

⁵The authors also estimate 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 and not in others.

of causality. Using artefactual field experiments, Anderberg and Morsink (2020) and Lenel and Steiner (2020) vary the availability of insurance in Ethiopian and Cambodian settings, respectively, and study its impact on voluntary, non-reciprocal transfers. In a natural field experiment, Takahashi et al. (2019) investigate the extent to which instrumented variation in demand for index insurance — varied in an encouragement design where discount coupons are distributed randomly — affects transfers among members of pre-existing informal risk-sharing groups of pastoralists in Ethiopia. They find that greater demand for index insurance does not compromise pre-existing informal social arrangements, and they provide suggestive evidence that it may even strengthen them.

Our findings suggest that interventions that strengthen, build on, or encourage risk sharing of idiosyncratic losses can increase demand for index insurance and overall insurance coverage. Take-up rates for index insurance vary greatly, but the majority of studies find a relatively low take-up rate between 5–20% (Casaburi and Willis, 2018; Belissa et al., 2019; Ahmed et al., 2020). Take-up rates for weather-index products range from 0.5–50% (Cole et al., 2013; Karlan et al., 2014; Ahmed et al., 2020) with rates varying greatly, partly dependent on the mark-up over the actuarially fair premium. Take-up rates are often higher, between 12–70% for indices based on more sophisticated data than just weather data, such as those that track pasture or soil quality (Jensen et al., 2016; Casaburi and Willis, 2018; Belissa et al., 2019; Takahashi et al., 2019). Mobarak and Rosenzweig (2013), Hill et al. (2016) and Jensen et al. (2018) find that demand falls significantly with an increase in basis risk, suggesting that there is scope for the combination of index insurance with informal risk-sharing arrangements.

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

2. THEORY

The key predictions of the theoretical framework presented in this section 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 aggregate shocks, covered by index insurance, rather than idiosyncratic shocks, covered by indemnity insurance. The first, ‘skeletal’, model shows this in the simplest possible context, and the second model shows it in a standard insurance model with binary loss states. The third model, which we see as our primary theoretical contribution, generates the key predictions in a 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. Therefore, variability in x captures aggregate risk, while variability in y_i conditional on x captures idiosyncratic risk. In this way, personal realized income captures both aggregate risk (through variation in x) and idiosyncratic risk (through variation in y_i conditional on x).

We also assume the existence of an index that is able to track x , the aggregate risk, perfectly and thus underpin an index insurance product against shocks to x . Conceptually, this could be thought of as a satellite-based average crop-quality index or an area-yield index that captures the average income of the total area cultivated by the farmers, in which case the index would exactly equal x . However, the assumption would also hold approximately if there is a sufficiently positive correlation between the index and average income x , such as may be the case with weather-based indices.⁶

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 index and indemnity insurance, which compensate for variation in average income, x , and individual income, y_i , respectively.

In principle, risk sharing might happen on the basis of gross (agricultural) income, or net income, that is, income after deducting any insurance premia and adding any insurance payouts. However, our pilot, anecdotal evidence and experience with agricultural insurance programs indicate that risk sharing generally happens based on gross income, or yield, and we believe this is related to observability. Yield is generally observable through the quality of crops and the number of bags collected, while insurance decisions and outcomes are generally not observable, especially given that an increasing proportion of financial products purchases and claim payments happen through bank accounts

⁶Index-based insurance is thought to have become increasingly accurate with a move from reliance on scattered, ground-based rainfall gauges to estimates of crop quality based on satellite images, and as image frequency, resolution and analysis techniques have improved. Breustedt et al. (2008) compare the effectiveness of area yield index insurance and weather index insurance for wheat producers in Kazakhstan. Whilst we assume perfect correlation between the index and average income for tractability, the key theoretical insight does not depend on this: Dercon et al. (2014) find that risk sharing and index insurance are complements in a more general model that allows for imperfect correlation between index and average losses. However, they do not model indemnity insurance and their framework does not permit a closed-form solution.

or mobile money accounts. Incentives to hide financial circumstances when possible are well established in the literature (Jakiela and Ozier, 2016; Boltz et al., 2019). For these reasons, even though risk-sharing networks tend to be close-knit communities, our theory and experiment assume that farmers share agricultural risk, while insurance transactions are private.

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.

Proof. Risk-averse agents will insure fully when premia are actuarially fair. The optimal insurance purchases are therefore those that will provide the farmer with the unconditionally expected income, with certainty. Since index insurance covers variation in x while indemnity insurance covers variation in y_i , risk-sharing farmer i obtains full insurance by purchasing θ units of index insurance and $1 - \theta$ units of indemnity insurance. \square

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 or not 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 , an aggregate 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. This risk framework captures both aggregate (the drought either occurs or not) and idiosyncratic risk (variation in farmer outcomes conditional on drought outcome), without requiring these two types of shock to be additively separable.

For simplicity we assume that farmer personal losses are independent conditional on drought outcome.⁷ In other words, all factors affecting the probability of personal losses, aside from the aggregate shock, are assumed to be idiosyncratic and uncorrelated. However, unconditionally, losses are positively correlated, since each farmer is more likely to incur the loss if there is a drought.

This personal 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. As mentioned in the previous section, this would be the case if the index was an area-yield or satellite-based crop index based on the total area cultivated by the farmers. 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.⁸

As in the skeletal framework above, the farmers engage in costless risk sharing: irrespective of individual income, 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 income, 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 before considering insurance. If there is a drought and the farmer incurs a loss, which happens with probability qP ,

⁷The 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.

⁸Innovations 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).

the outcome after risk sharing is:

$$\theta(y - PL) + (1 - \theta)(y - L) \quad (1)$$

If there is a drought, but the farmer does not incur a loss, which happens with probability $q(1 - P)$, the outcome after risk sharing is:

$$\theta(y - PL) + (1 - \theta)y \quad (2)$$

With probability $(1 - q)p$ there is no drought but the farmer incurs the loss. In this case the outcome after risk sharing is:

$$\theta(y - pL) + (1 - \theta)(y - L) \quad (3)$$

Finally, if there is no drought and the farmer does not incur the loss, which happens with probability $(1 - q)(1 - p)$ the outcome after risk sharing is:

$$\theta(y - pL) + (1 - \theta)y \quad (4)$$

Next, we introduce insurance. Each farmer has access to two types: 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.

Now consider the four states of the world with insurance. If there is a drought and the farmer incurs a loss, which happens with probability qP , the outcome after risk sharing and insurance is:

$$\theta(y - PL) + (1 - \theta)(y - L) - \alpha(qP + (1 - q)p)L - \beta qL + \alpha L + \beta L \quad (5)$$

If there is a drought but the farmer incurs no loss, which happens with probability $q(1 - P)$, the outcome after risk sharing and insurance is:

$$\theta(y - PL) + (1 - \theta)y - \alpha(qP + (1 - q)p)L - \beta qL + \beta L \quad (6)$$

With probability $(1 - q)p$ there is no drought but the farmer does incur a loss. In this

case the outcome after risk sharing and insurance is:

$$\theta(y - pL) + (1 - \theta)(y - L) - \alpha(qP + (1 - q)p)L - \beta qL + \alpha L \quad (7)$$

Finally, if there is no drought and the farmer does not incur the loss, which happens with probability $(1 - q)(1 - p)$, the outcome after risk sharing and insurance is:

$$\theta(y - pL) + (1 - \theta)y - \alpha(qP + (1 - q)p)L - \beta qL \quad (8)$$

Equation 6, where the index insurance pays out but there is no loss, represents the upside basis risk associated with index insurance. Equation 7, 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 decision variables (insurance purchases) 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 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 with actuarially fair premia, 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 predictive of losses, so that $P - p$ is an inverse measure of basis risk. Indeed, when $P = 1$ and $p = 0$, so that $P - p = 1$, there is no basis risk. In this case, 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. 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 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.⁹ However, in practice insurance is often not fully subsidised.

This section therefore examines the case of marked-up insurance premia. The model is otherwise exactly as above. However, to keep the model tractable, we now focus our attention on the constant relative risk aversion (CRRA) family of utility functions.

We will 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.

⁹For 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 (Ritchie et al., 2016).

The following additional quantities are defined for convenience:

$$\begin{aligned} r &= qP + (1 - q)p \\ s &= Mr \\ t &= Nq \\ w &= tP + (1 - t)p \end{aligned}$$

The first three are interpreted as follows: r is the unconditional probability of loss occurrence, s is the marked-up premium per unit of indemnity cover and t is the marked-up premium per unit of index cover. The quantity w does not have an equally straightforward interpretation, but it simplifies the analysis. It is defined analogously to r , but where the probability of drought occurrence q has been ‘marked up’ by N . The quantity w reduces to r when index insurance is actuarially fair ($N = 1$).

We will also 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. No rational agent would buy insurance if this did not hold.

Consumption levels in the four states of the world depend on the agent’s decision variables, α and β , and can be written as follows:

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

Omitting the arguments of c_1, \dots, c_4 for parsimony, the agent’s problem is to set the insurance levels α and β so as to maximise expected utility over these four states:

$$EU(\alpha, \beta) = qPu(c_1) + q(1 - P)u(c_2) + (1 - q)pu(c_3) + (1 - q)(1 - p)u(c_4)$$

Whilst we do not obtain explicit solutions for optimal α and β , our primary interest is in the dependence of insurance decisions on the risk-sharing parameter θ . We can therefore consider optimal insurance levels α and β as implicit functions of θ , and derive the following result.

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:

$$\frac{d\alpha}{d\theta} = -\frac{k+1-\alpha}{k+\theta}$$

$$\frac{d\beta}{d\theta} = \frac{(P-p)k+\beta}{k+\theta}$$

Here,

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

does not depend on θ , α or β .

The proof for Result 3 can be found in Appendix A.

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 $M > N$,

$$\begin{aligned} s - w &= M[qP + (1 - q)p] - [NqP + (1 - Nq)p] \\ &= (M - N)r + (N - 1)p \\ &> 0. \end{aligned}$$

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 optimizing expected utility over α and β with all other parameter values as in the experiment.¹⁰ 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.

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. In these remote rural communities, most adults are members of the local *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. Iddir tend to have 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).

¹⁰That is, $y = 200$, $L = 100$, $q = 1/6$, $P = 5/6$, $p = 1/6$, $M = 1.44$ and $N = 1.20$.

Iddir differ in the extent to which insurance contributions depend on the occurrence of a shock. One type collects contributions only 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.¹¹ 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 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. In 2017, 200,000–300,000 farmers had index insurance cover from these two companies. Payouts are made when a Normalised Difference Vegetation Index (NDVI), based on satellite data, falls below a defined threshold.

3.2 Recruitment and selection

In September 2013, the 400 farmers in our sample were randomly selected from ten randomly selected villages (kebele), and therefore, by construction, from ten iddir, in Hitalo Wajira district in Tigray. 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 and were therefore also members of the same iddir.¹² The DA informed the farmers that they were invited to participate in a survey and an experiment about

¹¹The exchange rate at the time of the experiment was 1 USD to 60 birr.

¹²Among our participants, 98% report being a member of their village iddir.

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 of 50–150 birr at the time.

3.3 The experiment

This section lays out the experimental design. Full experimental instructions are provided in Appendix B.

In September and October 2013, one experimental session was held at the Farmer Training Centre of each of the ten sample villages. As soon as participants entered, they were seated in private cubicles constructed from desks, chairs and cardboard dividers. All individuals in a session were from the same village and iddir, and received the same experimental treatment.¹³ The four treatments were randomly assigned to the ten sessions, and thus to the ten villages. Three sessions were dedicated to the indemnity insurance treatment without risk sharing; two sessions to the index insurance treatment without risk sharing; three sessions to the indemnity insurance with risk sharing; and two sessions to the index insurance with risk sharing. A key advantage of this design is that all individuals in one session were accustomed to sharing risk with each other, making the experiment a natural extension of the risk-sharing behaviour that our respondents typically engage in with this group.

At the session level, participants were provided with comprehensive training on the insurance experiment, and enumerators assisted the participants during the experiment in case they had additional questions. The same eight enumerators were present at all sessions, each enumerator in charge of training and collecting the responses for five

¹³Individual-level treatment assignment was considered and piloted, but was deemed infeasible. Given that many respondents were illiterate, experimental instructions were provided orally and accompanied by group-level illustrations, demonstrations and trial runs, provided by the same researcher. Individual-level treatment assignment would have required either i) explanation and demonstration of all four treatments to all respondents, or ii) dividing farmers up into four separate groups and giving each group separate explanations in turn. The former would have required giving four complex and time-consuming explanations and raised concerns about mental load as well as disappointment over treatment assignment. The latter would also have been time-consuming, and would have made it more difficult to preserve anonymity in their assignment to a peer, and therefore the integrity of the experiment.

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 for all participants so that all participants could see and experienced the same aggregate shock.

Each participant then rolled a die to determine whether or not he or she would incur a personal 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 conditionally independent, 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 autarky was $200 - 2 \cdot (5/36) \cdot 100 \approx 172$ birr.

3.3.2 Risk sharing

As in the theory section, the extent of risk sharing was not left to individual participants but varied experimentally. In the treatment arms without risk sharing, participants could not share risk with each other during the session.¹⁴

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

¹⁴It 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.

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, and as discussed in Section 2.1, 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. This design is arguably realistic in an agricultural context where yield outcomes are observable to all, but insurance contracts are not. The setup also ensures that participants do not have an incentive to free-ride on their partner's insurance decision, which might otherwise depress demand for insurance and confound the mechanism studied here (De Janvry et al., 2014; Janssens and Kramer, 2016).

3.3.3 Insurance

Depending on their session, farmers received information about either indemnity insurance or index insurance. In the former case, they were told that insurance pays out in the case of an individual loss, irrespective of whether or not the drought occurs. In the latter case, they were informed that an insurance payout is made in case of a drought, irrespective of whether or not they incurred an individual 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 1 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 (Clarke and Kalani, 2011). 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 2 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 3 summarises the structure and timing of the game.

3.4 Descriptives

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 4 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 5) and with (column 6) 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 by risk sharing, and column 7 does the same for index insurance. Out of 16 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. As an additional test of balance we present normalized differences in columns 4 and 8. These differences also suggest that balance is achieved (Imbens and Rubin, 2015, Chapter 5), with only three normalised differences exceeding the suggested threshold of 0.25, and only one of these (again, marital status for index insurance) exceeding 0.3.

4. EXPERIMENTAL RESULTS

This section presents the results of the experiment. A simple comparison of means across treatment arms indicates that the theoretical predictions are borne out in the data, and regression confirms that these findings are statistically significant and robust to econometric specification, the inclusion of controls, and clustering the standard errors at the session level.

Our main results are discernible in a simple comparison of means across treatment arms. The first row of Table 5 presents, by treatment arm, the mean number of insurance units purchased. On average, the presence of risk sharing decreases the number of indemnity insurance units purchased from 1.73 to 1.26, while increasing the number of index insurance units purchased from 0.60 to 1.38. This is in line with our theoretical predictions.

The remaining rows of Table 5 detail the proportion of participants who buy 0 units (no cover), 1 unit (partial cover) or 2 units (full cover) of insurance, by experimental arm. The decline in mean demand for indemnity insurance is driven by a decline in the proportion of participants who buy full cover, while there is an increase in the prevalence

of both uninsured and partly covered participants. The increase in mean demand for index insurance is driven by a smaller proportion of participants buying no insurance, while the prevalence of both partial and full cover increases.

Table 6 presents OLS regressions of the number of units of insurance purchased (0, 1 or 2) on a binary variable indicating sessions with risk sharing. Columns 1 and 3 present the basic regression for indemnity insurance and index insurance, respectively, while our preferred specifications in columns 2 and 4 include controls for sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in tropical livestock units and land owned in hectares, as well as enumerator fixed effects. In all regressions, standard errors are clustered at the session level. Due to the small number of clusters we study the significance of the main risk-sharing coefficient using the wild bootstrap with Webb weights, and the associated p -value is presented in all columns. Risk sharing is associated with a reduction in indemnity insurance demand of 0.52 units ($p = 0.018$) and an increase in demand for index insurance of 0.52 units ($p = 0.057$) in the preferred specifications.

Table 7 presents probit regressions where the dependent variable is an indicator for whether the participant bought a positive amount (either 1 or 2 units) of insurance. The specifications are otherwise as above. Again columns 2 and 4 represent our preferred specification which includes controls, and p -values for the main coefficient, estimated with the wild bootstrap with Webb weights, are presented. The associated marginal changes in average insurance units purchased are also presented. In line with the predictions, risk sharing reduces the probability of buying indemnity insurance (by 18 percentage points, $p = 0.035$), while it raises the probability of buying index insurance (by 32 percentage points, $p = 0.064$).

We also implement a double probit hurdle model. Here, one probit is used to model the decision regarding whether to buy a positive amount of insurance or not (the hurdle), while a second probit models the decision regarding whether to buy 1 or 2 units of insurance, conditional on buying at least one. The advantage of the hurdle model over OLS is that it does not assume linearity in the insurance-purchase decision. It therefore allows us to study, in a regression format, whether the mean effects are driven by the extensive (buying insurance or not) or intensive (how many units to buy) margin.

Table 8 presents maximum-likelihood estimates for the hurdle model. In column 2, which is the preferred specification including controls, it is shown that the probability of buying any indemnity insurance, and also the probability of buying two units of indemnity insurance conditional on buying at least one, both decrease significantly with risk sharing. Furthermore, and with reference to the lower part of the table, risk sharing increases the predicted probability of buying 0 or 1 unit of indemnity insurance while decreasing the probability of buying 2 units.

The preferred specification for index insurance is presented in column 4. Risk sharing

is associated with an increase in the propensity to buy at least one unit of insurance ($p = 0.077$). The effect on the propensity to buy two units of insurance, conditional on buying at least one, is also positive but smaller and not significant. The lower panel suggests that risk sharing decreases the probability of buying no index insurance by 43 percentage points at the margin, and that it increases the probability of buying both 1 and 2 units.

Overall, the hurdle model indicates that, in the experiments, risk sharing reduces demand for indemnity insurance along both the extensive and the intensive margins. For index insurance, the positive effect of risk sharing is only significant along the extensive margin.

Table 9 presents an additional robustness check where the sessions are left out of the regression one by one. All specifications are double-probit hurdle models and include controls. Standard errors are clustered at the session level, and p -values for the main (risk-sharing) effect are estimated using the wild bootstrap with Webb weights to account for the small number of clusters. Dropping the indemnity insurance sessions one by one yields p -values of no more than 0.094. For index insurance, dropping sessions one by one yields p -values of no more than 0.13 for the extensive margin, while there is no evidence of an effect on the intensive margin.¹⁵ Dropping data in this way does, as expected, weaken the power of the tests, but overall the robustness check appears to confirm that our results are not driven by any one outlying session.

5. CONCLUSION

We present a parsimonious theoretical framework and show 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, who share risk in real life, the theoretical predictions are borne out. This highlights the importance of considering pre-existing risk-sharing arrangements when introducing formal insurance contracts and, in particular, the potential for risk-sharing networks to boost demand for index insurance. Intuitively, this is because risk sharing is best suited to cover idiosyncratic risk, while index insurance tends to cover aggregate shocks.

References

Afshar M.H., Foster T., Higginbottom T.P., Parkes B., Hufkens K., Mansabdar S., Ceballos F., Kramer B., 2021. Improving the Performance of Index Insurance Using Crop Models and Phenological Monitoring. *Remote Sensing* 13, 924.

¹⁵Note that the maximum-likelihood estimation did not converge when dropping session #3, one of the index-insurance sessions.

- Ahmed S., McIntosh C., Sarris A., 2020. The impact of commercial rainfall index insurance: experimental evidence from Ethiopia. *American Journal of Agricultural Economics* 102, 1154–1176.
- Albarran P., Attanasio O., 2003. Limited commitment and crowding out of private transfers: Evidence from a randomised experiment. *Economic Journal* 113, C77–C85.
- Anderberg D., Morsink K., 2020. The introduction of formal insurance and its effect on redistribution. *Journal of Economic Behavior and Organization* 179, 22–45.
- Aredo D., 1993. The iddir: A study of an indigenous informal financial institution in Ethiopia. *Savings and Development* 17, 77–90.
- Aredo D., 2010. The iddir: An informal insurance arrangement in Ethiopia. *Savings and Development* 34, 53–72.
- Arnott R., Stiglitz J.E., 1991. Moral hazard and nonmarket institutions: Dysfunctional crowding out or peer monitoring? *American Economic Review* 81, 179–190.
- Attanasio O., Ríos-Rull J.V., 2000. Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare? *European Economic Review* 44, 1225–1258.
- Belissa T., Bulte E., Cecchi F., Gangopadhyay S., Lensink R., 2019. Liquidity constraints, informal institutions, and the adoption of weather insurance: A randomized controlled Trial in Ethiopia. *Journal of Development Economics* 140, 269–278.
- Berg E., 2018. Funeral insurance: An inter-generational commitment device? *Journal of African Economies* 27, 321–346.
- Besley T., Coate S., Loury G., 1993. The economics of rotating savings and credit associations. *American Economic Review* 83, 792–810.
- Boltz M., Marazyan K., Villar P., 2019. Income hiding and informal redistribution: A lab-in-the-field experiment in Senegal. *Journal of Development Economics* 137, 78–92.
- Boucher S., Delpierre M., 2014. The impact of index-based insurance on informal risk-sharing arrangement. Luxembourg Institute of Socio-Economic Research (LISER) Working Paper 13.
- Breustedt G., Bokusheva R., Heidelbach O., 2008. Evaluating the potential of index insurance schemes to reduce crop yield risk in an arid region. *Journal of Agricultural Economics* 59, 312–328.
- Casaburi L., Willis J., 2018. Time versus state in insurance: Experimental evidence from contract farming in Kenya. *American Economic Review* 108, 3778–3813.
- Ceballos F., Kramer B., 2019. From index to indemnity insurance using digital technology: Demand for picture-based crop insurance. IFPRI Discussion Paper 1890 .
- Clarke D.J., 2016. A theory of rational demand for index insurance. *American Economic Journal: Microeconomics* 8, 283–306.
- Clarke D.J., Dercon S., 2016. *Dull disasters? How planning ahead will make a difference.* Oxford University Press.

- Clarke D.J., Kalani G., 2011. Microinsurance decisions: Evidence from Ethiopia. In: Insurance Design for Developing Countries. DPhil thesis, Daniel Clarke, University of Oxford.
- Cole S., Giné X., Tobacman J., Topalova P., Townsend R., Vickery J., 2013. Barriers to household risk management: Evidence from India. *American Economic Journal: Applied Economics* 5, 104–35.
- Cutler D.M., Gruber J., 1996. The effect of Medicaid expansions on public insurance, private insurance, and redistribution. *American Economic Review* 86, 378–383.
- De Janvry A., Dequiedt V., Sadoulet E., 2014. The demand for insurance against common shocks. *Journal of Development Economics* 106, 227–238.
- Dercon S., De Weerd J., Bold T., Pankhurst A., 2006. Group-based funeral insurance in Ethiopia and Tanzania. *World Development* 34, 685–703.
- Dercon S., Hill R.V., Clarke D., Outes-Leon I., Taffesse A.S., 2014. Offering rainfall insurance to informal insurance groups: Evidence from a field experiment in Ethiopia. *Journal of Development Economics* 106, 132–143.
- Fafchamps M., Lund S., 2003. Risk-sharing networks in rural Philippines. *Journal of Development Economics* 71, 261–287.
- Hill R.V., Robles M., Ceballos F., 2016. Demand for a simple weather insurance product in India: theory and evidence. *American Journal of Agricultural Economics* 98, 1250–1270.
- Hoddinott J., Dercon S., P., 2009. Networks and Informal Mutual Support in 15 Ethiopian Villages. *Institutional Economics Perspectives on African Agricultural Development* .
- Imbens G.W., Rubin D.B., 2015. Causal inference in statistics, social, and biomedical sciences. Cambridge University Press.
- Jakiela P., Ozier O., 2016. Does Africa need a rotten kin theorem? Experimental evidence from village economies. *Review of Economic Studies* 83, 231–268.
- Janssens W., Kramer B., 2016. The social dilemma of microinsurance: Free-riding in a framed field experiment. *Journal of Economic Behavior & Organization* 131, 47–61.
- Jensen N.D., Barrett C.B., Mude A.G., 2016. Index insurance quality and basis risk: Evidence from northern Kenya. *American Journal of Agricultural Economics* 98, 1450–1469.
- Jensen N.D., Mude A.G., Barrett C.B., 2018. How basis risk and spatiotemporal adverse selection influence demand for index insurance: Evidence from northern Kenya. *Food Policy* 74, 172–198.
- Kahneman D., Tversky A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–292.

- Karlan D., Osei R., Osei-Akoto I., Udry C., 2014. Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics* 129, 597–652.
- LeMay-Boucher P., 2009. Beninese and Ethiopian informal insurance groups: A comparative analysis. *Development Policy Review* 27, 333–347.
- Lenel F., Steiner S., 2020. Formal insurance and solidarity. Experimental evidence from Cambodia. *Journal of Economic Behavior and Organization* 174, 212–234.
- Mauri A., 1967. Il mercato del credito in Etiopia. A. Giuffrè.
- McIntosh C., Povel F., Sadoulet E., 2019. Utility, risk, and demand for incomplete insurance: Lab experiments with Guatemalan co-operatives. *The Economic Journal* 129, 2581–2607.
- Misra R., Ceballos F., Kannan S., Kramer B., 2020. Potential Impacts and Demand for Picture-Based Crop Insurance Qualitative research findings from Haryana State, India. Unpublished manuscript .
- Mobarak A.M., Rosenzweig M.R., 2012. Selling formal insurance to the informally insured. Yale Economics Department Working Paper 97.
- Mobarak A.M., Rosenzweig M.R., 2013. Informal risk sharing, index insurance, and risk taking in developing countries. *American Economic Review* 103, 375–80.
- Pankhurst A., 2008. The emergence, evolution and transformations of iddir funeral associations in urban Ethiopia. *Journal of Ethiopian Studies* 41, 143–185.
- Ritchie E.R., de Janvry A., Sadoulet E., 2016. Weather index insurance and shock coping: Evidence from Mexico’s CADENA program World Bank Policy Research Working Paper 7715.
- Segal U., 1988. Probabilistic insurance and anticipated utility. *Journal of Risk and Insurance* 55, 287–297.
- Takahashi K., Barrett C.B., Ikegami M., 2019. Does index insurance crowd in or crowd out informal risk sharing? Evidence from rural Ethiopia. *American Journal of Agricultural Economics* 101, 672–691.
- Townsend R.M., 1994. Risk and insurance in village india. *Econometrica* 62, 539–591.
- Wakker P., Thaler R., Tversky A., 1997. Probabilistic insurance. *Journal of Risk and Uncertainty* 15, 7–28.

A. PROOF RESULT 3

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:

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

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:

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

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

B. 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, treatments were randomized at the session level, and all forty participants in a session are members of the same iddir. In total there are 10 sessions, 3 sessions are assigned to indemnity insurance without risk sharing, 2 sessions are assigned to index insurance without risk sharing, 3 sessions are assigned to indemnity insurance with risk sharing, and 2 sessions are assigned to index insurance with risk sharing. Since most subjects are illiterate, each decision problem is explained with visual aids. All sessions start with a session-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 risk sharing, and either indemnity insurance or index insurance. Finally, the main experiment is conducted. Care is taken to explain that the final experiment involves real monetary payoffs.

B.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 insurance. If you buy insurance, 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 do not involve real money.

B.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.

B.3 Introducing indemnity insurance

In the indemnity insurance 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 ($5/18 \cdot 100$). 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.

B.4 Introducing index insurance

In the index insurance sessions, the session leader reads the following to the participants:

I will now introduce you to 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. 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.

B.5 Introducing risk sharing

In the risk sharing 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.

A practice round of the game with risk sharing is played, in the sessions with indemnity insurance for indemnity insurance, in the sessions with index insurance for index insurance. The enumerator records the outcomes and choices of the participants.

Practice round : 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 : Index 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 token draw is yellow.
8. Remind the participant that this was a practice round and reverse payments.

B.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: 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 2: 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 3: 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 steps only apply to treatment arms with risk sharing.

Table 4: Summary statistics and balance checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Indemnity insurance, without risk sharing	Indemnity insurance, with risk sharing	t -test (2) - (1)	Normalized difference (2) - (1)	Index insurance, without risk sharing	Index insurance, with risk sharing	t -test (6) - (5)	Normalized difference (6) - (5)
Female	0.37 (0.49)	0.39 (0.49)	0.025 (0.063)	0.052	0.47 (0.50)	0.42 (0.50)	-0.045 (0.080)	-0.091
Age	48.03 (13.58)	44.75 (13.49)	-3.28* (1.750)	-0.241	42.33 (13.91)	38.66 (10.11)	-3.68* (0.194)	-0.299
Married	0.84 (0.37)	0.82 (0.39)	-0.016 (0.050)	-0.042	0.80 (0.41)	0.62 (0.49)	-0.18** (0.074)	-0.386
Household head	0.88 (0.33)	0.88 (0.32)	0.001 (0.042)	0.003	0.91 (0.28)	0.97 (0.16)	0.062* (0.037)	0.266
Literate	0.55 (0.50)	0.59 (0.49)	0.037 (0.065)	0.074	0.40 (0.49)	0.50 (0.50)	0.10 (0.080)	0.200
Household size	5.41 (1.82)	5.57 (1.71)	0.161 (0.252)	0.091	5.36 (2.22)	4.97 (1.90)	-0.39 (0.37)	-0.187
Tropical livestock units	1.86 (2.42)	1.80 (1.81)	-0.061 (0.277)	-0.028	1.46 (1.60)	1.86 (1.66)	0.40 (0.26)	0.244
Land owned, in hectares	0.72 (0.39)	0.66 (0.38)	-0.063 (0.050)	-0.163	0.64 (0.31)	0.64 (0.39)	0.0047 (0.056)	0.013
Observations	120	120	.	.	80	80	.	.

Note: Summary statistics and balance checks. Columns 1, 2, 5 and 6 present means of variables by treatment arm, with standard deviations in parentheses. Columns 3 and 7 present the t -test statistic with its standard error in parentheses. Columns 4 and 8 present normalized differences. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Insurance purchases by experimental arm

	Indemnity insurance without risk sharing	Indemnity insurance with risk sharing	Index insurance without risk sharing	Index insurance with risk sharing
Mean units purchased	1.73	1.26	0.60	1.38
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

Table 6: OLS regressions

	Units of indemnity insurance purchased		Units of index insurance purchased	
	(1)	(2)	(3)	(4)
Risk sharing	-0.47 (0.16)	-0.52 (0.17)	0.77 (0.47)	0.52 (0.30)
Wild cluster bootstrap p -value	0.013	0.018	0.20	0.057
Controls	No	Yes	No	Yes

Note: OLS regressions. The dependent variable is the number of units of insurance purchased (0, 1 or 2). Standard errors, clustered at the session level, are in parentheses. Wild cluster bootstrapped p -values for the risk-sharing coefficient, estimated with Webb weights due to the small number of clusters, are presented. Columns 2 and 4 control for participant sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in tropical livestock units and land owned in hectares, as well as enumerator fixed effects.

Table 7: Probit regressions

	Indemnity insurance purchased		Index insurance purchased	
	(1)	(2)	(3)	(4)
Risk sharing	-1.32 (0.38)	-1.77 (0.35)	1.53 (0.91)	1.81 (0.76)
Wild cluster bootstrap p -value	0.048	0.035	0.20	0.064
Marginal change in the probability of buying any insurance	-0.16	-0.18	0.41	0.32
Controls	No	Yes	No	Yes

Note: Probit regressions. The dependent variable is a binary indicator for whether the participants purchased any insurance (either 1 or 2 units) or not. Standard errors, clustered at the session level, are in parentheses. Wild cluster bootstrapped p -values for the risk-sharing coefficient, estimated with Webb weights due to the small number of clusters, are presented, as are the marginal changes, associated with risk sharing, in the probability of buying any insurance. Columns 2 and 4 control for participant sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in tropical livestock units and land owned in hectares, as well as enumerator fixed effects.

Table 8: Double probit hurdle regressions

	Units of indemnity insurance purchased		Units of index insurance purchased	
	(1)	(2)	(3)	(4)
<hr/>				
Pr($Y \in \{1, 2\}$):				
Risk sharing	-1.32 (0.38)	-1.77 (0.35)	1.53 (0.91)	1.81 (0.76)
Wild cluster bootstrap p -value	0.036	0.033	0.19	0.077
Pr($Y = 2 Y \in \{1, 2\}$):				
Risk sharing	-0.73 (0.43)	-0.92 (0.43)	0.50 (0.092)	0.16 (0.32)
Wild cluster bootstrap p -value	0.13	0.019	0.052	0.72
<hr/>				
Controls	No	Yes	No	Yes
<hr/>				
Marginal change in the probability of buying . . .				
zero units	0.13	0.16	-0.46	-0.43
one unit	0.20	0.21	0.15	0.24
two units	-0.33	-0.37	0.31	0.18

Note: Double probit hurdle regressions. The dependent variable is the number of units of insurance purchased (0, 1 or 2). Standard errors, clustered at the session level, are in parentheses. Wild cluster bootstrapped p -values for the risk-sharing coefficients, estimated with Webb weights due to the small number of clusters, are presented, as are the marginal changes, associated with risk sharing, in the probabilities of buying 0, 1 and 2 units of insurance. Columns 2 and 4 control for participant sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in tropical livestock units and land owned in hectares, as well as enumerator fixed effects.

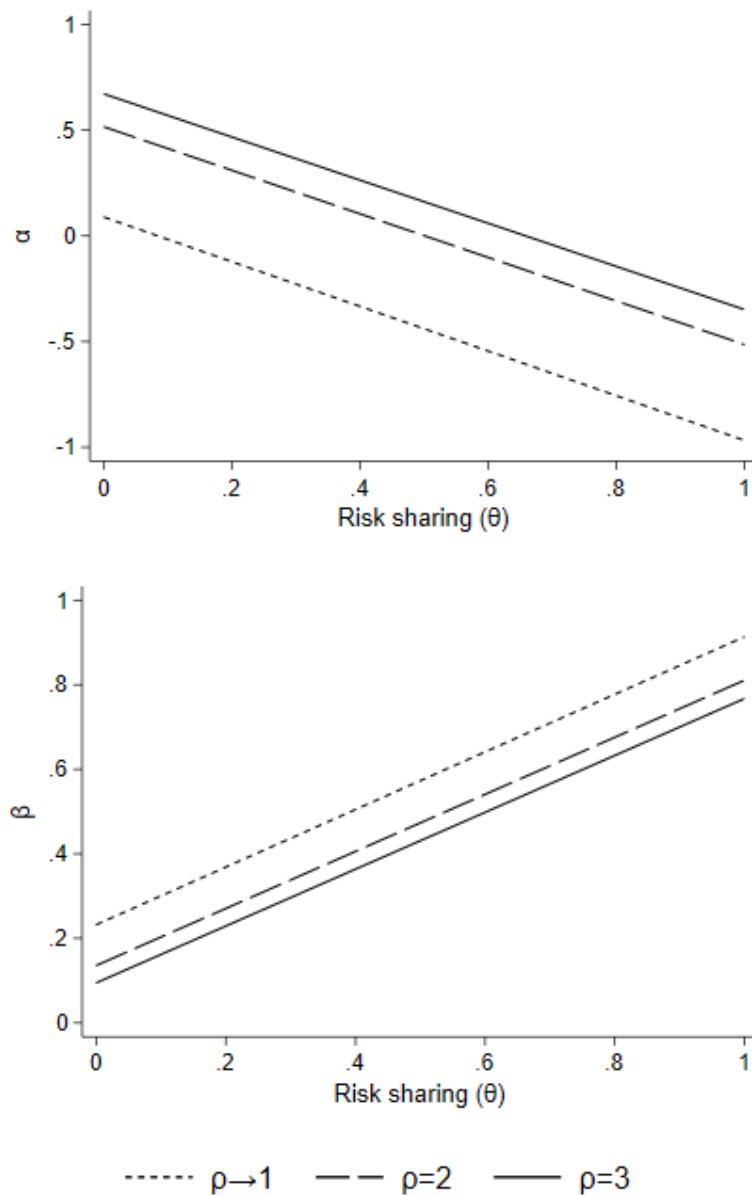
Table 9: Robustness check: Leaving out sessions one by one

	Units of indemnity insurance purchased						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Pr($Y \in \{1, 2\}$):							
Risk sharing	-1.77 (0.35)	-1.56 (0.33)	-1.66 (0.57)	-5.66 (0.84)	-2.53 (0.83)	-2.17 (0.34)	-1.67 (0.37)
Wild cluster bootstrap p -value	0.025	0.089	0.052	0.047	0.031	0.023	0.065
Pr($Y = 2 Y \in \{1, 2\}$):							
Risk sharing	-0.92 (0.43)	-0.58 (0.25)	-0.56 (0.31)	-1.25 (0.55)	-1.25 (0.45)	-0.99 (0.56)	-1.48 (0.60)
Wild cluster bootstrap p -value	0.016	0.094	0.12	0.024	0.037	0.081	0.025
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)
Pr($Y \in \{1, 2\}$):				
Risk sharing	1.81 (0.76)	2.64 (0.79)	1.47 (0.76)	2.53 (1.06)
Wild cluster bootstrap p -value	0.068	0.074	0.11	0.13
Pr($Y = 2 Y \in \{1, 2\}$):				
Risk sharing	0.16 (0.32)	0.47 (0.79)	0.41 (0.052)	0.16 (0.34)
Wild cluster bootstrap p -value	0.70	0.69	0.25	0.68
Controls	Yes	Yes	Yes	Yes
Session dropped	None	#4	#7	#8


Note: Double probit hurdle regressions, leaving out one session at a time. The dependent variable is the number of units of insurance purchased (0, 1 or 2). Standard errors, clustered at the session level, are in parentheses. Wild cluster bootstrapped p -values for the risk-sharing coefficients, estimated with Webb weights due to the small number of clusters, are presented. All regressions control for participant sex, age, age squared, marital status, household headship, literacy, household size, livestock owned in tropical livestock units and land owned in hectares, as well as enumerator fixed effects. Note that the estimation did not converge when dropping session #3, one of the index-insurance sessions.

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



Note: Numerically estimated optimal demand for indemnity insurance (α), in the upper panel, and index insurance (β), in the lower panel, 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.