

The demand for insurance under limited trust: Evidence from a field experiment in Kenya*

Stefan Dercon,[†] Jan Willem Gunning,[‡] and Andrew Zeitlin[§]

June 2015

Abstract

In spite of strong theoretical reasons to believe in the welfare-enhancing value of microinsurance products, demand for such products to date has been disappointingly low across a range of developing countries. In this paper we investigate the role of trust in the demand for indemnity insurance. First, we develop a theoretical model of insurance demand under limited trust to derive predictions for the way trust, risk aversion, and insurance premiums interact. Second, we test these predictions using field and laboratory-experimental data from a randomized controlled trial introducing a composite health insurance product to tea farmers in Kenya. Consistent with the theory, we find that not only low trust but also risk aversion is *negatively* associated with insurance demand, and that individuals with low trust are more responsive to experimental variation in premium costs. Third, we combine take-up decisions with subjective probability distributions for health costs to structurally estimate the model. Structural estimates reveal that choices are consistent with pessimistic and heterogeneous beliefs about the probability of insurance payouts for indemnified events. These estimates allow us to calculate welfare losses relative to counterfactual insurance products that are (perceived as) fully credible: expected losses from foregone insurance due to low trust exceed 31 percent of premium costs. Our results suggest that limited trust is an important barrier to the adoption of insurance, particularly among poor and risk-averse households who stand to benefit the most from such financial products.

*We thank Michal Matul (ILO), Ann Kamau and David Ronoh (CIC Kenya), Therese Sandmark (SCC), and Edward Kinyungu and Morris Njagi (Wananchi SACCO) for their support in the design and implementation of the project. Job Harms, Naureen Karachiwalla, and Felix Schmieding provided excellent research assistance. We are grateful to George Akerlof, Abigail Barr, Lori Beaman, Martin Browning, Gharad Bryan, Daniel Clarke, Mark Rosenzweig, John Rust, Laura Schechter, and participants at the Stanford Summer Institute for Theoretical Economics for helpful comments. Financial support from the ILO Microinsurance Innovation Facility is gratefully acknowledged.

[†]University of Oxford

[‡]VU University Amsterdam

[§]Georgetown University

1 Introduction

Risk, and its mitigation, are widely considered an important source of welfare losses in developing countries. Shocks to health and income appear to have long-lasting impacts (Alderman, Hoddinot and Kinsey 2006, Beegle, De Weerd and Dercon 2008). The mitigation of risk may lead to foregone investment opportunities with substantial expected returns (Rosenzweig and Binswanger 1993, Morduch 1995, Mobarak and Rosenzweig 2013, Karlan, Osei, Osei-Akoto and Udry 2014). High costs of participation in informal insurance networks can restrict resources available for investment in economic activities.

Seen in this light, recent empirical evidence of the demand for microinsurance is puzzling. Not only is demand for both indemnity and index-based insurance products often low (Banerjee, Duflo and Hornbeck 2014), but the likelihood of insurance purchases is *negatively* associated with measures of risk aversion in many contexts (Cole, Gine, Tobacman, Topalova, Townsend and Vickrey 2013). While in an index-based insurance context this may be explained by basis risk—such that low demand for such products is entirely rational¹—many relevant risks, such as health, are poorly addressed by index products, and the low demand for indemnity products such as health insurance remains a puzzle. Suggestively, in some studies measures of trust are positively associated with insurance demand (Cai, Chen, Fang and Zhou 2010, Cole et al. 2013).

In this paper, we develop and test a model of *limited trust* to explain the low uptake of indemnity insurance products. We define trust in the insurer as the potential policyholder’s perceived likelihood that a claim would be paid in the event of a loss. Clearly, a lack of trust can reduce the demand for insurance. We show that it can also explain the presence of a negative relationship between risk aversion and insurance demand. Intuitively, a reduction in trust increases the likelihood of the ‘worst-case’ outcome, in which an insurance premium is paid and a loss is suffered, but no claim is paid. This outcome is particularly threatening to the risk averse.

We confront the empirical implications of this model with data from a randomized, controlled trial offering *Bima ya Jamii*, a composite health insurance policy sold at close to actuarially fair prices to tea farmers in Kenya. This field experiment was a factorial design, cross cutting individually-randomized variation in premium costs with cluster-randomized treatments of basic marketing and marketing paired with financial literacy training. Our findings also show relatively low uptake, and little impact of financial literacy training. We combine data from the field experiment with data from two laboratory-type experiments conducted in the field at baseline, a trust game (Berg, Dickhaut and McCabe 1995) and a Holt and Laury gamble-choice game (Holt and Laury 2002), which provide measures of trust attitudes and risk preferences respectively. These allow us to test hypotheses about the interaction between experimental variation in prices and individual levels of trust.

Combining the data from the laboratory experiments and the field experiment allows us to test three key implications of the theory: insurance demand is negatively associated with risk aversion as measured in the lab, and positively associated with trusting behavior. Moreover, we show that individuals who exhibit low trust are more responsive to premium costs, as predicted by theory.

To illustrate the scope for policies to improve outcomes by raising trust, we combine these data with subjective probability distributions for hospitalization costs to structurally estimate our model of insurance demand under limited trust. Allowing for heterogeneity in trust levels, we find that consumers’ decisions are consistent with very low levels of trust in insurers: consumers decisions are consistent with perceived probabilities of indemnity payouts between 0.24 and 0.47. Using

¹See Clarke (2011) for a theoretical explanation. Low demand for index products is evident in Cole et al. (2013) and contrasts with the high uptake in Karlan et al. (2014). Mobarak and Rosenzweig (2013) provide direct experimental evidence of the importance of basis risk in these decisions.

these structural estimates to estimate willingness to pay under counterfactual trust levels, we find that the welfare costs of limited trust are substantial—amounting in expected value to at least 31 percent of the face value of the premium—and that these costs are borne disproportionately by the poor.

The contribution of this paper is fourfold.

First, the model we develop extends existing theoretical work on other contracts and products to the case of developing-country indemnity insurance. The model is related to Doherty and Schlesinger (1990), who study insurance demand with a possibility of insurer default, and Clarke (2011) and Mobarak and Rosenzweig (2013), who study the demand for index insurance in the presence of basis risk. In an indemnity (livestock) insurance context, Cai and coauthors (2010) make the related observation that farmers who believe such insurance is unlikely to pay out in the event of a loss are less likely to purchase insurance. We show that when indemnity payouts are uncertain—including from the subjective uncertainty arising from low trust—then demand for insurance can be decreasing in risk aversion. This offers an explanation for an empirical puzzle.

Second, our empirical results shed light on policy-relevant constraints to uptake of financial products. We find that financial literacy training has no effect on insurance demand, while our model and results give evidence that the perceived enforceability of claims for indemnified losses is an important constraint to insurance adoption. Thus we contribute to the increasingly mixed evidence of the limited efficacy of financial literacy training in this domain (Cole, Sampson and Zia 2011, Cole et al. 2013), a trend that suggests that either the curricula in these experiments have been poorly targeted, or that this constraint is not binding. On the other hand, insurers are well positioned to shape perceptions about the likelihood of paying claims,² and regulators have tools at their disposal that can improve confidence in these outcomes as well. Our model suggests that such policies will have attractive distributional properties, since limited trust is a particular deterrent to the insurance uptake of the risk averse—precisely those who stand to benefit from insurance the most.

Third, we contribute to a growing literature that quantifies the welfare costs of market failures in the provision of insurance (Einav and Finkelstein 2011). This literature comprises both ‘sufficient statistics’ and structural approaches (Chetty and Finkelstein 2013). The structural approach we take here allows us to simulate welfare losses relative to counterfactual trust levels, under alternative assumptions about the trustworthiness of the insurance product studied. In this respect, we contribute to recent research that has sought to quantify welfare losses arising from information frictions in the demand for insurance (Spinnewijn 2014, Handel and Kolstad forthcoming).

Fourth and more generally, our results shed light on the role of trust in financial sector deepening and economic development. Measures of trust are associated with growth rates across countries (Knack and Keefer 1995, Zak and Knack 2001). For example, in Africa evidence suggests that mistrust acts as a causal mechanism linking the slave trade to contemporary economic outcomes (Nunn 2008, Nunn and Wantchekon 2011). Yet the potential mechanisms through which trust matters for economic development are multiple—including both the strength of political incentives (Easterly, Ritzen and Woolcock 2006) and the costs of enforcing contracts (North 1990, Zak and Knack 2001). These mechanisms are difficult to distinguish in cross-country data. Using micro data from Peru, Karlan (2005) has shown that laboratory measures of *trustworthiness* are predictive

²Evidence from India suggests that endorsements by trusted authorities may have this effect (Cole et al. 2013). We take this evidence as suggestive, since trust is difficult to cleanly manipulate experimentally: product endorsements are potentially confounded by peer pressure, and proxies like past payouts (as considered by Cai et al. 2010 and Karlan et al. 2014) are confounded by income shocks. Our empirical approach, which combines incentive-compatible measures of generalized trusting behavior with experimentally induced variation in insurance premiums to test further theory-derived hypotheses, is complementary to such evidence.

of microcredit borrower behavior, but he finds no association between *trusting behavior* and these financial transactions. By contrast with microcredit, insurance places the burden of trusting on the part of the client rather than the financial institution. We show that trust limits the scope for such financial transactions.

The remainder of the paper proceeds as follows. Section 2 presents a simple model of indemnity insurance under limited credibility, and derives testable implications. Section 3 describes the field experiment, as well as survey and laboratory data collected. Section 4 tests empirical implications of the theoretical model, and discusses the robustness of these results to alternative explanations. Section 5 presents structural estimates of the theoretical model and their welfare implications. Section 6 concludes.

2 A model of insurance demand under limited trust

Here we develop the empirical implications of a model of indemnity insurance demand with limited trust. When agents have limited trust in insurers—such that the expected value of a policy is increasing in (subjective) trust levels—this generates a non-monotonic, and potentially negative, relationship between levels of risk aversion and insurance demand. This provides an explanation for the puzzle of low insurance demand among measurably risk averse individuals. To further test this model, and anticipating our experiment’s randomized variation in insurance premiums, we show that the price elasticity of demand for insurance is greater for individuals with low trust.

We consider an agent who has to decide whether or not to take indemnity insurance to protect himself against the risk that his wealth, w , is reduced by a fixed amount, c . The probability of this loss is p . Without insurance the agent’s welfare (expected utility) is given by

$$W = (1 - p)u(w) + pu(w - c). \quad (1)$$

The agent is risk averse so the utility function u is strictly concave.

Under insurance the agent pays a premium, π . If the loss occurs the insurer pays full compensation with probability q and otherwise defaults, paying nothing. This probability q is the subjective probability that a loss is covered—this means that it may be that the insurer never defaults but the agent is unclear on what is covered by the contract. With insurance the agent’s expected utility is therefore

$$\begin{aligned} \widetilde{W} &= (1 - p)u(w - \pi) + p[qu(w - \pi) + (1 - q)u(w - \pi - c)] \\ &= (1 - \tilde{p})u(w - \pi) + \tilde{p}u(w - \pi - c) \end{aligned} \quad (2)$$

where $\tilde{p} = p(1 - q)$. The probabilities in this compound lottery satisfy

$$0 < p < 1, 0 < q \leq 1.$$


The agent will accept the insurance contract if $\widetilde{W} > W$.

The probability q is a measure of the agent’s trust in the insurer. Under complete trust ($q = 1$) and actuarially fair insurance ($\pi = pc$) the probability \tilde{p} equals 0 and

$$\widetilde{W} = u(w - \pi) = u(w - pc) > (1 - p)u(w) + pu(w - c) = W$$

by Jensen’s inequality and the concavity of the utility function. This is, of course, the standard result that under full trust a risk averse agent will prefer insurance. Insurance raises the outcome in the bad case (from $w - c$ to $w - pc$) and reduces the outcome in the good case (from w to $w - pc$).

Figure 1: q^* locus as a function of price and the coefficient of relative risk aversion R



Note: The Figure plots the q^* locus for a premium with 0% (top), 10% (middle) or 25% (bottom) subsidy.

Since the premium is actuarially fair this amounts to the opposite of a mean preserving spread and is therefore obviously attractive to a risk averse agent.

Limited trust ($q < 1$) changes the attractiveness of insurance fundamentally. Insurance now reduces the probability of a loss (from p to \tilde{p}) but it makes the bad outcome worse: $w - \pi - c$ instead of $w - c$. It follows that a very risk averse agent may refuse an insurance contract which a less risk averse agent would accept.

The model is similar to that of Doherty and Schlesinger (1990). However, while they assume that the agent can choose the degree of insurance cover we rule out partial insurance: the loss c is either fully covered by insurance or not at all. In the context of health insurance in developing countries this specification is more realistic: insurance contracts (such as in the Kenyan *Bima ya Jamii* project studied here) typically offer indemnification for specific risks such as the cost of hospitalisation on an all-or-nothing basis. In this setting limited trust ($q < 1$) affects the decision to take up insurance whereas in the Doherty-Schlesinger model it affects the optimal insurance cover.³

Insurance will be accepted for $q > q^*$ where q^* solves $W(q) = \tilde{W}(q)$. It follows from (1) and (2) that

$$q^* = 1 - \frac{u(w - \pi) - [(1 - p)u(w) + pu(w - c)]}{p[u(w - \pi) - u(w - \pi - c)]}. \quad (3)$$

Figure 1 plots q^* as a function of the degree of relative risk aversion, R , for a numerical example with constant relative risk aversion (CRRA) and parameter values $p = 0.5$, $w = 100$ and $c = 50$. The plot is shown for various values of the premium $\pi = \delta pc$ where δ takes the values 1.0 (top

³Clarke (2011) uses a similar approach to consider the question why a rational agent might refuse index insurance. In this setting the key issue is basis risk: the index is imperfectly correlated with the agent's outcome variable (e.g. crop yield) so that he may get no compensation after suffering a loss or, conversely, get compensation when in fact he has not suffered a loss, making the demand for insurance with basis risk fundamentally different from the case of indemnity insurance (2011). Our case is asymmetric: while the agent may fail to be compensated for a loss he will not receive compensation in the absence of a loss. What is similar in the two models is that insurance may be rejected because it makes the worse outcome worse: under index insurance because of imperfect correlation, in our case of indemnity insurance because of limited trust.

curve), 0.9 (middle curve) or 0.75 (bottom curve).⁴ Note that for $\delta < 1$ the premium is subsidised. For $\delta = 1$ the premium is actuarially fair in the conventional sense ($\pi = pc$) but not in the sense of allowing for limited trust ($\pi = pqc$). While $\pi = pqc$ is obviously the relevant theoretical concept of actuarial fairness it would imply that the insurer lowers the premium to compensate for his clients' lack of trust in him; this would seem rather farfetched.

We assume that agents are heterogeneous in terms of risk aversion (R) and the subjective probability of default (q). It follows that an agent is characterised by (q, R) which defines a point in the Figure. Clearly, the agent will accept insurance only if that point lies above the locus.

Figure 1 shows that q^* can be non-monotonic in risk aversion: for the chosen parameter values q^* initially decreases with risk aversion, reaches a minimum and then increases. It is therefore possible that (for a given level of q) those with either very low or very high risk aversion accept insurance while those with an intermediate degree of risk aversion do not. This may explain the empirical 'puzzle' that more risk averse agents refuse a contract which less risk averse agents accept.

While the numerical example is instructive the result is more general. We show this for the class of CRRA utility functions. The key result is that for high values of risk aversion q^* increases with R . For low values of risk aversion it usually decreases with R but in special cases (a combination of a high subsidy and a high value of p) q^* increases with R even for very low values of R :

Proposition 1. *For CRRA utility functions and a premium $\pi = \delta pc$ the minimum trust level q^* is increasing in the degree of relative risk aversion R for high values of R . There exist values δ^* and p^* such that for low values of R (a) q^* increases with R for $0 < \delta < \delta^*$ and for $\delta^* < \delta \leq 1$ and $p > p^*$ and (b) q^* decreases with R for $\delta^* < \delta \leq 1$ and $0 < p < p^*$.*

Proof. The q^* locus is defined by $\widetilde{W} = W$:

$$(1 - \widetilde{p})u(w - \pi) + \widetilde{p}u(w - \pi - c) = (1 - p)u(w) + pu(w - c). \quad (4)$$

For the first part of the Proposition note that for large values of R the expected utilities on either side of this condition can be approximated by the second term (which describes the worse case) so that

$$\widetilde{p}u(w - \pi - c) \cong pu(w - c).$$

Using $\widetilde{p} = p(1 - q^*)$ this gives

$$q^* \cong 1 - \frac{pu(w - c)}{pu(w - \pi - c)} = 1 - \left(\frac{w - c}{w - \pi - c} \right)^{1-R}.$$

and hence

$$\frac{dq^*}{dR} > 0.$$

For the second part of the proposition consider values of risk aversion such that $R \neq 1$ and assume⁵

$$\ln(w - \pi - c) > 0.$$

Differentiating (4) using $\widetilde{p} = p(1 - q^*)$ and

$$\frac{dx^{1-R}}{dR} = -x^{1-R} \ln x = -\psi(x)$$

⁴Since $\pi = \delta pc$ it follows from (3) that when the utility function is linear ($R = 0$) then $q^* = \delta$.

⁵This technical assumption ensures that the function $\psi(x)$ defined below is positive and increasing for all values of x considered.

gives

$$\frac{dq^*}{dR} = \frac{A - B}{C} \quad (5)$$

where

$$\begin{aligned} A &= (1 - \tilde{p})\psi(w - \pi) + \tilde{p}\psi(w - \pi - c), \\ B &= (1 - p)\psi(w) + p\psi(w - c), \\ C &= p[(w - \pi)^{1-R} - (w - \pi - c)^{1-R}]. \end{aligned}$$

Hence

$$\frac{dq^*}{dR} > 0 \text{ iff } \begin{cases} A > B & \text{for } 0 \leq R < 1 \\ A < B & \text{for } R > 1 \end{cases}. \quad (6)$$

Note that B is linear and decreasing in p , that A is decreasing and convex in p , and that for $p = 0$ we have $A = B$ and

$$\frac{dA}{dp} < \frac{dB}{dp} < 0.$$

Therefore two cases can arise, depending on the values of p and δ : (i) for $0 < \delta < \delta^*$ we have $A < B$ for all p so that

$$\frac{dq^*}{dR} < 0$$

provided $0 \leq R < 1$; (ii) for $\delta < \delta^* \leq 1$ there is a value $0 < p^* < 1$ for which $A = B$ so that

$$\frac{dq^*}{dR} \begin{cases} < 0 & \text{for } 0 < p < p^* \\ > 0 & \text{for } p > p^* \end{cases}$$

provided, again, that $0 \leq R < 1$. The critical value δ^* solves $A = B$ for $p = 1$:

$$q^*\psi(w - \delta c) + (1 - q^*)\psi(w - \delta c - c) = \psi(w - c).$$

□

A subsidy shifts the q^* -locus downwards so that (for a given distribution of agents in (R, q) space) more agents will accept insurance. In particular, a risk neutral agent will now strictly prefer insurance at $q = 1$ because of the subsidy element. Note from Figure 1 that the minimum shifts to the left: the larger the subsidy the lower the degree of risk aversion beyond which q^* is increasing in risk aversion. For extreme values of p and δ the minimum does not occur for $R > 0$: in that case the locus is monotonically increasing in R .

The key part of the Proposition is that q^* is increasing in R for large R so that very risk averse individuals reject the insurance contract. We extend this to the effect of R on $\Delta \equiv \tilde{W} - W$, the difference in expected utility between the insured and uninsured states. We do so on the grounds that it is a desirable property of a stochastic choice model that the probability of becoming insured should be increasing in this expected utility differential.⁶ Our key result carries over to the effect of risk aversion on the expected utility differential, the effect is negative for large R :

Proposition 2. *For large values of R the expected utility differential Δ is decreasing in R .*

⁶As Wilcox (2008, 2011) notes, when changes in the expected utility differential arise from changes in preference parameters, rather than the (possibly subjective) probabilities and payoffs associated with events, then the consequent rescaling of the utility measure cannot be considered independently of the stochastic component of choice. In the present context the scaling proposed by Wilcox would require Proposition 2 to hold for the effect of R on $\Delta/[u(w) - u(w - \pi - c)]$ rather than on Δ . (Cf. equation(7) in Appendix B.) We show in Appendix D that this is the case.

Proof. As before we approximate expressions involving terms in x^{1-R} by using only the terms with the lowest values for x . Differentiating Δ with respect to R then gives:

$$\begin{aligned} \frac{d\Delta}{dR} &< \tilde{p}u(w - \pi - c) \left(\frac{1}{1-R} - \ln(w - \pi - c) \right) \\ &\quad - pu(w - c) \left(\frac{1}{1-R} - \ln(w - c) \right). \end{aligned}$$

Since $\tilde{p} < p$ and $u(w - \pi - c) < u(w - c) < 0$ (provided $R > 1$) a sufficient condition for the right hand side to be negative is

$$\frac{1}{1-R} - \ln(w - \pi - c) > \frac{1}{1-R} - \ln(w - c)$$

and this is true since $\pi > 0$. □

Anticipating the exogenous premium variation in our experiment, we now derive further predictions from the model for the interaction of premiums with lab-experimental measures of trust.

The expected utility differential is decreasing in the price of insurance, trivially; this is what generates a downward-sloping demand curve. We show that strict concavity of the utility function also implies that this expected utility differential is decreasing in price *more strongly* for individuals who have low trust in the insurer (low q).

Individuals will hold subjective beliefs about the credibility of a particular insurance policy. The trust parameter, q , is a composite of several factors, among them: the likelihood of the hospital agreeing to accept the insurance policy; the likelihood of the insurer continuing to be in business and agreeing to pay a claim; and—if the individual is required to make a cash payment at the time of hospital admission—the likelihood of reimbursement actually reaching the individual.⁷ Objective values of q are therefore likely to vary across individuals, who may have variable success in using the policy. Subjective beliefs about one's own value of q may introduce a further element of subjectivity, as they will depend (among other things) on trust in particular individuals and institutions. Proposition 3 shows that the effect of price on the expected utility gain from adopting insurance is greater among individuals with low trust.

Proposition 3. *Let the expected utility differential from insurance adoption be given by Δ , as defined above, and assume that individuals have strictly concave utility, defined over their net wealth. Then, trivially, $\partial\Delta/\partial\pi < 0$ and $\partial\Delta/\partial q > 0$. Moreover, $\partial^2\Delta/\partial\pi\partial q > 0$.*

Proof. Differentiation of Δ yields

$$\frac{\partial^2\Delta}{\partial\pi\partial q} = p \left(u'(w - \pi - c) - u'(w - \pi) \right),$$

where $u'(\cdot)$ denotes the first derivative of the utility function. By the strict concavity of $u(\cdot)$, this is strictly positive. □

These propositions yield a set of three empirical implications, which we test in our data combining lab-experimental measures of preferences with field-experimental variation in premium costs.⁸

⁷While the *de jure* policy is that no up-front payments should be made by Bima ya Jamii policyholders, individuals were in some cases required by hospitals to make such payments in the early stages of implementation (prior to the present study).

⁸Our model also implies that $\partial\Delta^2/\partial\pi\partial R > 0$; however, since this is also true in a model with complete trust ($q = 1$), this cannot be used to test the theory.

Prediction 1. When trust is incomplete ($q < 1$), the probability of insurance purchase is either decreasing in risk aversion throughout or is so for sufficiently high risk aversion (Proposition 2).

Prediction 2. At a given level of risk aversion, potential clients’ trust in the insurer has a positive effect on the probability that insurance is purchased (Proposition 3).

Prediction 3. The probability of insurance purchase is more responsive to price for potential clients with low trust in the insurer (Proposition 3).

The following section introduces the setting, experimental design, and data sources used to test these hypotheses.

3 Experimental design and baseline data

We take these testable implications to data from a field experiment conducted in Nyeri District, Kenya. The experiment offered a composite health insurance policy, *Bima ya Jamii*, to tea farmers belonging to the Wananchi Savings and Credit Cooperative Society. The field experiment was a factorial design, in which farmer-level variation in premium costs of the policy was cross-cut with cluster-randomized marketing and learning interventions.

Bima ya Jamii is a composite health insurance product offered by the Cooperative Insurance Company (CIC) of Kenya. This product bundles in-patient hospitalization cover for household members, provided by the National Hospital Insurance Fund (NHIF) to all public-sector employees, with cover for lost work during hospital stays and funeral insurance. There are no exclusions on the basis of prior conditions. At the time of the study, the cost of the policy was KShs 3,650 per year (roughly USD 50 using exchange rates at that time).

CIC typically partners with local financial institutions, who act as intermediaries in the delivery of the product. Our study focuses exclusively on their work with Wananchi Savings and Credit Cooperative Society, a cooperative comprised primarily of tea farmers in Nyeri District, Central Province. All Wananchi members hold bank accounts with this SACCO, and payments for their tea harvest is made through these accounts by the Kenya Tea Development Agency. In addition, Wananchi offers a range of loans to its members, but participation in these loans is fairly limited.

Wananchi’s members are divided into 162 tea-collection centres, which are grouped in 12 administrative zones. 120 of these centres form the basis for the cluster-randomized assignment to the treatment arms described below. In each centre, we randomly selected 9 farmers at random from Wananchi’s membership roll, together with the elected ‘delegate’ who represents the members in the co-op’s meetings, for inclusion in our baseline study. We analyze the decision to purchase insurance among this sample.

3.1 Field experiment

Against a backdrop of limited demand, the field experiment piloted and evaluated interventions designed in consultation with policymakers to address perceived constraints to insurance uptake: price, access, and knowledge. A basic marketing campaign was proposed and designed by CIC, while an NGO, the Swedish Cooperative Society (SCC), offered a more in-depth training in financial literacy, risk management and insurance (without ever mentioning the CIC product). Further, a persistent concern that costs may still be too high for poor farmers led to a pilot in an area

where poverty was moderate, with flexible payment terms and experimentally controlled variation in premium costs.⁹

To test the reduced-form impacts of these treatments and their interactions, we used a factorial design. First, we randomly assigned tea centres either to control or to one of two cluster-level treatment arms: basic marketing or marketing preceded by a financial literacy intervention. Second, we randomly assigned some individuals in each of the two treatment arms to receive vouchers that would reduce the premium cost of the policy, as described below and in Table 1. The entire implementation of the sales of the policies was done via the SACCO, as a trusted intermediary.¹⁰

Sixty of the tea centres in the study were allocated to a control treatment arm, given the interest of the project in studying the impact of the insurance product on health and economic outcomes. While all Wananchi members were technically eligible to purchase insurance, members of these control centres received no direct information about the product from Wananchi staff, and received no price discounts. In practice no policies were bought in these centres, and they are excluded from the analysis in the remainder of this paper.

The remaining 60 tea centres all attended a meeting in which basic information about the *Bima ya Jamii* product was provided by CIC marketing agents, who were accompanied by a representative from Wananchi. These meetings lasted between one and two hours. We refer to the 30 centres that attended these meetings but did not receive the educational treatment described below as receiving the *marketing only* treatment.

In our *study circles* treatment, the remaining 30 centres received education in financial literacy, with a focus on insurance. The ‘study circles’ modality used to deliver this educational training is a system practiced in the dissemination of agricultural technologies and other contexts by the Swedish Cooperative Center (SCC), an international NGO that administered this treatment. Its basic idea is to train someone in the community—in this case, the Wananchi Delegate—to lead regular study groups, in which they discuss written materials together with a small group of their peers. SCC developed the curriculum for these study circles together with Microfinance Opportunities, an NGO with extensive experience in financial literacy training. The topics covered were general, in that the *Bima ya Jamii* product was not mentioned by name, though the focus was primarily on indemnity insurance and health-related shocks. The resulting course consisted of 10 modules, which were undertaken on a weekly basis prior to the launch of the basic marketing treatment. In order to better position the study to capture any potential impact of this treatment, delegates were instructed to include the 9 other sample members in their centre in the first of the study groups they conducted, although they were also encouraged to repeat this curriculum with other members of their centre.

At the individual level, Wananchi members outside of the control centres were randomly allocated vouchers that reduced the costs of the Wananchi premium by values equivalent to 0, 10, or 20 percent of the original cost. These vouchers were drawn with equal probability during a

⁹Following CIC’s interest in piloting an alternative marketing channels, a further set of 30 tea-collection centres were allocated to a parallel scheme that offered payouts to purchasers of insurance who encouraged a second generation of adopters. This compromise scheme was the result of discussions between the research team and CIC, who were interested in exploring alternative marketing channels uncomfortably close to a pyramid scheme. The resulting intervention retained the feature that it was marketed not just as an insurance product but as a basis for financial returns. It suffered from low take-up, and for reasons related to the government’s reform of NHIF coverage, the sign-up window closed before members had a sufficient opportunity to engage in peer referrals. For these reasons, we exclude this arm in its entirety from the current study.

¹⁰As neither CIC, NIHF nor SCC were known in the community, we had to work via the local SACCO; while it may have been useful to isolate the role of the SACCO, for example to investigate trust in this institution, there would have been no viable payment system for premiums, so we did not consider this further. We randomly built this issue into the laboratory experiments, discussed below.

Table 1: Experimental design and sample sizes by treatment assignment

Centre-level treatment	Individual premium vouchers		
	No discount	10% discount	20% discount
Control (60)	597	0	0
Marketing only (30)	105	90	102
Marketing + study circles (30)	108	91	100

Notes: Table displays number of survey respondents, by centre-level treatment arm and discount voucher received. Number of tea centres assigned to each centre-level treatment reported in parentheses.

public lottery conducted during the marketing session common to all treatment arms.¹¹ Since not only sampled households but all Wananchi members in treatment centres were invited to these marketing sessions, even members who were not included in the baseline survey were eligible to participate in this lottery. The resulting distribution of vouchers across individuals in the baseline survey is shown, broken down by treatment arm, in Table 1.

3.2 Laboratory experiment

To test the empirical implications of the theory in Section 2, we combine exogenous field-experimental variation in prices with incentive-compatible measures of risk preferences and trusting behavior derived from a lab-type experiment conducted in the field at baseline—an *artefactual field experiment*, in the taxonomy of Harrison and List (2004). These experimental games were undertaken with the same sample of tea farmers who participated in the baseline survey. Games were played sequentially, in randomized order, with payoffs revealed after each game but not made until the end. We played variants of two standard games, a Holt and Laury *gamble-choice game* (Holt and Laury 2002, henceforth HL) to measure risk preferences, and a Berg, Dickhaut, and McCabe *trust game* (Berg et al. 1995, henceforth BDM) to measure trusting behavior.

We measure trust with a variant of BDM’s trust game (1995), played by Wananchi members with a peer from their tea centre. We interpret behavior in this game as determined, in part, by a generalized perception of the trustworthiness of others, which is likely to—and indeed, which we demonstrate does—carry over into insurance purchase decisions.¹² Such lab-type measures of trusting have been shown elsewhere to correlate with decisions in naturally occurring contexts (Camerer forthcoming). Barr and Serneels (2009) show that levels of trust among employees of Ghanaian manufacturing firms are associated with firm productivity. Closer to our application, Karlan (2005) shows that *trustworthiness* in this game correlates with microcredit repayment rates in Peru; our focus on *trusting* behavior is a natural reflection of the fact that the burden of trust in microinsurance is the mirror image of microcredit: in the latter case, it is the microinsurance client who is asked to place their trust in the insurer.

We adapt the specific implementation of the *trust game* from the design used in Zimbabwe by Barr (2003). The basic setup is as follows (further details of the protocol are provided in Appendix A). Subjects are assigned to one of two roles, Sender or Receiver. Both are endowed with KShs

¹¹Attendance at these marketing sessions by our sample participants was not universal. In order to ensure that the probability of receiving a voucher was uncorrelated with other possible determinants of insurance demand, we randomly assigned vouchers with the same probabilities to individuals who did not attend this marketing session. Delegates visited all sample members who did not attend the marketing session to notify them of the product and to deliver any non-zero vouchers.

¹²While it would have been ideal to measure trust in the insurer directly, it was not possible to do so in an incentive-compatible way without introducing other confounds.

200 at the outset of the game. The Sender can then decide to send a portion of their endowment to the Receiver (from zero to KShs 200, in increments of KShs 50). Any amount that is sent to the Receiver is tripled. The Receiver can then decide to return any portion of this tripled amount—possibly none—to the Sender, at which point the game concludes.

For the empirical analysis, we categorize individuals as exhibiting *low trust* if they invest less than 50 percent of the stake.¹³ Levels of trusting behavior in our study are somewhat higher than those observed in other contexts. Senders in our treatment locations send an average of 65 percent of their endowment to the Receiver, slightly higher than the 50 percent investment rate typical of Camerer’s (2003) survey. The modal Sender invests the full stake (just under 40 percent of respondents), which we note is in comparable with the substantial fraction of individuals investing the full stake in the cross-country study of Ashraf et al. (2006).¹⁴

A large literature explores the determinants of trusting behavior in the BDM trust game, which may be influenced not only by subjective perceptions of trustworthiness, but also by altruism and risk preferences (Barr 2003, Eckel and Wilson 2004, Ashraf, Bohnet and Piankov 2006). It should be noted that this is in some sense an inevitable feature of any incentive-compatible measure of trust, which requires strategic interactions relying on expectations of the uncertain trustworthiness of others. However, as discussed in Section 4.2, there are three reasons to believe that the data support an interpretation in terms of trust. First, we separately undertake and control for lab-experimental measures of risk-taking behavior. Second, we show that our measure of trusting behavior is uncorrelated with available measures of risk preferences. Third, the predictions that we take to the data are not plausibly explained by an altruism confound. And fourth, an interpretation of trust-game play as a proxy for risk preferences can only explain the full pattern of empirical results if one assumes that trust in the insurer is limited—thereby providing alternative confirmation for the model that we seek to test.

Measures of risk preferences will be used both as a test of the theory’s implications in their own right, and as a control for their possible confounding role in the interpretation of Sender behavior in the trust game as a measure of trust. We measure Wananchi members’ risk preferences with a Holt and Laury (2002) gamble-choice game, adapting the specific design of Barr (2007) to the Kenyan context. This game consists of a series of tasks, in each of which the subject chooses between two binary lotteries, one ‘safe’ lottery and one ‘risky’ lottery. Each lottery consists of a high-payoff outcome and a low-payoff outcome, which are held constant across tasks, while the probability of winning changes. Each subject’s decisions across these tasks is combined to a scalar measure of their risk aversion.

We played two series of this game, a gain-frame series and a loss-frame series. In the gain-frame series, subjects began with an initial endowment of zero, and had an opportunity to win either KShs 300 or KShs 0, if they chose the risky lottery, or KShs 100 or KShs 50, if they chose the safe lottery.¹⁵ Probabilities of winning ranged from 80 percent to 30 percent over the six tasks. In the loss-frame series, subjects were endowed with KShs 300 prior to play, with lottery outcomes framed as losses leading to the same reduced-form distribution of payoffs. The two series were played sequentially, with payoffs determined after both series were complete, based on a single task selected at random from across the two series.

From subjects’ choices in these series of tasks, we define several measures of their levels of risk aversion. To facilitate interpretation, since a substantial fraction of subjects exhibit multiple

¹³Reduced-form results are qualitatively unaffected when we use a linear function of the share contributed; the binary specification helps to keep the parameter space feasible for the structural estimates that follow.

¹⁴We discuss design specifics likely to contribute to this high investment rate in Appendix A.

¹⁵The prevailing exchange rate at the time of the laboratory experiment was KShs 75/USD, meaning that the maximum payout in this series is USD 4.

switching points (as is common in these games: see Hey 2002), we not only use the raw fraction of risky choices in the empirical analysis, but also follow Harrison and colleagues (2010) to estimate constant relative risk aversion (CRRA) parameters by maximum likelihood for each subject individually. Full details of this exercise are provided in Appendix B. To do so, we assume that preferences over (narrowly framed) outcomes in this lottery can be represented by a CRRA utility function of the form $u(x) = x^{1-R}/(1-R)$.

The mean value of R_{gain} , the CRRA coefficient in the gain-frame lottery, is 0.5 (standard deviation 0.19). Behavior in the loss-frame sequence is consistent with a greater degree of risk aversion; the mean estimated coefficient of relative risk aversion is 0.56 (0.19) in this series. This puts these estimates in the same range as those found in similar laboratory experiments in the field; for example, Harrison and coauthors (2010), in their EUT model assuming homogeneous preferences, estimate a population parameter of $R = 0.54$.

3.3 Survey data and balance

An extensive baseline household survey was collected among the sampled population, overlapping with those included in the laboratory games. In each tea centre, the study sampled nine farmers from the centre’s register, as well as the elected *delegate*, who serves as a liaison with the Wananchi SACCO.

Table 2 presents descriptive statistics for the population studied in this paper, to whom insurance was marketed, including both survey characteristics and participants’ behavior in the laboratory games. In column (1) of Table 1, we present means and standard deviations for this population. Columns (2) and (3) present regression coefficients and associated standard errors for a regression of the relevant characteristic on the proportional value of the premium discount received (either 0, 0.10, or 0.20 of the premium) and an indicator for whether the tea center received the financial literacy treatment.¹⁶ Consistent with the randomization, nearly all characteristics are found to be balanced, and point estimates are generally small in magnitude, such that the regression results of Section 4 are robust to inclusion of controls for all variables in there is evidence of chance imbalance.

Measured characteristics suggest a favorable population for an expansion of microinsurance. Households in our sample are poor, though not at the extremes of poverty. The tea farmers sampled are predominantly male, with average ages in their fifties, and household sizes between three and four individuals. They have some education, although only 39 percent of sampled farmers have more than primary education. Using prevailing exchange rates of KShs 75/USD from the time of the survey, mean per capita monthly consumption is approximately USD 154.

Beside their own resources, many respondents have access to both formal and informal insurance mechanisms. are members of informal support networks. On average, respondents report that they could turn to between 5 and 6 other friends or family members to help address a shock, and that they can borrow a total value of KShs 10,282 (USD 147) in such an event. Perhaps most surprisingly, a substantial fraction of households in the survey have purchased insurance in the past. Of the 36 percent of individuals who report their household having ever purchased in the past, more than half report having bought health insurance, and nearly all report that this cover remains in place at the time of the baseline. Although we suspect this may be overreported due to poor baseline levels of understanding of insurance, it should be noted that there is private provision of various forms of insurance in the study area.

Households in our survey have experienced medical expenses in the past year. Approximately 40 percent of households have experienced a non-zero medical expenditure, and average household

¹⁶So, for a given characteristic x , we report the coefficients (β_d, β_f) and associated standard errors from a regression of the form $x = \beta_0 + \beta_d \text{discount} + \beta_f \text{fin. literacy} + e$, with standard errors clustered at the tea-centre level.

Table 2: Survey characteristics, by discount voucher and financial literacy treatments

	Full sample	Discount	Fin. lit
$\mathbf{1}$ [primary respondent female]	0.33 (0.47)	-0.00 (0.00)	-0.00 (0.04)
age, primary respondent	56.40 (14.84)	-0.01 (0.02)	-0.06 (1.53)
HH size	3.40 (1.67)	-0.00** (0.00)	0.02 (0.18)
HHH post-primary education	0.39 (0.49)	0.00 (0.00)	-0.01 (0.05)
value of HH consumption, last month, KShs	39,334.88 (88,897.71)	-145.76 (104.11)	-7,844.71 (8,009.89)
value of HH assets, KShs	83,807.23 (148,580.37)	-30.82 (217.44)	4,421.11 (12,956.48)
informal network size	5.50 (14.08)	-0.01 (0.02)	-1.14 (1.08)
informal network value, KShs	10,282.12 (27,934.48)	26.85 (22.59)	-235.01 (2,325.26)
$\mathbf{1}$ [HH medical expenditure > 0]	0.40 (0.49)	-0.00 (0.00)	0.00 (0.04)
HH medical expenditure, past year, KShs	2,676.14 (17,291.71)	-14.35 (29.48)	-43.46 (1,436.23)
$\mathbf{1}$ [HH inpatient costs > 0]	0.10 (0.30)	-0.00 (0.00)	-0.03 (0.02)
HH inpatient costs, past year, KShs	1,117.57 (7,825.24)	4.86 (12.22)	-1,214.65** (597.41)
subjective Pr[hospital cost > 0]	0.47 (0.25)	-0.00 (0.00)	0.01 (0.03)
subjective E[hospital cost]	43,297.83 (126,472.60)	-74.92 (209.40)	-4,527.26 (9,580.47)
$\mathbf{1}$ [HH ever bought insurance]	0.36 (0.48)	-0.00 (0.00)	0.03 (0.05)
trust game: share sent	0.64 (0.32)	-0.00 (0.00)	0.02 (0.04)
low trust	0.52 (0.50)	0.00 (0.00)	-0.04 (0.07)
CRRA (gain frame)	0.50 (0.19)	0.00* (0.00)	-0.02 (0.02)
CRRA (loss frame)	0.56 (0.19)	-0.00 (0.00)	0.01 (0.02)

Notes: Column (1) presents means and standard deviations of each characteristic for the full sample of farmers. Columns (2) and (3) present regression coefficients and associated standard errors from a regression of the specified baseline characteristic on two measures of subsequent treatment assignments, as denoted by column headings. ‘Discount’ is defined as the fraction of the full premium (0, 0.10, or 0.20) offered as a discount. ‘Fin. lit.’ is an indicator for the study circles financial literacy treatment. Standard errors are clustered at the tea centre level. Asset values exclude land and business capital. All financial amounts in KShs.

medical expenditures in the past year were approximately KShs 2,676 (USD 38).¹⁷ 10 percent of the households had a member spending time in hospital for inpatient treatment. Mean household expenditures of KShs 1,118 (USD 16) imply that these episodes, when they occur, have direct costs of the same order as the monthly consumption of one household member. Looking forward over the next year, households expect to incur in-patient expenditures with 47 percent probability, and the (unconditional) expected value of these costs as reported by respondents is KShs 43,298 (USD 618).¹⁸ While the latter value is higher than would be expected based on past costs, these reflect genuine fears on the part of respondents regarding health risks, the consequences of which are perceived to be potentially beyond the reach of informal insurance mechanisms. In-patient hospitalization costs are particularly relevant to demand for *Bima ya Jamii*, which primarily covers in-patient medical costs. Historical and perceived hospitalization risks are high, and perceived values of the associated costs are large, though unconditional average realized in-patient costs are approximately one third of the premium of the product on offer.

To summarize, we are dealing with an area in which health insurance could thrive. The target farmers are not destitute and have a reasonable level of education. They have had some insurance exposure in the past. Finally, the target group experienced medical emergencies and expenses, and expect to experience them also in the future. Actual and likely hospitalization costs are relatively high, and they appear too high to be covered by informal insurance systems. Therefore, the target farmers may well value the opportunity to insure themselves against health shocks. The product on offer appears to be reasonably well (and close to actuarially fairly) priced given recent experiences with health costs and expectations for future costs.

4 Evidence from the reduced form

4.1 Tests of empirical predictions

Here we present the main empirical results of the study. We first show the reduced-form results of the field experiment, showing that demand is highly price-elastic but unresponsive to financial literacy training. We then proceed to the primary tests of our model, by combining this field-experimental variation with lab-experimental measures of preferences and beliefs. There we show that risk aversion and low trust are both negatively associated with insurance demand, and that the purchase decisions of individuals with low trust are significantly more sensitive to price.

We begin by presenting estimates of the reduced-form effect of our experimental treatments on demand. Given that the model to be estimated consists of a set of binary treatment indicators, we estimate a linear probability model, where the dependent variable is a binary indicator for insurance purchases. The results are presented in Table 3.

In the first column of this table, we present results for the basic effects of our treatment arms, without allowing for treatment interactions. Two results are notable here. The ‘study circles’ financial literacy intervention had no measurable effect on demand. This may of course be attributable to a failure of the usefulness or execution of this particular curriculum, but it accords with the mixed results of the literature on financial literacy training and insurance demand (Cole et al. 2011, Cole et al. 2013); this failure is notable in part because the intervention studied here engendered sustained engagement over a ten-module course. We also test, and fail to reject, the linearity of demand as a function of price. Premium discount vouchers have an economically and

¹⁷Figures for realized medical expenditures include inpatient, outpatient and traditional medicine.

¹⁸These subjective expectations were elicited following the approach of Manski (2004), as advocated by Delavande and coauthors (2009).

Table 3: Demand for insurance by experimental treatment

	(1)	(2)
voucher, 10%	0.0726** (0.03)	0.0622 (0.05)
voucher, 20%	0.112*** (0.04)	0.127** (0.06)
study circles	-0.0180 (0.04)	-0.0141 (0.05)
voucher, 10% \times study circles		0.0205 (0.07)
voucher, 20% \times study circles		-0.0296 (0.08)
Constant	0.129*** (0.03)	0.127*** (0.04)
Obs	623	623
H_1 : p-value	0.569	
H_2 : p-value		0.764

Notes: Linear probability model. Dependent variable = 1 if respondent completed application. Explanatory variables defined as indicators for treatment arms and their interactions. Robust standard errors, clustered by tea-collection center. Test statistics for hypotheses that (H_1) coefficient on voucher of 20 percent is twice coefficient on voucher of 10 percent; and (H_2) interaction effects are jointly insignificant.

statistically significant effect on demand: starting at full price, a ten percent reduction in price increases demand by approximately 7 percent.

The second column of Table 3 tests for interactions between price and financial literacy treatments. These interaction terms are statistically insignificant, with economically small point estimates. We comfortably fail to reject the hypothesis (labelled H_2) that the interactions are jointly insignificant, and on this basis we will focus on the average effect of prices and its interaction with lab-based measures of risk-taking and trusting behavior (while controlling for exposure to financial literacy treatments).

To test the model of Section 2, we now combine these field-experimental data with the measures of trusting and risk-taking behavior in the baseline lab experiment. Specifically, we test the three primary empirical implications of that model: the relationship between risk aversion and trust is either strictly decreasing or inverted-U shaped; insurance demand is lower for individuals with low trust; and insurance demand is more sensitive to price for individuals with low trust.

We report coefficients from a probit model of the decision to purchase insurance in Table 4, where the sample is defined as individuals in our treatment sample who played the ‘Sender’ role in the BDM trust game. Randomly assigned premium prices are expressed in units that correspond to shares of the full price. R_{gain} is the estimated coefficient of relative risk aversion from the gain-frame HL gamble-choice game. Cluster-randomized assignment to a tea center where financial literacy training was conducted via ‘study circles’ is denoted by the variable *study circles*. Controls for zone are included in all specifications, though their omission does not substantially alter the results.

Column (1) reports the basic findings of the experiment, introducing measures of risk preferences and trusting behavior. Consistent with empirical Predictions 1 and 2, demand is increasing in the amount of the discount, decreasing in measured risk aversion, and decreasing among individuals with low trust. This reproduces the basic stylized facts, observed across various contexts, that motivate our model (Cole et al. 2013, Cai et al. 2010).

In column (2) we test for a nonlinear association between risk aversion and demand, as suggested by Prediction 1. To do so we include the square of the measure of risk aversion, less its mean in the estimating sample. Although we cannot reject equality of the coefficient on this quadratic term with zero, the resulting point estimates are consistent with this prediction. For all weakly risk

Table 4: Risk, trust, and price in insurance demand

	(1)	(2)	(3)
price	-3.043*** (0.89)	-3.022*** (0.88)	-2.499** (1.02)
low trust	-0.425** (0.18)	-0.431** (0.19)	-0.856*** (0.23)
R_{gain}	-0.866** (0.38)	-0.950** (0.42)	-0.864** (0.38)
R_{gain}^2		-2.405 (2.10)	
price \times low trust			-3.444* (1.88)
study circles	-0.114 (0.17)	-0.124 (0.17)	-0.109 (0.17)
Observations	458	458	458

Notes: Probit coefficients reported. Dependent variable equals unity if respondent purchased Bima ya Jamii insurance policy. Robust standard errors reported, clustered at tea-centre level. Controls for individual characteristics include logs of household asset values, household size, and respondent age, as well as indicators for the gender of the respondent and whether any household member has post-primary education.

averse individuals, small increases in their levels of risk aversion are associated with decreases in insurance demand.¹⁹

Column (3) tests Prediction 3, that price variation should have a stronger effect on those who hold low values of q , as proxied by low-trust behavior in the trust game. This is supported in the data. Estimated coefficients imply that trust is associated with substantial differences in the marginal effect of price on the probability of insurance purchase. For a high-trust individual with the characteristics of the mean individual in the sample, the marginal effect of a change in price is -0.76 (with standard error 0.36), starting from a base value of 0.8 times the full price. This implies that an increase to 0.9 times the full price would reduce demand by 7.6 percentage points. By contrast, the estimated marginal effect for a low-trust individual is much larger, at -1.72 (0.70). Not only is the low-trust individual less likely to purchase insurance at this initial price, but a further increase in price to 0.9 times the full price is estimated to cause a 17 percentage point reduction in the probability of insurance purchase.

To summarize, we find empirical results that are broadly consistent with the model outlined in Section 2. Demand is increasing in the laboratory measure of trust and decreasing in the laboratory measure of risk aversion. In line with Proposition 1, we find modest evidence of non-monotonicity of demand in the measure of risk aversion.

¹⁹At the floor value of 0.22 at which our risk aversion measures are censored—corresponding to cases in which individuals chose the risky option for all gambles in the HL series—the implied marginal effect of R_{gain} is -0.57 (standard error 0.347), and the predicted probability of insurance purchase (using mean values of other variables) is 0.14. At the ceiling value for this risk aversion measure of 0.82, the predicted probability insurance demand is substantially lower, at 0.00081, and the corresponding marginal effect of R_{gain} is -0.05 (standard error 0.14).

4.2 Robustness and alternative hypotheses

Behavior in the trust game is interpreted above as proxying for an element of trust in the insurance product. An important concern, as mentioned above and widely discussed in the literature, is that trusting behavior in this game may depend not only on beliefs about the trustworthiness of others, but may also be a function of altruism and/or risk aversion (Ben-Ner and Putterman 2001).

Here, after reviewing the basis in the literature for concerns that risk preferences may confound the interpretation of our empirical tests, we present three arguments in support of the interpretation put forward in our model.²⁰ First, we show that risk aversion explains little of our trust measure. Second, we show that the trust-price interactions of Table 4 are robust to alternative measures and functional forms of risk attitudes. And third, we point out that a risk-aversion confound can only explain our empirical results in a model in which trust is limited ($q < 1$): if trust does not constrain insurance demand, then measures of risk aversion should be *positively*, rather than negatively, associated with purchase decisions. Thus one can only argue that the risk confound drives our empirical results by accepting the premise that the theoretical model of Section 2 is relevant, and trust is a barrier to adoption.

Across a number of studies, empirical support for the role of risk aversion in trust-game behavior is mixed. Eckel and Wilson (2004) fail to find a correlation between trusting and a range of survey-based and incentivized of risk measures. Karlan (2005) interprets relatively high microfinance default rates among high-trust individuals as evidence that they are more prone to take risks. Ashraf and coauthors (2006) are unable to detect a statistically significant relationship between sender decisions in a trust game and decisions in a gamble-choice game. Given that measured risk attitudes explain very little of the variation in trusting behavior they observe, they show that expected trustworthiness is quantitatively most important in determining trusting behavior. Schechter (2007) uses a measure of risk attitudes derived from a risky investment game explicitly designed to mimic the structure of the trust game, and finds decisions in this risk game to significantly predict trusting behavior for men, but not for women. Schechter argues for the importance of controlling for risk attitudes when interpreting trust game decisions as trust.

Table 5 confirms that measures of risk aversion explain remarkably little of trusting behavior in our lab experiments—particularly notable since our measure of risk aversion *does* explain insurance purchase decisions. There, we show means and standard deviations for each measure of risk aversion, for samples broken down by trusting behavior. Risk aversion, as derived from both the gain- and loss-framed gamble-choice games, differs by less than 0.02 across trust levels.

Table 5: Are trusting and risk aversion correlated?

	Low trust	High trust	<i>p</i> -value
R_{gain}	0.48 (0.19)	0.50 (0.20)	0.16
R_{loss}	0.55 (0.19)	0.56 (0.19)	0.87

Notes: Table reports means and standard deviations for alternative measures of behavior in Holt and Laury gamble-choice game, by level of trust. R_{gain}, R_{loss} give fitted coefficient of relative risk aversion from gain- and loss-frame gamble-choice tasks, respectively. *p*-values from test of equality in means, with standard errors clustered by experimental session.

²⁰We focus on risk preferences due to the lack of a plausible theoretical model under which an altruism confound, in which low trusting behavior would proxy for low altruism, would reproduce our empirical predictions.

Table 6: Probit coefficients and marginal effects, with alternative controls for risk and risk-price interactions

	R_{gain}	R_{loss}	F_{gain}
<i>Probit coefficients</i>			
price	0.24 (1.78)	-2.53 (1.79)	-0.76 (1.58)
low trust	-1.10*** (0.28)	-1.05*** (0.26)	-1.05*** (0.28)
price \times low trust	-5.19*** (2.01)	-4.36** (1.84)	-4.35** (2.05)
<i>Marginal effects</i>			
high trust: $\partial E[Y]/\partial \pi$	0.06 (0.43)	-0.65 (0.46)	-0.18 (0.39)
low trust: $\partial E[Y]/\partial \pi$	-0.67** (0.30)	-0.91*** (0.34)	-0.67** (0.32)
Observations	453	453	453

Notes: The dependent variable is an indicator equal to one if the individual purchased insurance. Probit coefficients and standard errors are reported in the first three rows. The marginal effects of a change in price, π , on the probability of insurance purchase are reported in the fourth and fifth rows, for a high- and low-trust individual, respectively. The marginal effects are computed for a discount level of 10%, at the mean of the remaining characteristics. Robust standard errors are shown in parentheses, clustered at tea-center level. All specifications include controls for zones, marketing treatment, and individual characteristics as in Table 4. Each column controls for a fourth-order polynomial in a measure of risk aversion, and its interaction with price. These are the fitted coefficient of relative risk aversion from the gain-frame (column 1) and loss-frame (column 2) HL series, and the fraction of safe lotteries chosen in the gain-frame HL series (column 3).

Further, the results of the preceding section are robust to inclusion of a variety of measures of risk aversion. Recall that column (4) of Table 4 showed that the point estimate and statistical significance of both trust and the trust-price interaction are robust to controls for the level of risk aversion and its interaction with price. In Table 6, we show that this remains the case for a wider array of measures of risk aversion, as derived from the gamble-choice games. We employ flexible functional forms for these measures and their interactions with price: in each column of Table 6, we control for a fourth-order polynomial in a measure of risk aversion, R , and its interaction with price. Columns (1)–(3) employ as measures of risk aversion, respectively, the fitted coefficient of relative risk aversion from the gain-frame lottery task; the fitted coefficient of relative risk aversion from the loss-frame task; and the fraction of safe lotteries chosen in the gain-frame task. The interaction term in the probit model is always statistically significant at the five percent level. We separately compute marginal effects of a price change for low- and high-trust types, for mean characteristics and a base price corresponding to a ten percent discount, to illustrate the greater price elasticity of low-trust individuals.

These results do not fully preclude the possibility that, since trust is measured in an incentive-compatible lab experiment, rather than experimentally manipulated, the theoretical object for which we interpret this as a proxy (trust parameter q) may be confounded by risk preferences. However, this alternative explanations cannot explain the full set of empirical results. Suppose, for example, that individuals categorized here as ‘low trust’ were in fact just particularly risk averse.

Then, the results of Table 4, Column (1) would show that two distinct measures of risk aversion are negatively associated with insurance purchase decisions. But in a world of perfect trust, risk aversion would be positively associated with insurance purchases. Consequently, a risk confound can only yield this pattern of results if we accept the premise that limited trust is an empirically relevant constraint to insurance adoption.

Taken together, these findings provide support for the view that the observed trust-price interaction is unlikely to be driven by confounding risk attitudes. Even if it this were the case, the implied re-interpretation of the empirical results would still support the model of Section 2.

5 Structural estimates and welfare implications

The reduced-form results presented thus far provide evidence that limited and heterogeneous trust levels discourage insurance purchases. But in order for this to be informative for policy purposes, there must be scope for substantive improvement in trust. If trust levels are economically low, then the welfare costs associated with foregone insurance transactions can be weighed against the potential costs of policy interventions that would strengthen consumers' faith in insurer payouts.

To address these issues, we estimate a structural model of consumers' demand for insurance. This enables us to calibrate levels of trust and risk aversion that are consistent with consumers' purchase decisions, beliefs about the distribution of potential hospitalization costs, and baseline consumption levels.²¹ Experimentally induced variation in premium costs helps to identify the model, while we allow for heterogeneous 'types' with respect to trust, according to the lab-based measure of trusting behavior used in the preceding section. Behavior is consistent with quite low subjective beliefs about the probability of insurer payouts conditional on hospital events: we estimate that 'high trust' types believe payouts will occur with 47 percent probability, while the behavior of 'low trust' types is consistent with a belief that payouts occur a mere 24 percent of the time. This suggests substantial scope for policies to reinforce trust in health insurance in Kenya.

These structural estimates also allow us to calculate welfare losses relative to alternative, counterfactual trust levels. As Chetty and Finkelstein (2013) have argued, one virtue of such a structural approach is that—by contrast with a 'sufficient statistics' approach to welfare analysis—it allows the calculation of welfare losses relative to contracts not offered in the observed equilibrium. In this sense, our approach differs from the seminal work of Akerlof (1970) and from recent empirical work by Einav and coauthors (Einav, Finkelstein and Cullen 2010, Einav and Finkelstein 2011), since the welfare losses considered here are relative contracts that are (perceived to be) fundamentally different than those observed in the market. This is crucial to the question considered here: limited trust changes the characteristics of a given insurance product, so that consumers' willingness to pay *at low trust* is not a sufficient statistic for the welfare loss associated with potential improvements in trust.

Here, we consider losses relative to a pair of polar assumptions about the *true* trustworthiness of the insurance product considered here. First, we consider the possibility that the insurance product is in fact fully trustworthy (in the notation of our theoretical model, $q = 1$), so that consumers' limited trust represents an overly pessimistic belief about this particular product. This approach is similar in spirit to Spinnewijn (2014) and Handel and Kolstad (forthcoming), who consider information frictions that drive a wedge between consumers' revealed preferences and the welfare gains that they would experience from insurance purchases. In this case, the welfare losses from improvements in trust are represented by the willingness to pay for a fully credible product, among

²¹Other recent work has shown the usefulness of stated beliefs in estimating welfare losses associated with adverse selection in insurance markets: see Hendren (2013).

the subset of individuals who would change their purchase decisions if they held full trust. Second, we consider the possibility that trust levels reflect the true, current (and therefore consumer-specific) trustworthiness of the insurer. In such a case, welfare losses relative to full trustworthiness also include the *inframarginal* losses arising from the poor quality of the product experienced by those who purchased insurance even under the status quo. Our results suggest that these welfare losses are substantial.

5.1 How much (mis)trust is there? Structural estimates of risk preferences and trust

We estimate the trust levels and coefficient of relative risk aversion consistent with health insurance decisions by maximum likelihood. To do so, we build on the theoretical model of Section 2 as follows.

We specialize to a constant relative risk aversion utility function over consumption outcomes, with the form $u(x) = \frac{x^{1-R}}{1-R}$ for household consumption level x . We take households' baseline consumption levels²² to be their (annualized) consumption levels in the survey data; utility is defined over consumption net of any insurance premiums paid and hospitalization expenses incurred. Because some individuals report a distribution of hospital costs that would leave them with negative consumption in some states of the world, we choose to define a minimum level of consumption as the lowest household consumption observed in our sample. This may be motivated as a limited liability constraint: hospitals may be unable to collect from patients beyond the point at which doing so threatens their survival.

Our estimation procedure uses survey data on consumers' subjective probability distributions for health costs. In the simplified version of the theoretical model presented in Section 2, consumers' health costs were modeled as taking on a value of zero, if they did not go to the hospital, or a known constant, c , if they suffered the event. Individual-specific variation in the distribution of health costs helps to identify the parameters of the model. In our survey, respondents, i , were first asked the probability that their household would experience a hospitalization in the next 12 months, p_i . They were then asked, *conditional on having non-zero hospitalization costs*: what was the lowest possible hospitalization cost that their household might incur, a_i ; (b) what was the highest possible hospitalization cost that their household might incur, b_i ; and what was the probability of a hospitalization cost above the average of these two points. Defining $F_i(c; a_i, b_i)$ as the individual-specific CDF for the distribution of hospital costs conditional on hospital costs strictly greater than zero, the answer to the latter question may be written as $1 - F_i((a_i + b_i)/2; a_i, b_i)$.

Following the empirical tests of the model in Section 4, we allow for decision-makers to vary in their levels of trust—their subjective belief, q_i , about the probability of an insurance payout, conditional on incurring a hospitalization cost. Using data from the baseline laboratory experiment, we estimate distinct values $q_{\text{high}}, q_{\text{low}}$ for individuals with high- and low trust in the laboratory trust game, respectively.

To map individuals' subjective expected utilities into probabilities of insurance purchase, we assume choices follow a logit choice rule:

$$\Pr[D_i = 1 | w_i, \pi_i, w_i, p_i, q_i, F_i(\cdot)] = \frac{\exp\{\theta(\mathbb{E}_i[U_{1i}] - \mathbb{E}_i[U_{0i}])\}}{1 + \exp\{\theta(\mathbb{E}_i[U_{1i}] - \mathbb{E}_i[U_{0i}])\}}. \quad (7)$$

Here, θ , a parameter to be estimated, captures the extent of determinism in individuals' choices. $\mathbb{E}_i[U_{0i}]$ denotes individual i 's subjective expected utility in the absence of insurance, where the subscript in the expectations operator implies that the expectation is taken over the subjective

²²This corresponds to w in the model of Section 2.

Table 7: Structural estimates of utility function parameters

	Estimate	Std. Err.
R	1.82	(0.031)
q_{high}	0.47	(0.001)
q_{low}	0.24	(0.001)
θ	1.64	(0.005)

Notes: Standard errors calculated by nonparametric block bootstrap, with blocks drawn at tea centre level.

probability of hospitalization, p_i , and distribution of hospital costs conditional on such an event, $F_i(\cdot)$. Defined analogously, $E_i[U_{1i}]$ denotes individual i 's subjective expected utility in the presence of insurance, which now depends not only on beliefs about hospitalization costs, but also on the subjective trust parameter, q_i .

Putting these building blocks together yields the estimated parameters reported in Table 7. Our benchmark model assumes a common level of relative risk aversion for all members of our sample; we estimate this parameter as 1.82. These estimates imply that the choices of individuals who exhibit high trust in the lab are consistent with a subjective probability q_{high} of 0.47 that the insurer will pay out, conditional on a hospitalization occurring—this is already quite low. Among individuals who exhibit low trust in the lab, the implied subjective probability of insurer payout, q_{low} , of only 0.24. These results suggest that the trust differences underlying variation in insurance demand are substantial in economic terms, and that even among individuals whose observed behavior in the lab suggests they have relatively high trust, there is considerable scope for policy interventions that improved faith in contract enforcement.

5.2 Welfare costs and the distributive burden of limited trust


To measure the welfare consequences of limited trust, we simulate demand under alternative, counterfactual trust levels. Following Einav, Finkelstein, and Cullen (2010), we quantify the welfare gains to a given consumer of receiving insurance by their willingness to pay for that product, as defined by its attributes (including the true probability, q , of a payout conditional on hospitalization, which in our case need not be the probability implied by revealed preference). Since consumers' willingness to pay for insurance depends not only (inter alia) on their trust in the insurer, the implied welfare losses relative to any counterfactual scenario will depend on the hypothesized trustworthiness of the product in that scenario. We consider welfare losses relative to two such counterfactuals.

We illustrate willingness to pay for these counterfactual products in Figure 2. That figure displays two demand curves, with consumers as a unit mass on the horizontal axis, and log prices on the vertical axis.²³ The lower, solid demand curve illustrates consumers' willingness to pay for the insurance product at estimated levels of trust and risk aversion, and given their baseline consumption levels and beliefs. Denoting the willingness to pay of individual i as a function of their (estimated or imposed) trust level q by $\pi_i^*(q)$, this demand curve depicts $\pi_i^*(\hat{q}_i)$ for estimated trust levels $\hat{q}_i \in \{\hat{q}_{\text{low}}, \hat{q}_{\text{high}}\}$. The upper, dashed demand curve represents consumers' willingness to pay for a counterfactual insurance product in which they have full trust (so that $q_i = 1$ for all

²³Observations below the 5th percentile and above the 95th percentile are trimmed from the figure for clarity of exposition.

i). Note that heterogeneity in trust levels implies that the rank ordering of consumers' willingness to pay is not the same for the two products.

Figure 2: Welfare costs



Notes: Figure illustrates demand curves under status quo ($q_i = \hat{q}_i$) and counterfactual, full ($q_i = 1$) trust levels, together with welfare losses from two counterfactual policy simulations. Region A provides a lower bound on welfare losses from biased beliefs, while Regions A+B represent welfare losses from low trust when these beliefs represent the true value of the product.

First, we consider the welfare losses relative to a scenario in which the insurance product is perceived to be fully credible, under the assumption that the *true* probability of insurance payout, q , is in fact one even under the status quo. In so doing, we follow an emergent literature that considers the role of ‘information frictions’ in insurance demand; in such a setting, where misperceptions of product attributes affect consumer choices, revealed willingness to pay may diverge from the welfare losses associated with failing to purchase insurance (Spinnewijn 2014, Handel and Kolstad forthcoming).

Welfare losses from limited trust in this scenario are borne by the fraction of individuals who would change their purchase decisions when $q_i = 1$. This is true for 16 percent of our sample. Among this subpopulation, average willingness to pay for a fully credible product is KShs 7,082.²⁴ Consequently we estimate the expected welfare loss relative to a perception of full trustworthiness as KShs 1,132, or approximately USD 15 at prevailing exchange rates during the study. This welfare loss is substantial: to put it in economic perspective, the expected loss per person is approximately 31 percent of the premium cost of the product.

Region A in Figure 2 provides a lower bound for the welfare losses arising from the first counterfactual scenario considered. This region illustrates the foregone willingness to pay of consumers

²⁴Note that since consumption under fully credible insurance is always $w - pi$, willingness to pay for a fully credible product has closed-form solution $w_i - [(1 - R)EU_i^0]^{1/(1-R)}$, where w_i is individual i 's consumption in the absence of any premium or hospitalization costs, EU_i^0 is subjective expected utility in the absence of insurance, and R is the coefficient of relative risk aversion.

who would be induced by complete trust to switch their insurance purchase decision *if those who switch are the eventual purchasers with the lowest willingness to pay*. Because some of those induced to switch—i.e., some of those individuals who have $\pi_i^*(\hat{q}_i) < \pi$ and $\pi_i^*(1) \geq \pi$ —will have greater willingness to pay under full trust than some of those who bought under partial trust (i.e., some of those with $\pi_i^*(\hat{q}_i) \geq \pi$), this region represents a lower bound on the actual welfare loss.

A second alternative is to consider the welfare losses associated with limited insurance under the alternative assumption that estimated trust levels for high- and low-trust types correspond to the true (and therefore heterogeneous) probabilities of insurer payout. Fixing counterfactual trustworthiness at one relative to this scenario implies that, in addition to the gains attributed to newly-insured individuals considered above, there would be inframarginal gains as the expected utility of already insured individuals improves. The latter is captured by the expected difference in willingness to pay for consumers who purchase under the status quo, multiplied by the fraction of such consumers in the sample.²⁵ Naturally, the resulting welfare losses attributable to limited trust are greater. Combining marginal and inframarginal effects, we estimate expected welfare costs relative to full trust in this scenario as KShs 17,479—nearly five times the premium cost of KShs 3,650.

The welfare loss arising from the second counterfactual considered is given by the combined areas of Regions A and B in Figure 2. In this case, total welfare losses comprise both the foregone willingness to pay of those induced to purchase by the improvement in the product, as well as the added value of the improved insurance product for those inframarginal consumers who would have purchased the product even under limited trust. In this case, the figure represents a point estimate for the welfare loss, rather than a lower bound, since exchanging individual consumers' willingness to pay under either scenario does not affect the difference in the integrals of the demand curves (where each integral is taken up to the point of intersection with the price).

The structural estimates also shed light on the distributive burden of this limited trust. This is perhaps seen most starkly in the first counterfactual exercise, where the burden of limited trust is concentrated among those who would change their insurance purchase decisions if their beliefs changed from $q_i \in \{q_{\text{low}}, q_{\text{high}}\}$ to $q_i = 1$. A simple first-pass measure of the distribution of the burden of low trust can be obtained as follows. First, we note that the average monthly household consumption level in the estimating sample is KShs 42,410, or about USD 565. Respondents who exhibit low trust in the lab have slightly lower mean consumption of 35,370 on average. But the 16 percent of sampled individuals for whom the optimal insurance choice switches from abstaining to purchasing when trust levels for all agents are increased to one, average monthly consumption is less than half that of that in the sample as a whole, at KShs 18,403. This suggests that if either informational interventions, legal support, or other policy mechanisms could address the problem of limited trust, the benefits of such interventions would be concentrated among the poorest households in this sample.

²⁵Willingness to pay under partial trust, q_i , is solved numerically for each individual, as the premium cost π_i^* that equates expected utilities with and without insurance:

$$\begin{aligned} \Pr_i[c = 0] \cdot u(w_i - \pi_i^*) + (1 - \Pr_i[c = 0]) \int_{\underline{c}}^{\bar{c}} \{q_i u(w_i - \pi_i^*) + (1 - q_i) \min[u(w_i - \pi_i^* - c), u(\underline{w})]\} dF_i(c) \\ = \Pr_i[c = 0] \cdot u(w_i) + (1 - \Pr_i[c = 0]) \int_{\underline{c}}^{\bar{c}} u(w_i - c) dF_i(c) \end{aligned}$$

where $F_i(c)$ is individual i 's subjective distribution of hospitalization costs in the case of an event, $\Pr_i[c = 0]$ is i 's subjective probability of no event; q_i is the subjective probability of insurer payout conditional on an event, and w_i is baseline consumption.

6 Conclusions

This paper has developed and tested a model of demand for indemnity insurance when the target population has limited trust in the insurer. The model reproduces two emergent stylized facts of the demand for insurance: that demand for insurance can be decreasing in measures of risk aversion, and that demand for insurance is increasing in measures of trust. We take predictions of the model to field and laboratory-experimental data from Kenya, and find results consistent with the model. These findings lend support to the view that limited trust in insurers constrains the adoption of indemnity insurance, while illustrating the external validity of lab-type measures of attitudes toward risk and trust. The relationship we find between trust and financial sector development provides microeconomic evidence of one potential mechanism for the cross-country relationship between trust and growth.

From a practical perspective, an important policy consideration is the manipulability of trust in the insurer. Other papers have suggested that trust may indeed be affected by policy, either through endorsements by trusted third parties (Cole et al. 2013), or through direct observation of insurance payouts (Cai et al. 2010). There would appear to be substantial scope for government regulation to improve the objective enforceability of insurance contracts. But as our findings confirm, there are no easy fixes through financial literacy interventions: limited trust may arise not from a lack of understanding, but rather from endemic problems in the delivery of health services.

The structural model's implications for the welfare costs and distributive incidence of limited trust make this task particularly urgent. Our findings suggest that those who are more averse to absolute levels of risk—likely the poorest and least well protected by informal networks—will be most deterred by the fear that insurers will default on their claims. These are precisely the households who stand to gain most from such financial products. The incidence of increased insurance demand arising from improved trust is likely to benefit this vulnerable population in particular.

References

- Akerlof, George A**, “The market for ‘lemons’: Quality uncertainty and the market mechanism,” *Quarterly Journal of Economics*, August 1970, *84* (3), 488–500.
- Alderman, Harold, John Hoddinot, and Bill Kinsey**, “Long term consequences of early childhood malnutrition,” *Oxford Economic Papers*, July 2006, *58* (3), 450–474.
- Ashraf, Nava, Iris Bohnet, and Nikita Piankov**, “Decomposing trust and trustworthiness,” *Experimental Economics*, 2006, *9* (3), 193–208.
- Banerjee, Abhijit, Esther Dufo, and Richard Hornbeck**, “Bundling health insurance and microfinance in India: There cannot be adverse selection if there is no demand,” Unpublished, MIT March 2014.
- Barr, Abigail**, “Trust and expected trustworthiness: experimental evidence from Zimbabwean villages,” *Economic Journal*, July 2003, *113* (489), 614–630.
- , “Attitudes toward risk in urban Ghana,” Unpublished, Oxford University June 2007.
- and **Pieter Serneels**, “Reciprocity in the workplace,” *Experimental Economics*, 2009, *12*, 99–112.
- , **Jean Ensminger, and Jeffrey C. Johnson**, “Social networks and trust in cross-cultural economic experiments,” in K. Cook, M. Levi, and R. Hardin, eds., *Whom can we trust?*, New York: Russell Sage Foundation, 2009.
- Beegle, Kathleen, Joachim De Weerd, and Stefan Dercon**, “Adult mortality and consumption growth in the age of HIV/AIDS,” *Economic Development and Cultural Change*, January 2008, *56* (2), 299–326.
- Ben-Ner, Avner and Louis Putterman**, “Trusting and trustworthiness,” *Boston University Law Review*, 2001, *81*, 523–551.
- Berg, Joyce, John Dickhaut, and Kevin McCabe**, “Trust, Reciprocity, and Social History,” *Games and Economic Behavior*, 1995, *10*, 122–142.
- Cai, Hongbin, Yuyu Chen, Hanming Fang, and Li-An Zhou**, “Microinsurance, Trust and Economic Development: Evidence from a Randomized Natural Field Experiment,” BREAD Working Paper No. 279 January 2010.
- Camerer, Colin**, *The Methods of Modern Experimental Economics*, Oxford University Press, forthcoming.
- Camerer, Colin F**, *Behavioral Game Theory*, Princeton, NJ: Princeton University Press, 2003.
- Chetty, Raj and Amy Finkelstein**, “Social insurance: Connecting theory to data,” in A Auerbach, R Chetty, M Feldstein, and E Saez, eds., *Handbook of Public Economics*, Vol. 5, Elsevier, 2013, pp. 111–193.
- Clarke, Daniel**, “A theory of rational hedging,” Unpublished, Oxford University 2011.

- Cole, Shawn A, Thomas Sampson, and Bilal Zia**, “Prices or Knowledge? What drives demand for financial services in emerging markets?,” *Journal of Finance*, December 2011, *66* (6), 1933–1967.
- , **Xavier Gine, Jeremy Tobacman, Petia Topalova, Robert M Townsend, and James Vickrey**, “Barriers to household risk management: Evidence from India,” *American Economic Journal: Applied Economics*, 2013, *5* (1), 104–135.
- Delavande, Adeline, Xavier Giné, and David McKenzie**, “Measuring subjective expectations in developing countries: a critical review and new evidence,” Unpublished, World Bank December 2009.
- Doherty, Neil A and Harris Schlesinger**, “Rational Insurance Purchasing: Consideration of Contract Nonperformance,” *Quarterly Journal of Economics*, 1990, *105* (1), 243–253.
- Easterly, William, Jozef Ritzen, and Michael Woolcock**, “Social cohesion, institutions, and growth,” *Economics and Politics*, July 2006, *18* (2), 103–120.
- Eckel, Catherine and Rick K Wilson**, “Is trust a risky decision?,” *Journal of Economic Behavior and Organization*, December 2004, *55* (4), 447–465.
- Einav, Liran, Amy Finkelstein, and Mark R Cullen**, “Estimating welfare in insurance markets using variation in prices,” *Quarterly Journal of Economics*, August 2010, *125* (3), 877–921.
- and —, “Selection in insurance markets: Theory and empirics in pictures,” *The Journal of Economic Perspectives*, Winter 2011, *25* (1), 115–138.
- Handel, Benjamin R and Jonathan T Kolstad**, “Health insurance for ‘humans’: Information frictions, plan choice, and consumer welfare,” *American Economic Review*, forthcoming.
- Harrison, Glenn W and John A List**, “Field Experiments,” *Journal of Economic Literature*, December 2004, *42* (4), 1009–1055.
- , **Steven J Humphrey, and Arjan Verschoor**, “Choice under uncertainty: Evidence from Ethiopia, India, and Uganda,” *Economic Journal*, March 2010, *120* (543), 80–104.
- Hendren, Nathaniel**, “Private information and insurance rejections,” *Econometrica*, September 2013, *81* (5), 1713–1762.
- Hey, John D**, “Experimental economics and the theory of decision making under risk and uncertainty,” *Geneva Papers on Risk and Insurance Theory*, 2002, *27*, 5–21.
- Holt, Charles A and Susan K Laury**, “Risk aversion and incentive effects,” *American Economic Review*, December 2002, *92* (5), 1644–1655.
- Karlan, Dean**, “Using Experimental Economics to Measure Social Capital and Predict Financial Decisions,” *American Economic Review*, December 2005, *95* (5), 1688–1699.
- , **Robert Osei, Isaac Osei-Akoto, and Christopher Udry**, “Agricultural decisions after relaxing credit and risk constraints,” *Quarterly Journal of Economics*, 2014, *129* (2), 597–652.
- Knack, S. and P. Keefer**, “Institutions and Economic Performance: Cross-Country Tests Using Alternative Institutional Measures,” *Economics and Politics*, 1995, *7* (3), 207–227.

- Manski, Charles F**, “Measuring Expectations,” *Econometrica*, September 2004, 72 (5), 1329–1376.
- Mobarak, A Mushfiq and Mark Rosenzweig**, “Informal risk sharing, index insurance, and risk-taking in developing countries,” *American Economic Review*, May 2013, 103 (3), 375–380.
- Morduch, Jonathan**, “Income smoothing and consumption smoothing,” *Journal of Economic Perspectives*, 1995, 9, 103–114.
- North, Douglass C.**, *Institutions, Institutional Change and Economic Performance*, Cambridge, U.K.: Cambridge University Press, 1990.
- Nunn, Nathan**, “The long term effects of Africa’s slave trades,” *Quarterly Journal of Economics*, February 2008, 123 (1), 139–176.
- and **Leonard Wantchekon**, “The slave trade and the origins of mistrust in Africa,” *American Economic Review*, 2011, 101 (7), 3211–3252.
- Rosenzweig, Mark and Hans Binswanger**, “Wealth, weather risk and the composition and profitability of agricultural investments,” *Economic Journal*, 1993, 103, 56–78.
- Schechter, Laura**, “Traditional trust measurement and the risk confound: An experiment in rural Paraguay,” *Journal of Economic Behavior and Organization*, February 2007, 62 (2), 272–292.
- Spinnewijn, Johannes**, “Heterogeneity, demand for insurance and adverse selection,” Unpublished, London School of Economics July 2014.
- Vyrastekova, Jana and Sander Onderstal**, “The Trust Game behind the Veil of Ignorance: A Note on Gender Differences,” Tinbergen Institute Discussion Paper no. 10-063/1 2005.
- Wilcox, Nathaniel T**, “Stochastic models for binary discrete choice under risk: a critical primer and econometric comparison,” in James C Cox and Glenn W Harrison, eds., *Risk aversion in experiments*, Vol. 12 of *Research in Experimental Economics*, Emerald Group Publishing Ltd, 2008.
- , “‘Stochastically more risk averse:’ A contextual theory of stochastic discrete choice under risk,” *Journal of Econometrics*, May 2011, 162 (1), 89–104.
- Zak, Paul J and Stephen Knack**, “Trust and growth,” *Economic Journal*, April 2001, 111, 295–321.

Appendix A Lab-experimental protocols

Appendix A.1 Recruitment and timing

The same household members covered by the baseline surveys participated in a pair of laboratory experiments conducted in the field at baseline. These laboratory experiments were employed to provide measures of attitudes toward risk and trust that could be used to test hypotheses about the demand for insurance. The laboratory games were played sequentially, in randomized order. Payoffs were revealed after each game, but actual payoffs were not made until the end, with the payoff of one randomly selected game paid out to the participant.

Appendix A.2 Gamble-choice game

To obtain a measure of Wananchi members’ preferences toward risk, we employed a standard gamble-choice game, based on the instrument of Holt and Laury (2002, henceforth HL). Our specific design is adapted from Barr (2007) to the Kenyan context. We employ the script and relative payoff values that Barr developed for Ghana, translated into comparable expected returns for Kenyan tea farmers. This is reported below.

Appendix A.2.1 Procedures

This game consists of a series of tasks, in each of which the subject chooses between two binary lotteries, one ‘safe’ lottery and one ‘risky’ lottery. Each lottery consists of a high-payoff outcome (H) and a low-payoff outcome (L). Payoffs from either of these lotteries are constant within the series. In any given task, the probability of the high-payoff outcome is the same in both the risky and safe lotteries; this probability varies across tasks.

Table A.1: Gamble-choice game: payoffs and probabilities in gain- and loss-framed series.

Task	Pr(H)	Gain frame				Loss frame				E[$\pi_r - \pi_s$]
		Risky		Safe		Risky		Safe		
		H_r	L_r	H_s	L_s	H_r	L_r	H_s	L_s	
1	0.8	300	0	100	50	0	-300	-200	-250	150
2	0.7	300	0	100	50	0	-300	-200	-250	125
3	0.6	300	0	100	50	0	-300	-200	-250	100
4	0.5	300	0	100	50	0	-300	-200	-250	75
5	0.4	300	0	100	50	0	-300	-200	-250	50
6	0.3	300	0	100	50	0	-300	-200	-250	25

Notes: Table shows probability of high payoff, H, in risky and safe lottery choices, together with high and low payoff values, for gain- and loss-frame HL series. $E[\pi_r - \pi_s]$ denotes difference in expected return from risk versus safe lottery. All payoff values expressed in Kenya Shillings. Subjects endowed with KShs 300 prior to participation in loss-frame series.

We played two series of this game, as shown in Table A.1: a *gain-frame* series, and a *loss-frame* series. In the gain-frame series, subjects began with an initial endowment of zero, and had an opportunity to win either KShs 300 or KShs 0, if they chose the risky lottery, or KShs 100 or KShs 50, if they chose the safe lottery.²⁶ In the loss-frame series, subjects were endowed with KShs 300

²⁶The prevailing exchange rate at the time of the laboratory experiment was KShs 75/USD, meaning that the maximum possible payout in this series is USD 4.

prior to play, so that the reduced-form payoffs in each task of the loss-frame series are equivalent to those in the gain-frame series. These were played sequentially, with payoffs made after both series were complete. Monetary payoffs were based on a single task selected at random from across the two series. This payoff mechanism was explained to participants in advance. To the extent that the loss frame changes subjects' reference point, we may expect that differences in risk preferences manifested in the gain-frame and loss-frame tasks are driven in part by loss aversion.

Appendix A.2.2 Script

The following script was adapted, with permission, from Abigail Barr, based on her work in Ghana (Barr 2007). Session leaders referred to the visual aids in Figures A.1 and A.2. Instructions for enumerators (not to be read aloud), are provided below *in italics*.

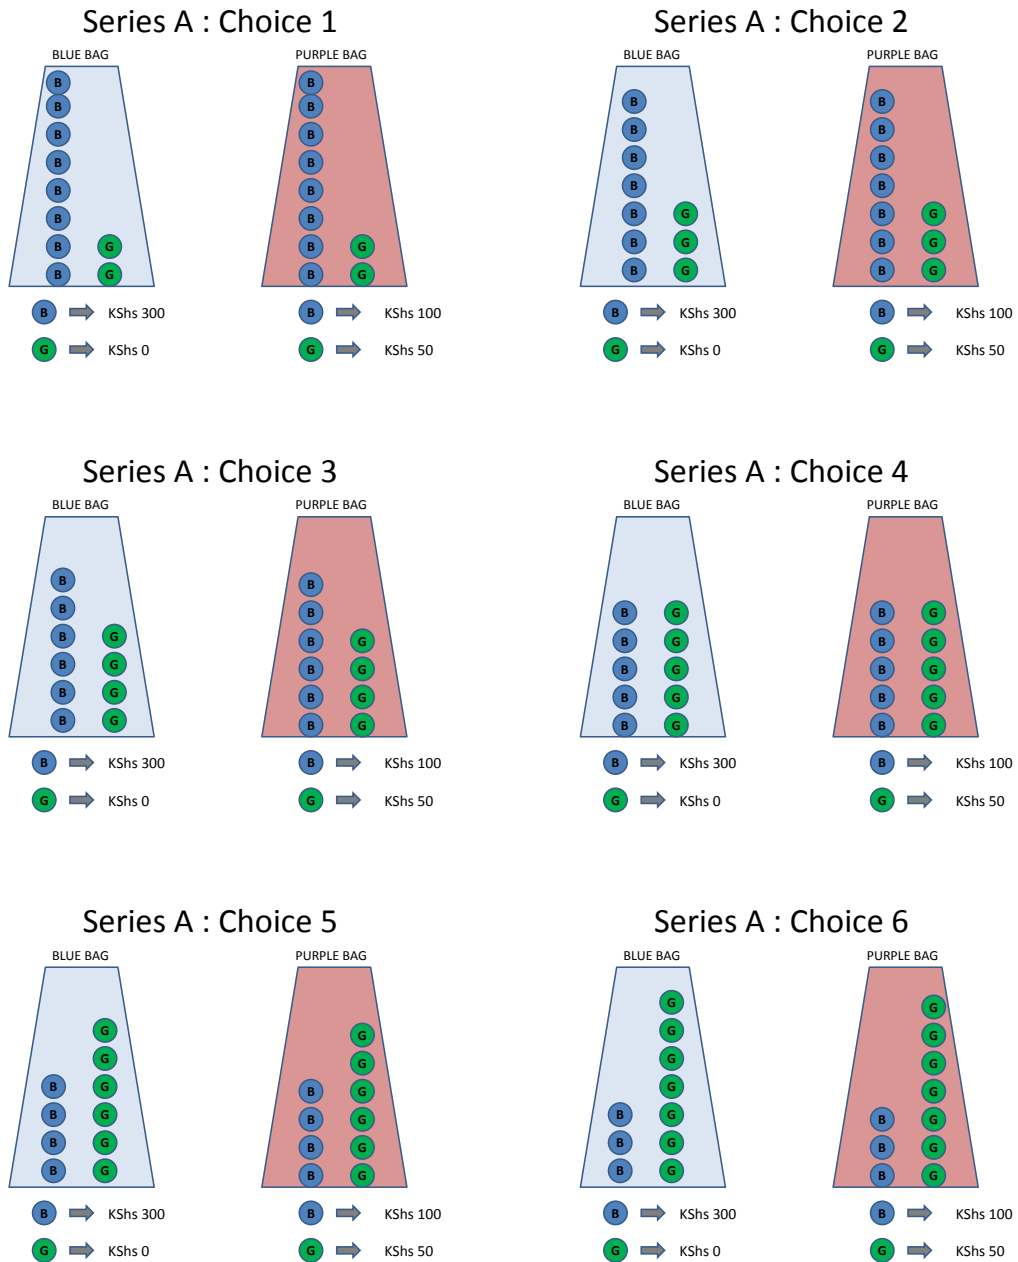
1. **Welcome and introductions:** Thank you everyone for joining us here today. Before we start, let us introduce ourselves. We work for Steadman, a survey firm, and we are working together with researchers from Oxford University.
2. **The purpose of today's meeting:** The activity you are about to take part in is linked to the survey you have just completed, but it does not involve us asking you questions in the way that we have done before. There will be two activities in this session. Each will involve us playing some games, and the games played in each of the two activities will involve an opportunity to earn real money. In the games you will be asked to make some decisions. One of the decisions you make will determine the amount of money you will get to take home. However, you will not know which decision is going to determine your pay until the very end. So, you need to take all the decisions seriously. There are no right or wrong decisions. You need to make the decisions that you feel are right for you. The decisions you make will help us understand how people in Kenya feel about risky situations and that, in turn, will help to design ways of helping you and other people in Kenya cope with the risks you face as part of your everyday life.
3. **Plan for the first activity:** This is how this activity will be carried out. After this introduction, we will describe the type of decisions we are going to ask you to make. We will work through a couple of examples and you will have a chance to ask questions. It is important to you and to us that you understand the decisions you are making. So, please listen carefully.

Once everyone feels that they understand, we will start the game. We will describe the first decision to you, you will make your decision and the enumerator next to you will record that decision on the form you have been given next to the number 1 (*point at form*). Then we will describe the second decision to you, you will make your decision and the enumerator will record it next to the number 2 (*point at form*). We will carry on until you have each made 12 decisions and they are all recorded on the form.

Then, each of you will meet with us one at a time. In that meeting, we will randomly select one of the 12 decisions you have made in a way that we will describe later on. Then you will play the game according to that decision. Some of you could win as much as KShs 300. Some of you may go home with nothing.. On average, the people playing these games will win about KShs 150 each. But, remember, it depends on the decisions you make.

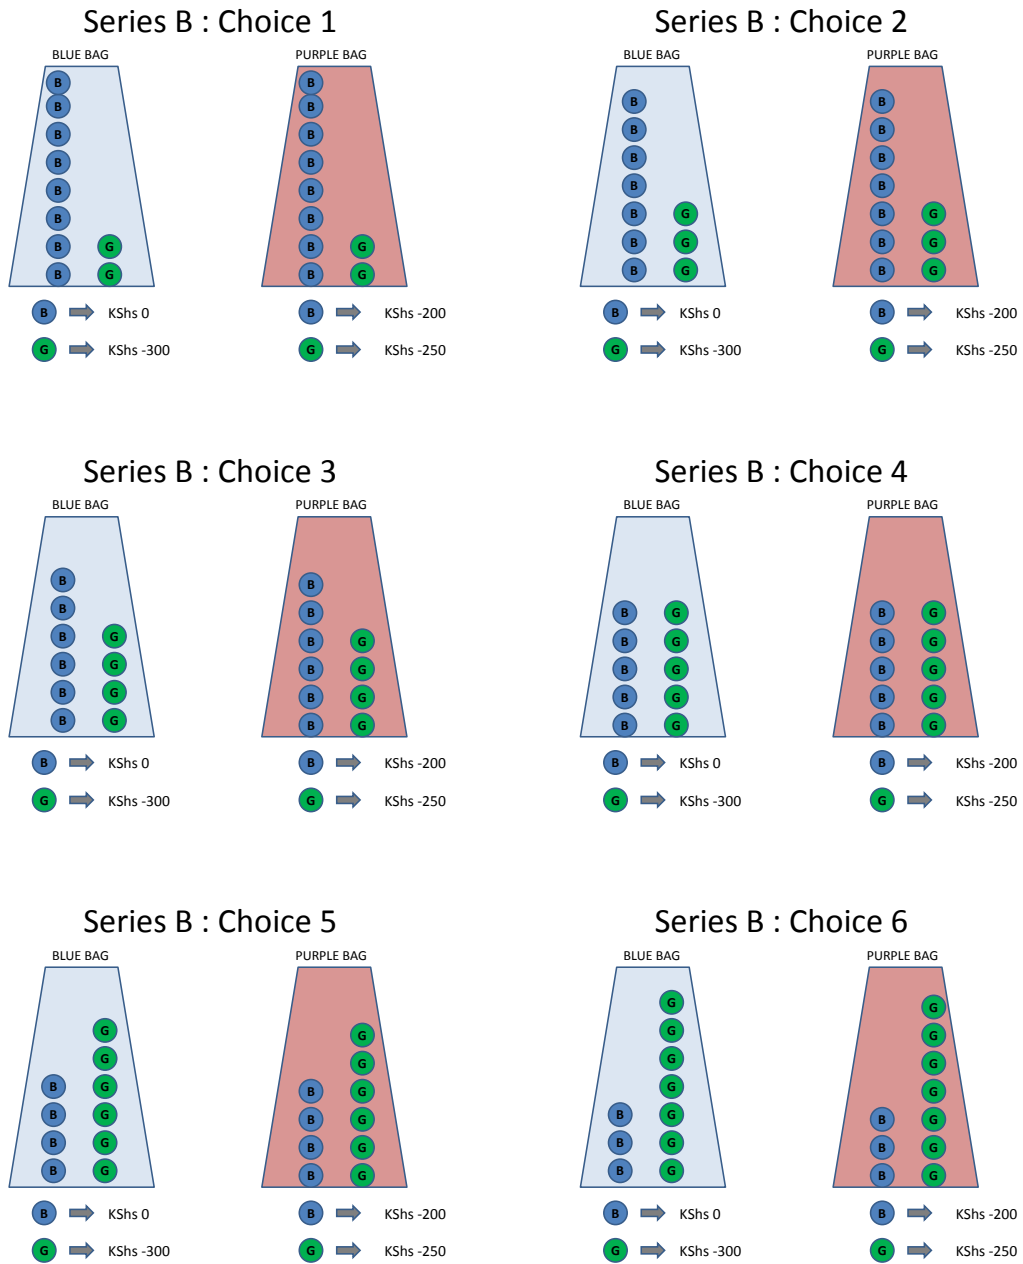
Once you have been paid, you will be free to go.

Figure A.1: Visual aids for HL Gamble Choice Game, gain-frame series



Notes: Figures adapted from protocol of Barr (2007).

Figure A.2: Visual aids for HL Gamble Choice Game, loss-frame series



Notes: Figures adapted from protocol of Barr (2007).

4. **General rules:** Please do not try and talk to each other. We are interested in the decisions you make on your own. And, remember, there are no right and wrong decisions. You need to make the decisions that feel right for you. It is very important that you do not talk to each other. If there is talking, we will have to stop the meeting and no one will end up getting paid. If you have questions ask one of us. PLEASE DO NOT TALK TO EACH OTHER.
5. **Free to leave:** We hope you are going to enjoy the games. However, if at any time you decide that you do not want to take part, you are free to leave.
6. **Now we are going to describe the decisions you are going to be asked to consider.** Each decision task will involve choosing one of two lotteries. (*Enumerators give subjects the relevant decision card to look at.*) Let me tell you what we mean by a lottery.

7. **Example 1 [Refer to Series A: Choice 1]:** Suppose we put 8 blue blocks and 2 red blocks in this blue bag, and I tell you that if you pull a blue block out you will get KShs 300, but if you pull a red block out you will get KShs 0. This is a lottery.

Now suppose that we put 8 blue blocks and 2 red blocks in this purple bag, and I tell you that if you pull a blue block out you will get KShs 100, but if you pull a red block out you will get only KShs 50. This is another lottery.

Now, which lottery would you choose? Would you prefer to pick a block out of the blue bag or the purple bag? I should add, you are not allowed to look when picking the block.

Enumerator: walk through an actual choice for this case, pick a block, and explain the payoffs.

I am going to work through another example, but first, do you have any questions?

8. **Example 2 [Refer to Series B: Choice 6]** Let's work through another example. (*Enumerators place the correct decision card in front of the subjects.*) In this example, you are given KShs 300 right at the start. Then, as before, you will be asked to choose whether you wish to select a block from the blue bag or the purple bag.

The blue bag contains 3 blue blocks and 7 red blocks. If you pull a blue block out you get to keep the 300 KShs. If you pull a red block out you have to give back all of the KShs 300.

Note on the decision card in front of you that the red block states **minus** 300 on it—the minus indicates that it is what you have to **give back** (*enumerators point to appropriate place on decision cards*). If you pull out a blue block you don't have to give anything back (*enumerators point to appropriate place on decision card*).

Now we need to look at the purple bag. The purple bag also contains 3 blue blocks and 7 red blocks. If you pull a blue block out you have to give back KShs 200, but if you pull out a red block, you have to give back KShs 250. See, the **minus** KShs 250 (*enumerators point*).

Now which lottery would you choose? Would you prefer to pick a block out of the blue bag or the purple bag? Remember, you are not allowed to look when picking the block.

Enumerator: walk through an actual choice for this case, pick a block, and explain the payoffs.

All of the decisions you are going to be asked to make are like one or other of these examples. The number of blocks and the colours of the blocks will change and the amounts of money associated with each colour will change, but all other aspects will be like one or other of these examples.

Does anybody have any questions?

9. **Randomized pay again:** Remember, you are going to make 12 decisions. Then we are going to randomly select one of the 12 decisions you have made to work out your winnings from the games. We are going to do this by putting 12 counters in this bag. The counters will be numbered 1 to 12. We will ask one of you to volunteer to come up and pull one of the counters out of the bag without looking. The number on the counter pulled out will tell us which decision to use to work out your winnings.

Then, each of you will come to us one at a time. The blocks will be put in the bags just as they were when the picked decision was first described to you and you will be asked to pick a ball from the bag you chose and the enumerator wrote on your form. You will not be allowed to look in the bag as you picked. The colour of the block you pull out will determine your winnings just as we described when you made your decision.

Does anybody have any questions?

Now, lets start.

10. **The tasks**

Enumerator: Refer to visual aids for the number/color of blocks and the payoffs associated with each.

Gain (Series A):

For Tasks A.2 A.6, state the following at the outset of the task:

This game is similar to the previous one. However, the blue bag now contains ... blue blocks instead of ... (one blue block LESS than in the previous game) and ... red blocks instead of ... (one red block MORE than in the previous game). The purple bag contains ... blue blocks instead of ... (one blue block LESS than in the previous game) and ... red blocks instead of ... (one red block MORE than in the previous game).

For Tasks A.1 A.6, then describe the task as follows:

This blue bag contains ? blue blocks and ? red blocks. if you pull a blue block out you will get ... KShs, but if you pull a red block out you will only get ... KShs.

This purple bag contains ? blue blocks and ? red blocks. If you pull a blue block out you will get ... Kshs, if you pull a red block out, you will get ... KSHs.

Which bag do you want to pull a block out of? (*enumerators record*)

Loss (Series B):

The next six games will have the same structure as the first six games: in the blue bag, the blue blocks will progressively decrease, while the red blocks will progressively increase. In the purple bag, the blue blocks will progressively decrease while the red blocks will progressively increase. (*Enumerators: quickly show the cards of the five games ahead*)

For tasks B.2 B.6, begin by saying the following:

This game is similar to the previous one. However, the blue bag now contains ... blue blocks instead of ... (one blue block LESS than in the previous game) and ... red blocks instead of ... (one red block MORE than in the previous game). The purple bag contains ... blue blocks instead of ... (one blue block LESS than in the previous game) and ... red blocks instead of ... (one red block MORE than in the previous game).

For tasks B.1 B.6, say the following:

This decision task starts will me giving you KShs 300.

Now, this blue bag contains ? blue blocks and ? red blocks. If you pull a blue block out

- you have to give back ... KShs OR
- you have to give back all of the ... KShs OR
- you dont have to give anything b

If you pull out a red block

- you have to give back ... KShs OR
- you have to give back all of the ... KShs OR
- you dont have to give anything back.

The purple bag contains ? blue blocks and ? red blocks. If you pull a blue block out

- you have to give back ... KShs OR
- you have to give back all of the ... KShs OR
- you dont have to give anything back.

If you pull out a red block

- you have to give back ... KShs OR
- you have to give back all of the ... KShs OR
- you dont have to give anything back.

Which bag do you want to pull a block out of? (enumerators record)

Remind them repeatedly about the money: Remember, this could be the decision that matters. If this is the decision task that is picked at the end of the meeting, then you will be picking a block out of either the blue or the purple bag to determine your winnings. And you will have to pick the block out without looking. Which bag do you want to be picking the block out of, the blue one or the purple one?

Appendix A.3 Trust game

In the second part of our baseline laboratory experiments, we sought to elicit a measure of trust.²⁷ To do so we use a variant of the Trust Game, originally designed by Berg, Dickhaut, and McCabe (1995, henceforth BDM). We adapt the design employed in Zimbabwe by Barr (2003) as described below.

Appendix A.3.1 Procedures

The basic setup of the Trust Game is as follows. Players are assigned to one of two roles, Sender or Receiver. Both are endowed with KShs 200 at the outset of the game. The Sender can then decide to send a portion of their endowment to the Receiver (from zero to KShs 200, in increments of KShs 50). Any amount that is sent to the Receiver is tripled. The Receiver can then decide to return any portion of this tripled amount—possibly none—to the Sender, at which point the game concludes.

²⁷Ideally, this measure would directly capture the extent of trust in the insurer. However, given practical constraints and a desire to avoid contamination of the field experiment with ‘experimenter demand’ effects, we settled on a more general measure of trust.

We adapt this basic setup in three ways. First, we elicit the decisions of the Receiver by strategy method.²⁸ Second, since our interest is primarily in Sender behavior, and since we wished to abstract from issues of learning and repeated interaction, we divide the ten participants in each laboratory settings into two groups of five at random. One individual in each group was selected to play the role of the Receiver, while the remaining four individuals played the role of the Sender. The Receiver’s strategy profile determined payoffs for all Senders, while one Sender’s decision was chosen randomly and anonymously to determine the payout of the Receiver.²⁹ Third, we enforce that the Wananchi Delegate was *always* the Receiver in one of the two sessions played for each tea centre. Since the Delegate represents an authority figure associated with the SACCO and so, potentially, with the credibility of the insurance product being marketed, comparison of trusting behavior of trusting behavior individuals playing with ordinary Wananchi members and those playing with the Delegate is expected to be informative about trust in financial institutions.

We find that Senders invest slightly less on average when they have been randomly paired to play with the Delegate in their center. A regression of the share invested on an indicator for whether the Receiver is a Delegate gives a coefficient of -0.05 (standard error 0.03, clustered by tea centre; this difference is significant at the 10 percent level). Caution is required in interpreting this difference as an absence of trust in institutions, since differences in altruism, or distributional preferences over outcomes that include income outside the game, may also drive the difference in observed play.

Appendix A.3.2 Script

We adapt the script employed by Barr, Ensminger, and Johnson (2009), who conduct a version of the BDM Trust Game both among the Orma of eastern Kenya and in urban Accra, Ghana. The following script draws directly on their work, and is provided here for transparency.

Instructions

Now we will describe the second game we will play together. This game is played by pairs of individuals. Each pair is made up of a Player 1 and a Player 2.

Each of you will play this game with someone from your own village. In this game, Player 2 will always be *[insert name of Player 2.]*

However, *[insert name of Player 2]* will never be able to tell what decisions you personally have made. Only the research team knows what you do, and we will never tell anyone else.

The research team will give KShs 200 to each Player 1 and another KShs 200 to each Player 2. Player 1 then has the opportunity to give a portion of their KShs 200 to Player 2. They could give KShs 200, or KShs 150, or KShs 100, or KShs 50, or nothing.

[Note: It is important to allow only 5 options for dividing the money—this is to simplify the game and to create the same focal points across sites.]

²⁸See, inter alia, Ashraf et al. (2006), and Barr, Ensminger, and Johnson (2009) for a related uses of the strategy method in the Trust Game. While it is a general concern that use of the strategy method elicits different behavior than sequential play in laboratory experiments, Vyrastekova and Onderstal (2005) provide evidence that the strategy method does not substantially affect Sender behavior, which is the focus of this paper.

²⁹To the extent that the Receiver has other-regarding preferences, this setup is likely to lead to more generous behavior on her part. Anticipating this, forward-looking Senders may exhibit relatively more trusting behavior than in the standard, two-person Trust Game. Because we are interested not in the comparison of behavior in this game with other implementations of the Trust Game in the literature, but rather in the relationship between variation in behavior in this game to insurance purchase decisions, there is little reason that this modification likely to confound our results.

Whatever amount Player 1 decides to give to Player 2 will be tripled by the research team before it is passed on to Player 2. Player 2 then has the option of returning any portion of this tripled amount to Player 1.

Then, the game is over.

Player 1 goes home with whatever he or she kept from their original KShs 200, plus anything returned to them by Player 2. Player 2 goes home with their original KShs 200, plus whatever was given to them by Player 1 and then tripled by the research team, minus whatever they returned to Player 1.

Here are some examples.

[You should work through these examples by having all the possibilities laid out in front of people, with Player 1's options from KShs 200 to KShs 0 and a second column showing the effects of the tripling. As you go through each example demonstrate visually what happens to the final outcomes for each Player. Be careful to remind people that Player 2 always also has the original KShs 200]:

1. Imagine that Player 1 gives KShs 200 to Player 2. The research team triples this amount, so Player 2 gets KShs 600 (3 times KShs 200 equals KShs 600) over and above their initial KShs 200. At this point, Player 1 has nothing and Player 2 has KShs 800. Then Player 2 has to decide whether they wish to give anything back to Player 1, and if so, how much. Suppose Player 2 decides to return KShs 150 to Player 1. At the end of the game Player 1 will go home with KShs 150 and Player 2 will go home with KShs 650.
2. Now let's try another example. Imagine that Player 1 gives KShs 150 to Player 2. The research team triples this amount, so Player 2 gets KShs 450 (3 times KShs 150 equals KShs 450) over and above their initial KShs 200. At this point, Player 1 has KShs 50 and Player 2 has KShs 650. Then Player 2 has to decide whether they wish to give anything back to Player 1, and if so, how much. Suppose Player 2 decides to return KShs 0 to Player 1. At the end of the game Player 1 will go home with KShs 50 and Player 2 will go home with KShs 650.
3. Now let's try another example. Imagine that Player 1 gives KShs 100 to Player 2. The research team triples this amount, so Player 2 gets KShs 300 (3 times KShs 100 equals KShs 300) over and above their initial KShs 200. At this point, Player 1 has KShs 100 and Player 2 has KShs 500. Then Player 2 has to decide whether they wish to give anything back to Player 1, and if so, how much. Suppose Player 2 decides to return KShs 150 to Player 1. At the end of the game Player 1 will go home with 250 and Player 2 will go home with KShs 350.
4. Now let's try another example. Imagine that Player 1 gives KShs 50 to Player 2. The research team triples this amount, so Player 2 gets KShs 150 (3 times KShs 50 equals KShs 150) over and above their initial KShs 200. At this point, Player 1 has KShs 150 and Player 2 has KShs 350. Then Player 2 has to decide whether they wish to give anything back to Player 1, and if so, how much. Suppose Player 2 decides to return KShs 100 to Player 1. At the end of the game Player 1 will go home with KShs 250 and Player 2 will go home with KShs 250.
5. Now let's try another example. Imagine that Player 1 gives nothing to Player 2. There is nothing for the research team to triple. Player 2 has nothing to give back and the game ends here. Player 1 goes home with KShs 200 and Player 2 goes home with KShs 200.

Note that the larger the amount that Player 1 gives to player 2, the greater the amount that can be taken away by the two players together. However, it is entirely up to Player 2 to decide what he should give back to Player 1. The first player could end up with more than KShs 200 or less than KShs 200 as a result. We will go through more examples with each of you individually

when you come to play the game. In the mean time, do not talk to anyone about the game. Even if you are not sure that you understand the game, do not talk to anyone about it. This is important. If you talk to anyone about the game while you are waiting to play, we must disqualify you from playing.

[Bring in each Player 1 one by one. Use as many of the examples below as necessary.]

6. Imagine that Player 1 gives KShs 200 to Player 2. The research team triples this amount, so Player 2 gets KShs 600 (3 times KShs 200 equals KShs 600) over and above their initial KShs 200. At this point, Player 1 has nothing and Player 2 has KShs 800. Then Player 2 has to decide whether they wish to give anything back to Player 1, and if so, how much. Suppose Player 2 decides to return KShs 300 to Player 1. At the end of the game Player 1 will go home with KShs 300 and Player 2 will go home with KShs 500.
7. Now lets try another example. Imagine that Player 1 gives KShs 150 to Player 2. The research team triples this amount, so Player 2 gets KShs 450 (3 times KShs 150 equals KShs 450) over and above their initial KShs 200. At this point, Player 1 has KShs 50 and Player 2 has KShs 650. Then Player 2 has to decide whether they wish to give anything back to Player, and if so, how much. Suppose Player 2 decides to return KShs 50 to Player 1. At the end of the game Player 1 will go home with KShs 100 and Player 2 will go home with KShs 600.
8. Now lets try another example. Imagine that Player 1 gives KShs 100 to Player 2. The research team triples this amount, so Player 2 gets KShs 300 (3 times KShs 100 equals KShs 300) over and above their initial KShs 200. At this point, Player 1 has KShs 100 and Player 2 has KShs 500. Then Player 2 has to decide whether they wish to give anything back to Player 1, and if so, how much. Suppose Player 2 decides to return KShs 0 to Player 1. At the end of the game Player 1 will go home with KShs 100 and Player 2 will go home with KShs 500.
9. Now lets try another example. Imagine that Player 1 gives KShs 50 to Player 2. The research team triples this amount, so Player 2 gets KShs 150 (3 times KShs 50 equals KShs 150) over and above their initial KShs 200. At this point, Player 1 has KShs 150 and Player 2 has KShs 350. Then Player 2 has to decide whether they wish to give anything back to Player 1, and if so, how much. Suppose Player 2 decides to return KShs 100 to Player 1. At the end of the game Player 1 will go home with KShs 250 and Player 2 will go home with KShs 250.
10. Now lets try another example. Imagine that Player 1 gives nothing to Player 2. There is nothing for the research team to triple. Player 2 has nothing to give back and the game ends here. Player 1 goes home with KShs 200 and Player 2 goes home with KShs 200.

Now, can you work through these examples for me:

11. Imagine that Player 1 gives KShs 150 to Player 2. So, Player 2 gets KShs 450 (3 times KShs 150 equals KShs 450) over and above their initial KShs 200. At this point, Player 1 has KShs 50 and Player 2 has KShs 650. Suppose Player 2 decides to return KShs 250 to Player 1. At the end of the game Player 1 will have how much? *[the initial KShs 200-KShs 150 (given to Player 2)=KShs 50+return from player 2 of KShs 250=KShs 300. If they are finding it difficult, talk through the maths with them and be sure to use demonstration with the actual money].* And Player 2 will have how much?

[Their original KShs 200+KShs 450 (after the tripling of the KShs 150 sent by Player 1)-KShs 250 they return to Player 1= KShs 400, if they are finding it difficult, talk through the maths with them].

12. Imagine that Player 1 gives KShs 50 to Player 2. So Player 2 gets KShs 150 (3 times KShs 50 equals KShs 150) over and above their initial KShs 200. Then, suppose that Player 2 decides to give KShs 50 back to Player 1. At the end of the game Player 1 will have how much?

[The initial KShs 200-KShs 50 (given to Player 2)=KShs 150+return from player 2 of KShs 50=KShs 200. If they are finding it difficult, talk through the maths with them and be sure to use demonstration with the actual money]. And Player 2 will have how much? *[Their original KShs 200+KShs 300 (after the tripling of the KShs 150 sent by Player 1)-KShs 50 they return to Player 1=KShs 300, if they are finding it difficult, talk through the maths with them].*

Game play

First player: You are Player 1. Here is your KShs 200. *[At this point KShs 200 is placed on the table in front of the player.]* While I am turned away, you must hand me the amount of money you want to be tripled and passed on to Player 2. You can give Player 2 nothing, KShs 50, KShs 100, KShs 150, or KShs 200. Player 2 will receive this amount tripled by me plus their own initial KShs 200. Remember the more you give to Player 2 the greater the amount of money at his or her disposal. While Player 2 is under no obligation to give anything back, we will pass onto you whatever he or she decides to return. *[Now the player hands back whatever he or she wants to have tripled and passed to player 2.]*

[Note to researcher: Finish all Player 1s and send them to a third holding location—they must not return to the group of Player 1s who have not played and they must not join the Player 2s. Once all Player 1s have played you can begin to call Player 2s. Player 2s can be paid off immediately after they play and sent home.]

Second player (strategy method): You are Player 2. First, here is your KShs 200. *[Put the KShs 200 in front of Player 2.]* Lets put that to one side. *[Move the KShs 200 to one side but leave it on the table.]*

This pile represents Player 1s initial KShs 200. *[Put this KShs 200 in front of the researcher.]*

Player 1 can either give you KShs 0, 50, 100, 150, or 200. Then the research team will triple this amount. This means that you will receive either 0, 150, 300, 450, or 600 from player 1.

In a moment we will find out which of these amounts you have received from the Player 1 that you are playing with. But before we do that, I would like you to tell me—for each of the possible amounts that you can receive from Player 1 and that have then been tripled by the research team—how much of this money you would like to give back to Player 1. So...

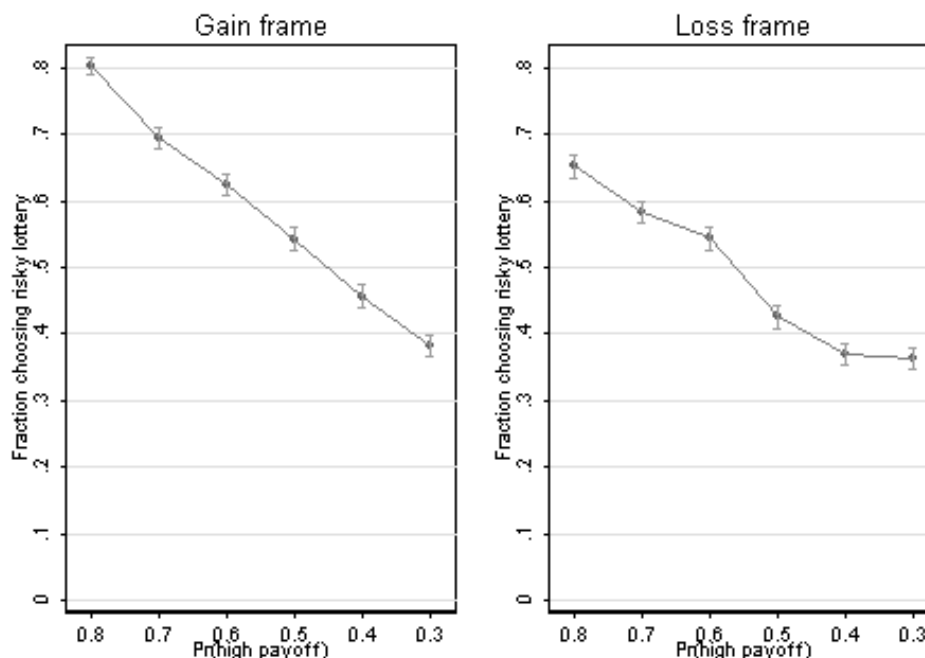
- If player 1 gives you 0, there is nothing for the research team to triple, and nothing for you to give back.
- If player 1 gives you 50, and the research team triples this amount to make it 150, you can choose to give back either 0, 50, 100, or 150 KShs.
- If player 1 gives you 100, and the research team triples this amount to make it 300, you can choose to give back either 0, 50, 100, 150, 200, 250, or 300 KShs.
- If player 1 gives you 150, and the research team triples this amount to make it 450, you can choose to give back either 0, 50, 100, or 150, 200, 250, 300, 350, 400, or 450 KShs.
- If player 1 gives you 200, and the research team triples this amount to make it 600, you can choose to give back either 0, 50, 100, or 150, 200, 250, 300, 350, 400, 450, 500, 550, or 600 KShs.

I will now ask you what you would like to do in each of these situations, should it be the one that actually reflects what Player 1 has done. Remember, you will play the game with one of these people in the role of Player 1, and they will actually choose one of these amounts to give to you. So the choice you make will count—and you will not be able to change your decision once you learn what Player 1 actually did give.

Appendix B Parameterization and estimation of risk preferences

Figure B.1 displays population frequencies for the choice of the risky lottery, by task, for both the gain-frame and the loss-frame lotteries. For ease of interpretation, tasks are indexed by the probability of the high-value outcome in that task, which ranges from 0.8 to 0.3. As expected, the fraction of individuals choosing the risky lottery declines as the probability of the high payout decreases. This reflects in part the change in the expected income difference between the risky and safe lotteries, which falls from KShs 150 to KShs 25 as the probability of the high payoff falls from 0.8 to 0.3.

Figure B.1: Decisions in gain-frame and loss-frame lottery-choice experiment



Notes: Figure shows the population frequency of individuals choosing the risky lottery in each task. This is illustrated for all participants in the lottery-choice experiment in treatment villages.

A rational, expected-utility maximizing individual, with weakly risk averse preferences should switch from choosing the risky to the safe lottery at most once over the course of the six tasks of the gain-frame series. Assuming that the individual's preferences over outcomes in this lottery can be represented by a constant relative risk aversion (CRRA) utility function of the form $u(x) = x^{1-R}/(1-R)$, then for such individuals it is possible to use their observed decisions to place bounds on the CRRA coefficient R . An individual choosing the risky lottery in all tasks must have $R \leq 0.22$, whereas an individual choosing the safe lottery in all tasks must have $R \geq 0.82$. Those who switch from risky to safe between tasks 1 and 6 will have an R that can be bounded within a strict subset of the interval $(0.22, 0.82)$.

In practice, only 53 percent of individuals make decisions in the gain-frame series that are can

be rationalized as the deterministic choices of an individual with weakly risk averse preferences.³⁰ Such inconsistent behavior is common among decision problems over risky choices in laboratory experiments. For example, Hey (2002) reports that 30 percent of subjects in a laboratory setting make different decisions when faced with the *same* task twice. To provide an individual-specific measure of risk preferences at the individual level for all individuals, we estimate CRRA parameters by maximum likelihood for each person individually (see, e.g., Harrison et al. 2010).³¹

To do so, we assume that people place values on each lottery according to expected utility theory, with CRRA utility defined as above. To allow for the possibility of errors, we use the ‘contextual choice’ specification of Wilcox (2008, 2011). Defining $EU_1(R)$, $EU_0(R)$ as the expected utility of an agent with CRRA coefficient R in the risky and safe lotteries, respectively, we model agents’ binary decision, D , to choose the risky lottery as

$$D = \mathbf{1} \left\{ \frac{EU_1(R) - EU_0(R)}{u(\bar{x}; R) - u(\underline{x}; R)} + e \right\}, \quad (8)$$

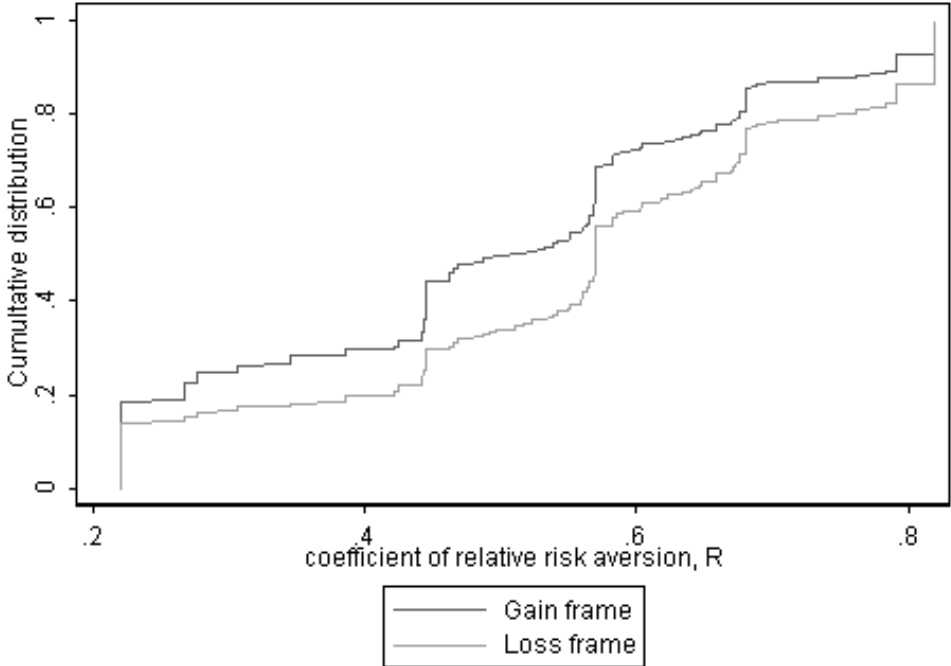
where $e \sim (N, \sigma_e)$, and $u(\bar{x}; R)$, $u(\underline{x}; R)$ are utility of highest/lowest payoffs that can occur in either lottery (that is, $\bar{x} = 300$; $\underline{x} = 0$). This specification differs from familiar probit models of choice only in that the expected utility differential is scaled by the difference in utility between the highest and lowest outcomes. As Wilcox argues, this specification ensures that an increase in R makes an agent less likely to choose the risky alternative, over all values of R —a desirable property that is *not* satisfied by linear index models that take the expected utility difference as their index.

The resulting distribution of estimated risk preferences is illustrated in Figure B.2 for the gain-frame and loss-frame lotteries. The mean value of R_G , the CRRA coefficient in the gain-frame lottery, is 0.49 (standard deviation 0.19). As evident in the raw data, behavior in the loss-frame sequence is consistent with a greater degree of risk aversion; the mean estimated coefficient of relative risk aversion is 0.55 (0.19) in this series. This puts these estimates in the same range as those found in similar laboratory experiments in the field; for example, Harrison and coauthors (2010), in their EUT model assuming homogeneous preferences, estimate a population parameter of $R = 0.54$.

³⁰A comparable number (52 percent) exhibit such consistent preferences in the loss-frame series.

³¹For individuals who choose either all risky or all safe lotteries, R cannot be estimated by maximum likelihood. We impose values of $R = 0.22, 0.82$ in these polar cases. These reflect the highest and lowest values of R , respectively, consistent with observed choices.

Figure B.2: Distribution of fitted coefficient of relative risk aversion



Appendix C Financial literacy treatment

The financial literacy treatment implemented as part of this experiment had no statistically significant impact on health insurance demand. This finding is informative to the extent that the curriculum was relevant, and that that participation rates among the study sample were high for this treatment.

The financial literacy intervention that we study was designed by Microfinance Opportunities and the Swedish Co-operative Centre and Vi Agroforestry (SCC-Vi). It combined materials developed by Microfinance Opportunities with a peer-led ‘study circles’ teaching method that had been used by SCC to promote the understanding and adoption of agricultural technologies.

Study circles consisted of a ten-session course, conducted over as many weeks. Each was led by a member of the tea centre (typically, the Wananchi ‘delegate’) and included the nine other sampled individuals as group members. Group leaders received initial training in both the pedagogical approach and the substantive content, and were provided with a workbook of instructional materials. Following this they convened a series of ten meetings to discuss the core topics of this workbook.

Table C.1: Financial literacy study circles curriculum

Session	Topic
1	Introduction to study circle methodology
2	Risks
3	Risk management tools
4	Savings
5	Introduction to insurance
6	How does insurance work?
7	Different types of insurance products
8	How to submit a claim
9	How to find the best insurance product for you and your family
10	How, when, and why to renew?

Source: Microfinance Opportunities and Swedish Co-operative Centre and Vi Agroforestry, 2009, *Knowing your risks and how to manage them: A study guide*.

The financial literacy curriculum, outlined in Table C.1, covered a variety of topics related to the management of risks, with a focus on insurance. Materials were translated and included substantial use of illustrations for purposes of accessibility; moreover, the peer-to-peer learning model relied heavily on discussions that did not exclude those with limited literacy. It did not mention the *Bima ya Jamii* product by name, but it did cover several concepts of particular relevance to this type of product. Distinctions between inpatient and outpatient care, eligibility of claims and dependents, whether cash payments should be made in case of hospitalization, and processes for filing claims.

Attendance records maintained by Wananchi Delegates suggest reasonably high participation rates in the study circles curriculum. The average individual in our sample attended 75 percent of these sessions, and 96 percent of individuals attended at least half of all sessions. While these attendance rates are not independently verified, this is at least suggestive that low attendance alone is unlikely to explain the absence of impacts from this financial literacy intervention.

Figure C.1: Sample illustration from financial literacy materials

In-patient care is any care that requires hospital admission and overnight stay. For example a woman who goes to the hospital to deliver a baby receives in-patient care.



John needs to spend a night at the hospital; he is in - patient

Source: Microfinance Opportunities and Swedish Co-operative Centre and Vi Agroforestry, 2009, *Knowing your risks and how to manage them: A study guide*.

Appendix D Proof of Proposition 2 with scaled welfare differential

We modify Proposition 2 so that it applies to the effect of risk aversion on the utility differential scaled in the way proposed by Wilcox (2008, 2011), rather than to the unscaled differential:

Proposition 4. *For large values of R the scaled expected utility differential*

$$\frac{\Delta}{u(w) - u(w - \pi - c)}$$

is decreasing in R .

Proof. Denote $U = u(w) - u(w - \pi - c)$. As before, for large values of R we use the approximation

$$\Delta \cong \tilde{p}u(w - \pi - c) - pu(w - c).$$

Recall that

$$\frac{du(x)}{dR} = u(x) \left(\frac{1}{1 - R} - \ln x \right),$$

that $u(x) < 0$ for $R > 1$ and that

$$\tilde{p} = p(1 - q) \leq p.$$

The proposition is true iff

$$LHS = \Delta \frac{dU}{dR} - U \frac{d\Delta}{dR} > 0.$$

Note that

$$\begin{aligned} LHS &= [\tilde{p}u(w - \pi - c) - pu(w - c)] [-u(w) \ln(w) + u(w - \pi - c) \ln(w - \pi - c)] \\ &\quad - [u(w) - u(w - \pi - c)] [-\tilde{p}u(w - \pi - c) \ln(w - \pi - c) + pu(w - c) \ln(w - c)] \\ &= \tilde{p}u(w)u(w - \pi - c) [\ln(w - \pi - c) - \ln(w)] \\ &\quad + pu(w - c) [u(w) \ln(w) - u(w - \pi - c) \ln(w - \pi - c) \\ &\quad - u(w) \ln(w - c) + u(w - \pi - c) \ln(w - c)] \\ &\geq pu(w - c)u(w) [\ln(w) - \ln(w - c)] \\ &\quad + pu(w - c)u(w - \pi - c) [\ln(w - c) - \ln(w - \pi - c)] \\ &> 0 \end{aligned}$$

since the \ln function is strictly increasing. □