

Utility, Risk, and Demand for Incomplete Insurance: Lab Experiments with Guatemalan Cooperatives *

Craig McIntosh[†], Felix Povel[‡], Elisabeth Sadoulet[§]

January 13, 2015

JEL Codes: C93, O22, I25

Keywords: Risk, Index Insurance, Utility Estimation

Abstract

We play a series of incentivized laboratory games with risk-exposed cooperative-based coffee farmers in Guatemala to understand the demand for index-based rainfall insurance. By varying risk in systematic ways over a large number of games, we can use numerical techniques to estimate a flexible utility function for each player, and simulate demand under alternate scenarios. We show that the introduction of a new risk that makes insurance probabilistic triggers a large decrease in willingness to pay that cannot be explained by expected utility but is consistent with prospect theory. Exploiting the group structure of the cooperative, we investigate the possibility of using group loss adjustment to smooth idiosyncratic risk. Our results help to explain why index insurance products have struggled to generate demand when they are introduced in environments with multiple risks to output, and suggest that while group loss adjustment can overcome this problem it introduces frictions that are unattractive to group members.

*Thanks to BASIS/USAID for financial support, for extensive discussions to Alain de Janvry, and for comments to seminar participants at PACDEV.

[†]University of California San Diego, ctmcintosh@ucsd.edu

[‡]Kreditanstalt für Wiederaufbau (KfW), Frankfurt, Felix.Povel@kfw.de

[§]University of California Berkeley, esadoulet@berkeley.edu

1 Introduction

The past decade has seen numerous attempts to extend risk protection to farmers in low-income countries through the use of index insurance. These products make pay-outs based on the value of a pre-defined index (of local rainfall, vegetation, or a more multi-dimensional set of factors) and hence can provide insurance against aggregated shocks without creating moral hazard (Barnett and Mahul, 2007). From a perspective motivated by the Townsend (1994) model of village-level risk pooling, these products appear ideal in that they insure precisely the correlated shock that cannot be smoothed locally. Yet, almost universally these products have met with disappointing demand when introduced in the field. Uptake rates in pure index insurance products have typically been very low, and several studies have found that interlinking index insurance with credit products actually dampens the demand for credit (Giné and Yang, 2009; Banerjee, Duflo, and Hornbeck, 2014). Understanding why such an apparently attractive financial service has met with such low demand is of increasing urgency as global warming amplifies the level of weather-induced risk that vulnerable farmers face.

A primary candidate, commonly discussed both in the empirical literature on index construction and in the theoretical/experimental literature on risk preferences, is that indexes provide only partial insurance. The use of an index necessarily means that the product fails to cover the full spectrum of risks to income, and the spatial aggregation of the index means that fine-grained variation will go uncovered. Unlike full insurance, the demand for incomplete insurance will display a complex relationship to risk (Pratt and Zeckhauser, 1987; Doherty and Schlesinger, 1983) and risk aversion (Gollier and Pratt, 1996), even in an expected utility framework. Clarke (2011) has suggested that the possibility of contract non-performance can drive precisely the most risk averse farmers away from incomplete insurance contracts if the worst state of the world can occur without a payout.

The behavioral literature has identified two distinct ways in which insurance products can be incomplete: *partial* insurance, in which the insurance always pays a fraction of the loss when the shock occurs, versus *probabilistic* insurance, in which it is possible that the insurance does not pay out at all when the shock occurs. Hypothetical surveys of insurance demand have found that demand for partial insurance conforms relatively well to expected utility theory, but drops substantially more than would be predicted when faced with a probabilistic insurance (Kahneman and Tversky, 1979; Wakker, Thaler, and Tversky, 1997). This would be consistent with the findings from prospect theory that many people (1) overweight small probabilities relative to their true likelihood and (2) are loss-averse,

given a small possibility of large losses.¹

The distinction between partial and probabilistic insurance is directly relevant to the design of index products. To the extent that the lack of demand is driven by the fact that the insurance is *partial*, offering products with higher premiums and more complete payouts should be welfare improving. Group insurance emerges as a particularly attractive option in this case because the localized variation in losses conditional on a shock will be hard to insure with an aggregated index, and yet this variation may be smoothed locally. This would allow the group to conduct loss-adjustment, thereby decreasing the extent to which payouts are partial. If, on the other hand, the drop-off in demand is coming from the *probabilistic* nature of the insurance, this suggests that the very idea of single peril insurance is the inhibiting factor, and insurers must do more to provide multi-peril (or directly loss-adjusted) products in order to see meaningful demand. To the extent that amplification of the underlying risk in either dimension has deleterious effects on demand, this would suggest that index insurance products will struggle in contexts with severe or multi-peril risks.

To better understand the barriers to index insurance demand, we conducted a set of controlled lab-in-the-field experiments with a very risk-exposed group, cooperative-based coffee farmers in Guatemala. During the course of an incentivized day-long exercise, we presented farmers with a way of visualizing the weather-driven risks to their farms. We presented them with numerous scenarios, in each case recording the willingness to pay (WTP) for an excess rainfall index insurance product whose attributes were held constant while the risk environment shifted. In every scenario the rainfall insurance was both partial (in that it covered only some of the yield risk in insured states) and probabilistic (in that it was possible to face a loss without a payout).

The starting point for our analysis is a set of seven exercises in which only the magnitude of the losses of the insured states differ across games. Wakker, Thaler, and Tversky (1997) provides evidence that the response to the partial-ness of insurance can be explained relatively well with expected utility theory. We therefore use numerical optimization methods to estimate a flexible utility function for each farmer using the variation in WTP across these scenarios. These utility functions can then be used to predict demand in any counterfactual scenario in the same state space, and so provide a straightforward way of distinguishing the effects in subsequent games that can be explained by expected utility theory versus those that cannot.

We then introduce a set of games featuring an increasingly likely and severe probabilistic

¹These are also consistent with the long-standing observation that many individuals are averse to ambiguous risks (Ellsberg, 1961; Segal, 1987; Bryan, 2010) and to complex compound lotteries (Budescu and Fischer, 2001; Elabed and Carter, 2014)

shock, which we frame as the advent of a ‘drought’ risk in an environment that has previously suffered only from excess rainfall risk. We find dramatic effects of even very low probability risks, a finding consistent with prospect theory but not with the predicted WTP estimated for each individual. The response to more severe risk, including those states in which the worst possible state can occur without a payout, is even stronger but proves to be in line with the responses we would estimate from expected utility theory. The behavioral puzzle in our results is thus an over-response to low probability risks.

Finally, we test the promise of group loss-adjustment as a potential solution to the problem of incomplete insurance. Variance in yields within the insured state can be seen as analogous to localized basis risk that occurs within the area covered by one rainfall station. To this extent, an alternative to attempting to design more localized indexes is to delegate loss adjustment to the group. If a group has low information costs and the ability to dynamically enforce contracts, internal loss adjustment will emerge as a welfare-enhancing contract relative to individual index insurance held by all members (Dercon et al., 2006). However, the delegation of loss adjustment to the group induces difficult questions of dynamic enforcement (Ligon, Thomas, and Worrall, 2002) and division, and in the end may prove too costly.

We begin by eliciting willingness to pay for a set of different group insurance options when the loss adjustment rules of the cooperative are externally set. We then relax the highly framed nature of the exercise, and hand decision-making power over to the real cooperative leadership to decide, and with this agreement in place again solicit demand. This juxtaposition of tightly defined individual demand parameters with real group deliberation provides a lens on the ways in which idealized versus actual group loss adjustment drives insurance group composition. Our results indicate that farmers understand the possibility and value of this type of loss adjustment, that they expect the cooperative to in fact conduct some loss adjustment if offered a group insurance product. Despite this expected benefit, dislike of the group modality is sufficiently strong as to depress net demand for group insurance.

The remainder of the paper is organized as follows: Section 2 provides the background and setting for the games, and a detailed description of the exercise. Section 3 uses the basic games to estimate the best-fit utility function for the data, a control structure that is then used throughout the paper. Section 4 provides results on the individual insurance games, Section 5 on the group insurance games, and Section 6 provides robustness checks, and Section 7 concludes.

2 Setting and Game Design

Coffee is by far the most important export sector in Guatemala, and yield in the coffee sector is quite variable with excess rainfall and hurricanes posing the primary source of weather risk exposure. In early 2010 we conducted a cooperative survey of the coffee sector in Guatemala. That survey attempted a census of every registered first-tier coffee cooperative in the country, and included data on 1,440 individuals in 120 cooperatives. For this exercise, we then selected from that population the 71 cooperatives that reported being vulnerable to excess rainfall risk (the product that this project is intended to pilot) and devised a set of games to understand the nature of index insurance demand. For each of the selected cooperatives we then attempted to draw in 10 individual members to participate in the day of laboratory experiments (the actual number that attended varies between 4 and 13, with 10 as the modal number).

Subjects in the study were presented with a sequence of scenarios, each featuring a carefully designed graphic illustrating the probability distributions in the insured and uninsured state in order to help the subjects visualize potential outcomes (see Appendix A1 for the graphics used in the games). All scenarios feature an excess rainfall index insurance product paying out 1,400 Quetzales in 1 year out of 7. The index is based on cumulative rainfall over the fruiting and flowering period for coffee as measured at the nearest government-administered rainfall station. This insurance product is partial for two reasons. First, the rainfall index is imperfectly correlated with yields on farmers' plots, thus providing some risk that is covered by the insurance product and some that is not (often referred to as basis risk in this literature). Furthermore the payout is calibrated to cover average input cost and not a losses. For each scenario we hold the basic attributes of the insurance itself constant (likelihood of payout, size of payout), and so all variation in the stated WTP across games arises from variation in the nature of the risk and from the specific modality used for the group insurance. The demands are incentivized by paying out experimental 'yields' that are 1/100th of the outcomes in a randomly chosen group of scenarios.

The games were typically played in the cooperative offices. The survey team that ran the games was comprised of a *presenter* who ran the sessions and read the scripts, an *enumerator* who would sit with the subjects and help them to fill in their sheets if they required assistance (25% of the respondents reported never having been to school), and two additional assistants. The heart of the variation across games was represented by a set of flip charts that represented the distribution of possible outcomes in as clear and realistic a fashion as possible. These graphics had quantities that were carefully calibrated based on information about average yields and typical losses from the baseline household survey, and

we repeatedly field tested the distributions in the graphics to find the most intuitive way of presenting the variation across games.

Upon arriving, subjects were walked through an intake survey asking a set of typical questions about household composition, wealth, education, risk exposure of the farm, as well as a set of behavioral questions focusing on risk aversion, ambiguity aversion, discounting, and present bias. The first and last exercises of the day were a marketing exercise that measured the WTP for a real commercial product on their own farm, using an index based on rainfall at the nearest rainfall station.

The remainder of the exercises are based on a tightly framed risk distribution, and consist of a series of 35 games grouped into eight exercises each of which focuses on a different issue. For each exercise subjects were asked to record their willingness to pay for the product, with the actuarially fair price remaining fixed at 200 Quetzales (\$31.73). These exercises completed during the course of the day were therefore the following:

1. ‘Variation in Insured States’ (I1-I7). Seven scenarios in which the variance or the severity of losses in the insured state are varied.
2. ‘Background Risk’ (I8-I13). Six scenarios in which the likelihood and severity of an uninsured loss is varied.
3. ‘Group, no rule’ (G1-G3). A repeat of scenarios I5-I7 in which variance of insured losses is increased, group loss adjustment is possible but the extent of loss adjustment is not specified.
4. ‘Group, with rule’ (G4-G6). Repeats of scenarios I5-I7 in which a specific loss adjustment rule is specified for each sub-game.
5. ‘Group Heterogeneity’ (G7-G11). Asking individuals to think of themselves as members of heterogeneous coops, altering the distribution of expected losses.
6. ‘Group Deliberation’ (G12-13). A semi-structured deliberation exercise in which groups are asked to debate and decide upon the loss adjustment rule they would use in practice.
7. ‘Real Values’ (I14-I16). A repeat of scenarios I1-I3 in which the payoffs are stated directly in their true monetary value (upside risk, 1/100th of the framed amount in the rest of the games).

Graphics illustrating the variation within and between games are provided in Appendix A; scripts for each of the games are given in the online appendix, and more detail on each game is provided as it is discussed in the text.

To protect the study against ordering effects and framing effects, we randomized two dimensions of the way in which the study is conducted. First, the brackets for the WTP worksheets were randomized at the cooperative level; half of the respondents spent the day using sheets that presented values between 40 and 320 Quetzales, and the other half between 80 and 360. This lets us examine the extent to which the framing of the price altered the resulting WTP. Secondly, we randomized the order of the games to the maximum extent possible. While the marketing exercises were always first and last, and while the ‘real values’ and the ‘deliberation’ rounds were always last within their respective groups (I or G), we randomized the ordering of the individual (I) and group (G) games, as well as the ordering of the games within the ‘variation in insured states’ set (I1-I7), as well as the ordering of group with rule (G4-G6) and group heterogeneity (G7-G11) games, leading to 8 possible ordering cells for the day’s games (Appendix Table 2).

Appendix Table A3 presents the results of an analysis of these ordering effects. The price bracketing did not lead to large variation in WTP (the \$6.35 difference between the ‘high’ and ‘low’ brackets led to an insignificant \$1.90 difference in average WTP), but the game ordering did have significant effects. Using fixed effects for games to identify the ordering parameter only off of the randomized games, we find that having a game come later in the day by one exercise lowered WTP by \$.40 per exercise. By contrast, players like whichever of the (Group, Individual) games they saw second, and so the overall sign on the group versus individual choice depends on the order in which the game was played. This makes it critical that we remove ordering effects from the remaining analysis; therefore we calculate randomization propensity weights (equal to the inverse of the proportion of the sample size in each randomization cell) and use these throughout the rest of the paper so that the ordering effects are completely removed from our estimates.

A final empirical issue on which we attempt to shed light is the basic promise of the use of this type of incentivized experimental game in understanding actual insurance demand. There are important ways in which the usual analogy of lab behavior to real-world actions breaks down in the study of insurance. Firstly, to understand insurance we must understand the effects of very large swings in income, and these cannot be replicated in the lab. Secondly, the core utility motivation of insurance is based on downside risk, and this cannot ethically be recreated in the lab (one can alter payouts but cannot confiscate income from experimental subjects). As a way of examining this fundamental epistemological question, after a day in which all insurance demand had been framed in amounts equal to 100 times the actual payouts in the game (and with a discussion of hypothetical losses to income), we then played one round where the same exercise was repeated but the payouts were framed in the actual winnings to players. By examining the shift in demand between the framed and unframed

scenarios we gain a window on the extent to which the framing of the payouts is revealing an insurance demand consistent with exposure to meaningful downside risk.

3 Estimation of a Utility Function Based on Revealed Demand for Insurance

The objective of this section is to estimate a utility function for each player based on its revealed willingness to pay the incomplete insurance scheme in the seven individual games I1–I7 that present variation in risk exposure in insured states. The estimated utility function can then be used to predict the willingness to pay that each player ought to have in any alternative risk scenario.

3.1 Model

Preferences are characterized by the following utility function:

$$u(y; k, \beta) = -\frac{1}{k}e^{-k\frac{y^{1-\beta}}{1-\beta}} \quad (1)$$

Despite having only two parameters, this utility function is quite flexible. Absolute risk aversion $ARA = \beta\frac{1}{y} + ky^{-\beta}$ decreases with income for $(\beta > 0 \text{ and } k > -y^{\beta-1})$ or $(\beta < 0 \text{ and } k < -y^{\beta-1})$, and increases with income otherwise. It converges to the CRRA function $u(y) = -\frac{1}{k}y^{-k}$ with $RRA = k + 1$ when $\beta \rightarrow 1$, and is the CARA exponential utility $u = -\frac{1}{k}e^{-ky}$ with absolute risk aversion k when $\beta = 0$.

Each risk experiment g presented to the players is characterized by a set of probabilities p_x^g for the states of nature with income x and payout C_x^g that the insurance will pay if the player is insured. In a given game, the expected utilities for an individual with preference parameters (k, β) without and with insurance are:

$$EU_0^g(k, \beta) = \sum_x p_x^g u(x; k, \beta) \quad (2)$$

$$EU_I^g(k, \beta, \delta) = \sum_{x=2}^{10} p_x^g u(x - wtp + \delta C_x^g; k, \beta) \quad (3)$$

where, wtp is the premium for the insurance and $\delta \in [0, 1]$ is a trust parameter that the agent places on the insurance payout. The addition of the parameter δ is prompted by the fact that observed willingness to pay was in most cases inferior to the fair price, which is not

conceivable with a standard utility function.

The willingness to pay for the insurance is defined as the maximum premium an individual is willing to pay for the insurance, i.e., the premium that makes the individual indifferent between being insured and not insured.

$$wtp(g, \theta) = (wtp : EU_I^g - EU_0^g = 0) \quad (4)$$

where $\theta = (k, \beta, \delta)$ denotes the vector of parameters of the model.

3.2 Econometric method

We proceed now with the estimation of a vector of parameters θ for each individual. We assume that there is some additive measurement error on the willingness to pay, such that the observed willingness to pay by a given individual wtp_g is:

$$wtp_g = wtp(g, \theta) + \epsilon_g \quad g = 1, \dots, 7 \quad (5)$$

We also assume the usual regularity conditions on the error ϵ_g such that our estimator is consistent and efficient. Let $\mathbf{X}(\theta)_{G \times 3}$ denote the matrix with characteristic element $\partial wtp(g, \theta) / \partial \theta_j$, $j = 1, 2, 3$. For each individual i , we use a non-linear least squares estimator:

$$\hat{\theta}_i = \arg \min_{\theta \in \Theta} \sum_{g=1}^G (wtp_g - wtp(g, \theta))^2 \quad (6)$$

implying that $\hat{\theta}_i$ must satisfy the first order conditions

$$-2\mathbf{X}(\hat{\theta}_i)^T (\mathbf{wtp} - \mathbf{wtp}(\mathbf{g}, \hat{\theta}_i)) = \mathbf{0}$$

Equation (6) describes a typical non-linear least squares problem, except that in addition to being nonlinear, the function $wtp(g, \theta)$ is only defined implicitly by equation (4). Thus, the derivatives with respect to θ that define the moment equations, and that are critical to any gradient-based solution algorithms, require application of the implicit function theorem at each trial value of θ .

3.3 Estimated preferences and predicted WTP

We start by estimating a unique utility function for all 674 players. Results for the parameters, with robust standard errors clustered at the individual level in parentheses, are reported in Table 1, col. 1. Absolute risk aversion is only slightly decreasing over the range

of values of income, from 0.73 to 0.80, implying that relative risk aversion increases very steeply from 1.6 (for the worst income equal to 20% of the normal income) to 7.3 when there is no negative shock to income.

We next proceed with the estimation of θ for each individual player. Since we rely on a very small number of observations for each player (at most 7, and less for the 61 players that did not play all 7 games), estimated parameters can take some extreme values. We therefore report the median and the lowest and highest 5th percentile of the estimated parameters in Table 1, col. 2-4. We see large variations in estimated parameters across individuals, reflecting heterogeneity in preferences.²

For each individual with parameter $\hat{\theta}$, we can compute the predicted WTP, $\widehat{wtp}(g', \hat{\theta})$ that the player ought to have for any game g' . As above, this is the solution to (4) for that particular game characterized by $p_x^{g'}, C_x^{g'}$. The process converged for 621 players for the first 3 games and 666 players for all other games.

Since these \widehat{wtp} will be used as regressors in the analysis of the observed WTP in games with background risk or in group games, we will need some measure of precision on these predicted values to correct the standard errors in the estimations. This is done by implementing a wild bootstrap using the 6-point distribution proposed by (Webb, 2013).³ With equal probability, the residual for each observation is multiplied by $\pm\sqrt{0.5}$, ± 1 , or $\pm\sqrt{1.5}$ to construct a bootstrap replicate of reported willingness to pay values. For each replicate we then re-estimate the optimal parameter estimates, and in turn compute the predicted $\widehat{wtp}(g', \hat{\theta})$. The wild bootstrap here assumes that errors are independent across observations, but allows them to be heteroskedastic and non-normal. Notice that because it is computationally intensive to repeat the gradient-based search for each bootstrap replicate, the bootstrap parameter estimates rely on a grid search method. The bootstrapped willingness to pay predictions will be directly used in the estimations that use \widehat{wtp} as regressors.

4 Demand for Individual Insurance

4.1 The effect of introducing drought

In seven of the individual games, we introduced the possibility that drought might occur. Given an excess rainfall insurance product, a drought is a source of loss to the farm that

²The small number of observations imply that standard errors on parameters are extremely high, and the quality of fit of the estimation, measured by $SSR_{i,q} = \sum_{g=1}^7 (wtp_g - wtp(g, \theta_q))^2$, very low. We thus do not report this information.

³With fewer than 10 observations, the 6-point distribution by Webb is recommended over the more common 2-point distribution (Cameron, Gelbach, and Miller, 2011).

is not protected by the insurance. To investigate how this source of uninsured risk drives demand, we hold constant the attributes of the insured shock and the insurance product (1 in 7 chance of incurring a loss of \$793, or 50% of potential income, triggering a payout of \$222) and vary the frequency and intensity of this uninsured shock. Because these scenarios feature variation in the probability that the insurance contract fails to perform, they correspond to varying *probabilistic* insurance, while the previous scenarios isolated the differences in WTP that arise from the extent to which the insurance is *partial*.

The six drought games began with a framing of a mild drought risk, one which was both unlikely to occur (1/21) and small (loss equal to 20% of potential income). Holding the insurance product constant, the magnitude of the drought-induced loss was then increased to 40% and 80% of potential income, and the likelihood of the risk was increased to 1/7, for each shock magnitude level.

Critically, at 80% potential loss the drought-induced outcome replaces the insured outcome as the worst one that can obtain. This has strong negative effects on predicted WTP (because the marginal utility of income is now highest in the drought state, the product has the possibility of moving income from worse states to better states. This stands in distinction to the first games, in which variation always occurred in insured states and thus an increase in downside risk should increase the WTP because insurance moves money into the worst state.

The presence of subsistence or other ‘safety first’ constraints will cause the worst possible outcomes to carry a heavy weight in decision-making. As noted by Clarke (2011), the possibility of the worst state being uninsured can introduce non-monotonicity into the relationship between risk aversion and insurance demand. The Maximin Expected Utility framework used by Gilboa and Schmeidler (1989) and Bryan (2010) evokes a pessimism in which decisionmakers fixate on the worst thing that could happen; another context in which the effect of these extreme tail risks would be accentuated.

We can first approach the presence of the drought risk by using our demand model to predict WTP for these new games, based on the utility curves estimated above and expected utility theory. We begin by investigating the basic premise that WTP for insurance should increase with the downside risk in insured states and decrease with downside risk in uninsured states. Column 1 of Table 2 shows that the WTP predicted from the utility curves displays the predicted relationships, with shocks in insured states driving up WTP and shocks in uninsured states driving it down. Column 2 shows that, as a result, predicted WTP falls slightly for a small uninsured risk and strongly for a larger one.

The main use of the predictions, however, is as a way of analyzing the extent to which actual patterns in WTP can be explained using expected utility theory, or if instead there

are responses to drought that require a behavioral explanation. To test this, columns 3 and 4 repeat the previous analysis but using the actual WTP observed across games. The patterns are similar overall, but we see relatively stronger response to mild drought risk and weaker response to severe drought risk than seen in predicted WTP. Column 5 then includes the predicted WTP as a control in analyzing the actual WTP, meaning that it removes the component of variation that can be predicted using our utility model, leaving any unexplained effects of mild or severe drought to be picked up in the dummy variables.

Here, we see that almost all of the negative reaction to severe drought can be explained with expected utility and that the ‘behavioral puzzle’ is in fact the over-reaction to mild drought risk in the actual WTP for insurance. The lack of a strong response to the worst possible state being uninsured does not appear consistent with the Maximin Expected Utility framework, or with Safety First decisionmaking.

4.2 Investigating utility

To investigate this point further, we proceed in several steps. First, given that the behavioral puzzle we need to explain is the very substantial drop in WTP for games featuring an uninsured (drought) risk, Table 3 runs interactions between risk aversion and ambiguity aversion and these games to illustrate how these behavioral parameters differentially drive demand in these scenarios. As seen in column 1 and 2 risk aversion should increase WTP in these games even under expected utility, and columns 3 and 4 confirm that it does in reality. In contrast, there should be no relationship between ambiguity aversion and the WTP for insurance with this small uninsured risk (columns 2 and 3), but in reality being ambiguity averse causes a sharp drop in WTP. This is consistent with the idea that behavioral issues lead to this drop in demand.

Finally, we are interested in using these games to understand whether the prediction found in Clarke (2010) of an inverted-U shaped relationship between risk aversion and WTP for insurance that can require a premium payout in the worst state of the world. To ask this question we focus on the two games in which the drought risk is so severe that it is possible that this may occur, and we interact a dummy for the worst games with a dummy indicating the high and low risk aversion groups (with those of moderate risk aversion serving as the omitted category). In column 2 we see that even our predicted demand does not follow this inverted U-shape, displaying instead a strong monotonically decreasing WTP with risk aversion. When we examine actual WTP in Column 4, we again fail to see an inverted U; rather there is a slightly elevated WTP for those with the lowest risk aversion and no difference between the moderately and highly risk averse. Taken together, these results

suggest that expected utility does a very poor job of explaining the strong response to small uninsured risk, and that this behavior is mostly localized to the ambiguity averse.

4.3 Using lab experiments to study insurance demand

The analysis of insurance demand using laboratory experiments is subject to fundamental problems that would not be present in studies of more routine, high-frequency behavior. First, the purpose of insurance is to cover individuals from drastic changes in income, and these cannot easily be replicated in the relatively low-stakes environment of the laboratory. Second, we are primarily interested in responses to negative shocks, and economic experiments can only ethically induce upside risk. Third, insurance decisions play out over long periods of time, and markets are built as individuals learn about the qualities of various forms of coverage by trying them. The quick decisions and payoffs based off of a one-day experiment explicitly do not capture the learning that would occur in a real market, and are thus in some sense measuring demand for the initial market only.

Despite this, the use of games to explore insurance-related questions has nonetheless been expanding (Lybbert et al., 2009; McPeak, Chantarat, and Mude, 2010; Norton et al., 2012). The stated purpose of these field studies has been mixed; some have intended to use the games only as a tool for learning about insurance, some have seen it as a marketing exercise, and others have sought to learn about the nature of demand from these exercises. We attempted to address each of these issues in the design of the experiment. We selected the most risk-prone coffee cooperatives from a larger census in Guatemala, and repeatedly field-tested the scripts and graphics to provide as realistic and clear as possible a picture of the issues surrounding group index insurance. Further, the variation in the potential actual winnings during the day was large (numbers). We framed the games heavily around issues of vulnerability to excess rainfall, basis risk, and the tradeoffs inherent in group insurance. Agricultural profit and hurricane loss values were carefully calibrated using data from the baseline survey. We also explicitly incorporate group deliberation and capture the evolution of preferences for group insurance before and after discussion of the most strategically fraught issue, loss adjustment. The sharp responses to variations across games present throughout the paper suggest that participants were sensitive to experimental variation.

Just how successful were we in framing the small upside risk of the actual experiment to reflect the large downside risk of the scenarios? As a way of exploring this we played three rounds in which the payoffs were unframed; three scenarios they had already seen were repeated, but rather than presenting the payoffs in their framed values, the actual payoffs they would win for the day were presented. In this analysis we would hope to see some WTP

for insurance in the unframed games; a lack of demand for insurance would suggest that the stakes were not large enough to induce risk aversion. We would also hope to see additional WTP in the scenarios framed as large losses, suggesting outcomes spread further out the distribution of the individual's utility curve.

Table 4 presents this comparison, beginning with Predicted WTP as a simple way of showing what individuals should have been willing to pay for an insurance product that effectively protects them against a very small upside risk. Given the small amounts of money at risk, had individuals shifted completely to the unframed outcomes they would have been willing to pay only 21 cents for the unframed insurance, relative to \$29.42 for the framed insurance. Instead, we see in columns 3 and 4 that Actual WTP in the unframed scenarios was on average \$20, and the framing increased this by an additional \$9. Column 2 shows that all of the response to variance should have come in the framed shocks only, while column 4 shows that in actuality the framing only doubled the marginal response to variance relative to the unframed scenarios.

These results are heartening in that they provide support both for the fact that the incentives provided in the laboratory environment can capture risk-averse behavior, and that framing can push purchase decisions in the direction we would expect. They are also difficult to interpret in that we always played the unframed games after the framed games, and thus we actually have an experiment in unframing, not framing. Particularly given the discussion in the previous paragraph, the WTP in the unframed scenarios is consistent with imperfect unframing, meaning that a sample that was willing to pay \$29 to protect themselves against risks with framed standard deviation of \$282 and an actuarially fair price of \$31.73 were willing to pay \$20 to insure themselves against unframed risks with a standard deviation of \$1.97 and an actuarially fair price of \$1.40. Thus while our results indicate that framing is effective, we clearly failed to completely unframe the decision during these three rounds. We nonetheless take these results as confirming the idea that the analysis of framed games does capture risk aversion over quantities substantially larger than those actually at risk.

5 Incorporating the Group

Group insurance has a strong theoretical motivation, but is it something that farmers want? We begin by posing this simple question, and then undertake a series of exercises intended to decompose WTP for the group mechanism into three component parts: the expected insurance component, the raw preference for the group once this benefit has been netted out, and the desire for the group to serve as a redistributive mechanism.

We address this by attempting to decompose the demand for group insurance using a set of experimental exercises. First, we present the fundamental idea of the group insurance to players, discussing directly the possibility that the group can help to smooth individual risks. Then, three of the games from the individual section which feature variation in the severity of risk in insured states are replayed, this time in the context of group insurance with these risks being idiosyncratic. Then, the same games are repeated but we now provide an explicit rule for the extent to which idiosyncratic risks are to be shared.

By systematically varying the amount of loss adjustment the group will conduct, we can isolate demand for insurance through the group mechanism. A reference point for this analysis is to begin with simulating what the demand for these group insurance products would be if the only benefit of the group were pooling. Column 1 of Table 5 shows that predicted WTP for the no-pooling group is exactly the same as the individual insurance (by construction), and that the WTP should increase by \$7.19 as the group exercises its maximal risk pooling capacity. Column 2 shows that while the actual WTP reveals a demand for risk protection that is about half of what we would have expected \$4.41, as well as a pure dislike of the group mechanism (-\$5.21) that is sufficient to completely overwhelm the risk protective benefits, leaving overall demand for the most protective group product just equal with demand for the individual product. Extending the analysis to scenarios with larger variance of shocks within community in column 3 exhibits similar result.

Does this dislike of the group arise from distrust of the group mechanism as an agent for loss adjustment? To address this we interact features of the group games with a survey-based index of trust in the cooperative (a normalized sum of Likert responses to four questions on trust, transparency, and rule-following in cooperative decision making). Column 4 shows that indeed the penalty to group loan products is smaller in trusting groups (a one standard deviation increase in trust increases WTP by \$.93), but column 5 shows that the magnitude of the effect is unrelated to the degree of loss adjustment that the group will undertake. In other words, lack of trust that the group will be able to conduct loss adjustment does not appear to be the driver of the overall higher WTP in trusting groups.

Having understood how much the groups are willing to pay for loss adjustment, we want to understand the actual degree of loss adjustment that the players expect from their groups. In other words, not *can* then loss-adjust but *will* they loss-adjust? Column 6 shows the results of an exercise conducted before the explicit presentation of pooling rules was conducted, in which we asked group WTP with the degree of loss-adjustment left unstipulated.⁴ By comparing WTP in this game to those in which it was stipulated, we can measure expecta-

⁴Because these two games were always played in the same order within the overall randomization we cannot control for the possibility of sequencing effects between these games.

tions over pooling in a very exact way. The coefficient on this unstipulated game is -\$3.62, relative to a moderate pooling group WTP of -\$2.23 and a no-pooling WTP of -\$5.20. Column 7 pools data from all three risk levels and arrives at very similar conclusions. This suggests that farmers expect the groups would conduct roughly 25% of the possible pooling of idiosyncratic risk. Finally, we can ask whether a lack of group trust effects the extent of pooling that the members expect from the group. This is accomplished in column 8 by interacting group trust with a dummy for the game in which the sharing rule was not stipulated; here we see an insignificant effect.

These results point to a complex and contradictory set of factors at work in the demand for group insurance. On the one hand, farmers are willing to pay for the loss adjustment that groups are able to conduct. On the other hand, they do not expect groups to conduct much of it, and on the whole there is a dislike of the group that is roughly equal in magnitude to the WTP for the maximum extent of risk reduction that group loss pooling can achieve. The WTP for marginal risk reduction achieved through group loss adjustment is about half of what it should be based on utility structures. Group trust decreases the magnitude of the penalty levied on group insurance products, but the mechanism for this is neither the extent nor the credibility of loss adjustment to be conducted by the group.

5.1 What is the effect of heterogeneity on group insurance demand?

Having considered so far the implications of idiosyncratic, mean-zero variation on the demand for group insurance we now take on a different issue that will effect demand, namely the possibility that exposure to shocks is asymmetric across members. If certain people are subject to more extreme shocks (because, for example, they are insuring steep or flood-exposed farmland) then loss adjustment will introduce expected transfers towards these individuals from those who are less exposed to risk. This alters the actuarially fair premium. The greater the heterogeneity within a group in the exposure to these shocks, the more difficult we would expect group contracting to be.

To investigate this, we introduced five scenarios in which the group was presented as being composed of heterogeneous members with different potential loss severity. In the first, we merely presented the issue of heterogeneity, but the player's exposure to risk is the same as the average in the group. We then went through four games in which we asked players to consider their WTP for group insurance if their expected losses changed while the aggregate risk exposure in the group stayed constant. In each case we maintained that there would be partial risk pooling, and gave concrete amounts to be pooled for each scenario. These five

games give the basic dislike of heterogeneity, and the change in WTP as the expected losses to that individual change.

In Table 6 we show in column 1 that simply framing the group as consisting of heterogeneous membership drives down WTP by \$6.54, an amount greater than the overall penalty to group insurance. We then place the individual in different parts of the expected loss distribution, meaning that group loss adjustment would predictably serve as a transfer to or from that individual. As a way of understanding what this move in expected payouts should have done to demand, again utilize our utility structure to predict WTP. Column 2 shows that Predicted WTP from the utility models should have decreased by \$1.21 for each dollar to be transferred (this number is less than negative one because the money is transferred in the worst states), while column 3 shows that the actual WTP drops by only \$.62. Thus, the demand disutility of making transfers to other group members is only half of what it is when the transfers are to the insurance company. Columns 3 and 4 repeat this analysis showing each cell of the game separately; the results indicate that the divergence between the two types of WTP is particularly pronounced when an individual is the one least exposed to shocks.

The takeaway from this analysis is that while group heterogeneity depresses demand for group insurance, and individuals do respond in the predicted way to their own shock exposure relative to the rest of the group, these individuals are only half as unwilling to transfer money to each other to reduce inequality as they are to lose money to the insurance company.

5.2 Evolution of demand: results of the group deliberation exercise.

Up to this point, we followed a very structured approach to understanding group loss adjustment so as to be able to isolate this dimension of group insurance. We also wanted to understand the process of deliberation that would transpire as the group discussed the pros and cons of group insurance. We first presented the idea that groups could loss adjust, framed the pros (better risk protection) and the cons (tensions within the group), and asked players as individuals what degree of loss adjustment they would prefer (1 = none, 2 = moderate, 3 = as much as possible) if they were obtaining group insurance. We then asked them to discuss and decide upon this issue as a group, and a moderator recorded the nature of this discussion (length, conflict, unanimity in decision making, main issues discussed, and the specific individuals who talked and who explained concepts) and its final outcome. To mimic the incentive to renege on group risk sharing, we then asked each individual to

draw an actual rainfall shock (and thus a level of income) and to vote again on the group risk pooling decision. These three outcomes (pre-deliberation individual preference, group choice, and post-shock individual preference) provide a window directly into the desirability of this theoretically central feature of group insurance.

We provide a descriptive quantitative analysis showing correlation on how different types of individuals moved through this process from individual decision to group deliberation and decision, and then a final decision to collaborate once the individual shock has been realized. Column 1 of Table 7 shows that players that are risk averse or ambiguity averse have a lower preference for sharing, although the point estimates are small. As we move through the decision tree in column 2 to 5, we see that groups with a high fraction of women see decreased debate, increased unanimity, and group decisions with a higher preference for pooling. Groups with many ambiguity averse members, on the other hand, see decreased debate and quicker decisions that lead to less pooling. Trust in the cooperative shows little correlation with the outcome of the group deliberation process, and the educated like to debate. Risk aversion (as measured from predicted utility) displays the interesting relationship of depressing desire for pooling when asked as an individual, but elevating it in the final group decision.

Once the group deliberation had concluded, we had individuals draw a random shock from the distribution they were pooling over, and we privately asked again what their pooling preference would be now that they knew their shock draw. The extent to which preferences for sharing are altered in this interval provides an application of withdrawing the Rawlsian Veil of Ignorance, as agents who had previously not understood their exact position in a shock redistribution now know what they personally stand to win or lose. The magnitude of the change in behavior provides some evidence for the extent to which the inability to writing binding contracts will pose a constraint on pooling agreements that must be ex-post incentive compatible.

In columns 1 and 5 of Table 7, we see what are in effect balance checks since these are outcomes observed prior to the draw, and in column 6 we see how the realization of the randomized shock drives the change in decisions between the group and final individual round. Consistent with ex-post desire to renegotiate, as the shock realized by the individual goes up by 1,000 dollars the desire to engage in sharing rises by 6 percentage points. Hence, even in the relatively contrived setting of the laboratory and with very little time transpiring between the two decisions, we find evidence that ex-post renegotiation will put pressure on informal pooling agreements.

5.3 Evolution of demand: results of the group deliberation exercise.

The coefficients on the desired degree of risk sharing can be taken back to the coefficients from Table 5 in which the expected degree of risk pooling is estimated. Across all three exercise, participants report wanting ‘moderate’ risk sharing (50% of potential), and yet they expect that the groups will only provide 25% of the potential risk sharing. Given our evidence that the dynamic consistency of risk sharing is a problem in practice, the expectation that actual risk pooling will come in below the level desired may be well justified.

6 The Games as a Marketing Instrument

A final way of considering the effect of the games is to look at how the demand revealed by the marketing survey changes from the WTP measured at the beginning of the day to the final marketing survey once all experimental rounds have been played. While we place less credence in the absolute magnitude of the WTP, we believe that the changes in WTP from before to after the exercise, and particularly the comparison of the group to individual demand, are informative. Table 8 shows that overall the demand in the first marketing rounds was \$1.31 lower than in the game. Column 2 illustrates that demand after the day of games had increased by \$7.08, confirms the results from the framed exercise that WTP for group insurance is lower than for individual, and finds no differential change in WTP for group insurance as a result of the day’s activities. This increase in WTP translates into an increase in the share of the producers willing to buy the insurance at its fair price from 26% before the game to 51% at the end of the day (column 3). This result is heartening in the sense that individuals who had spent the day thinking about the advantages and drawbacks of index insurance saw a substantial increase in their demand for the product. Interestingly, though, the exhaustive consideration of the potential loss adjustment benefits from group insurance did not lead to any additional enthusiasm for group-based insurance. Nor has the day differentially influenced those with higher aversion to ambiguity (column 4).

7 Conclusion and Discussion

Using a set of artefactual field experiments, we investigate the demand for index insurance among coffee farmers in Guatemala. Willingness to Pay is in general lower than the actuarially fair rate, which is an initial piece of evidence that partial insurance products do not generate the kind of demand that we would expect from risk-averse agents if offered

perfect insurance. We use the lab context to decompose the potential reasons that insurance products may meet with limited demand, and to investigate several promising modalities to stimulate demand.

Using variation in how partial the insurance is we first estimate a flexible utility curve for each individual. This allows us to predict the demand under different scenarios in the same state space that we think should obtain if expected utility theory drove all WTP decisions. Taking these estimates of WTP to a set of scenarios featuring probabilistic insurance we find results consistent with the behavioral literature on this topic. Individuals respond far more strongly than we would expect to a slight increase in the possibility that a shock occurs and yet the product does not pay out. We investigate the mechanisms that might lead to this result even under expected utility theory and do not find support for them, and hence conclude that a behavioral dislike of probabilistic insurance is to blame. These results are all consistent with prospect theory, which says that individuals are loss averse and overweight small probabilities relative to large ones.

We then move on to analyze whether a group loss-adjustment mechanism could provide sufficient decrease in risk as to make index insurance palatable to farmers. Here we verify the mechanisms at play; farmers understand the risk pooling benefits of loss adjustment, and indeed they expect their cooperatives would provide about a quarter of the possible degree of risk pooling. Despite this, there is a secular dislike of the group mechanisms, increasing in the degree of distrust of the cooperative, that makes even a fully loss-adjusted group insurance product only just equal to individual insurance. Given the expected degree of loss adjustment, the average individual would prefer individual insurance to group. Heterogeneity in the group further damages the desire to use the group mechanism to cross-insure. Hence, while we verify that the underlying mechanisms making group insurance attractive are at play, in the end it is not.

In summary our results isolate several reasons for the low demand that index insurance products have met with in the developing world. Index insurance will struggle to generate demand in environments with multiple risks, and group insurance does not appear to be an attractive way to overcome this hurdle.

References

- Banerjee, Abhijit, Esther Duflo, and Richard Hornbeck. 2014. “(Measured) Profit is Not Welfare: Evidence from an Experiment on Bundling Microcredit and Insurance.”
- Barnett, Barry J. and Olivier Mahul. 2007. “Weather Index Insurance for Agriculture and Rural Areas in Lower-Income Countries.” *American Journal of Agricultural Economics* 89:1241–1247.
- Bryan, Gharad. 2010. “Ambiguity and insurance.” *Unpublished manuscript* .
- Budescu, David V and Ilan Fischer. 2001. “The same but different: an empirical investigation of the reducibility principle.” *Journal of Behavioral Decision Making* 14 (3):187–206.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller. 2011. “Robust inference with multiway clustering.” *Journal of Business and Economic Statistics* 29 (2).
- Clarke, Daniel J. 2011. *A theory of rational demand for index insurance*. Department of Economics, University of Oxford.
- Dercon, S., J De Weerd, T. Bold, and A. Pankhurst. 2006. “Group-based funeral insurance in Ethiopia and Tanzania.” *World Development* 34 (4):685–703.
- Doherty, Neil A and Harris Schlesinger. 1983. “Optimal insurance in incomplete markets.” *The Journal of Political Economy* 91 (6):1045–1054.
- Elabed, Ghada and Michael R. Carter. 2014. “Compound-Risk Aversion, Ambiguity and the Willingness to Pay for Microinsurance.” Tech. rep., University of California, Davis.
- Ellsberg, Daniel. 1961. “Risk, ambiguity, and the Savage axioms.” *The Quarterly Journal of Economics* :643–669.
- Gilboa, I. and D. Schmeidler. 1989. “Maxmin expected utility with non-unique prior.” *Journal of Mathematical Economics* 18 (2):141–153.
- Giné, Xavier and Dean Yang. 2009. “Insurance, credit, and technology adoption: Field experimental evidence from Malawi.” *Journal of Development Economics* 89 (1):1–11.
- Gollier, Christian and John W Pratt. 1996. “Risk vulnerability and the tempering effect of background risk.” *Econometrica* :1109–1123.
- Kahneman, Daniel and Amos Tversky. 1979. “Prospect theory: An analysis of decision under risk.” *Econometrica* :263–291.
- Ligon, E., J.P. Thomas, and T. Worrall. 2002. “Informal insurance arrangements with limited commitment: Theory and evidence from village economies.” *The Review of Economic Studies* 69 (1):209–244.

- Lybbert, T. J., F. Galarza, J. McPeak, C. B. Barrett, S. Boucher, M. R. Carter, and A. G. Mude. 2009. "Dynamic field experiments in development economics: Risk valuation in Morocco, Kenya and Peru."
- McPeak, John, Sommarat Chantarat, and Andrew Mude. 2010. "Explaining index based livestock insurance to pastoralists." *Agricultural Finance Review* 70 (3):333–352.
- Norton, M. T., D. Osgood, R. Diro, and M. Gebremichael. 2012. "Do experimental games increase take-up rates for index insurances? A randomized control trial approach." Annual Meeting Agricultural and Applied Economics Association, Pennsylvania.
- Pratt, John W and Richard J Zeckhauser. 1987. "Proper risk aversion." *Econometrica* :143–154.
- Segal, Uzi. 1987. "The Ellsberg paradox and risk aversion: An anticipated utility approach." *International Economic Review* :175–202.
- Townsend, Robert M. 1994. "Risk and insurance in village India." *Econometrica* :539–591.
- Wakker, Peter, Richard Thaler, and Amos Tversky. 1997. "Probabilistic insurance." *Journal of Risk and Uncertainty* 15 (1):7–28.
- Webb, M. D. 2013. "Reworking wild bootstrap based inference for clustered errors (." Working Paper 1315, Queen's Economics Department.

Tables

Table 1: Estimated parameters

Parameters	Overall U	Individual utilities		
	Coeff. (se)	Median	Lowest 5%	Highest 5%
$\hat{\beta}$	0.042 (0.120)	.720	-.959	3.64
\hat{k}	0.801 (0.194)	.849	-4.033	50.500
$\hat{\delta}$	0.156 (0.004)	.217	.0737	1.0347