Nature's Frames, Reference Lotteries and Truly Risky Choice: Evidence from a Ugandan Field Lab

Paul Clist[*], Ben D'Exelle & Arjan Verschoor
School of International Development,
University of East Anglia,
UK

8th October 2013

Investment and insurance decisions are conceptually identical; both are choices between a riskier option that has a higher expected value and a less risky alternative that has a lower expected value. Crucially however, they differ in their default. This difference does not stem from an artificial construct but rather nature's own framing. By modelling the default as a reference lottery, we can apply Koszegi & Rabin's (2007) framework to the new domain of truly risky choice (where both prospect(s) *and* reference are non-degenerate lotteries). Using this theoretical framework we test the effect of altering the reference lottery, mimicking nature. We present a single risky choice problem in one of three ways: where the initial starting position of coins are relatively risky (the insurance treatment), relatively safe (the investment treatment) or neutral. We find significant treatment effects, despite the subtle difference in framing. In the second round of the experiment information on the socially most popular option is given, which we model as a competing reference. We find this has a large effect on behaviour, trumping the initial framing effect. Results from a Maximum Likelihood model are able to explain behaviour in both rounds, with surprisingly consistent parameter estimates recovered from apparently inconsistent behaviour. The experiment was conducted using 292 randomly selected participants from a rural farming community in Eastern Uganda, and results are discussed in the context of two puzzles in development economics: underinsurance and underinvestment in such communities. Recent successes in increasing low levels of insurance/investment are also discussed.

PRELIMINARY DRAFT: PLEASE DO NOT QUOTE

[*]paul.clist@uea.ac.uk

1 Introduction

Does the prior expectation of risk decrease risk aversion? Decision-makers operating in uncertain environments are, in one way or another, *forced* to gamble. If a decision problem is naturally framed as facing a risky situation - say, a pest that may destroy one's crops - does this make one mentally prepared to put up with more risk than when the starting point is relatively safe, for instance when one considers giving up a reasonably secure job for starting one's own business?

Despite the plausibility that they may, mainstream economic theory has until recently implied that such expectations and starting positions don't matter for risk attitudes. In expected utility theory, the utility of outcomes is unaffected by how they compare to prior expectations. In original prospect theory (Kahneman and Tversky, 1979), outcomes are compared to a deterministic reference point, which in its very many applications has usually been thought of as the current status quo, or sometimes the lagged status quo (e.g. Thaler and Johnson, 1990). These applications of prospect theory cannot capture the implications of stochastic reference points.

In the disappointment aversion models of Bell (1985), Loomes and Sugden (1986) and Gul (1991), outcomes are compared to the certainty equivalent value of a reference lottery, which is assumed to be the lottery that is chosen. These models thus allow for stochastic reference points but do not predict that features of the expectational environment matter, as in the examples above (cf. the evaluation of these models in Koszegi and Rabin, 2007). In particular, there is no reason in these models that a given lottery should yield a different certainty equivalent value when it is naturally suggested in the environment compared to when it is not.

In the reference-dependent utility theory of Koszegi and Rabin (2006, 2007), outcomes are evaluated relative to the full distribution of a reference lottery, which is itself conceptualised as recently held beliefs about outcomes. Because an (uncertain) outcome is compared to each outcome believed to be possible (with discrepancies weighted according to their probabilities), features of stochastic references matter for risk attitudes. Direct evidence in support of this theory, as far as its unique predictions about risk attitudes are concerned, is so far patchy[1]. One of its predictions is a so-called endowment effect for risk: that the certainty equivalent value of a lottery is larger when one owns it than when one considers buying it. This effect has been found in the lab experiments of Knetsch and Sinden (1984), Kachelmeier and Shehata (1992) and Sprenger (2010), the latter an explicit test of the reference-dependent preferences model of Koszegi and Rabin (2006, 2007) versus the disappointment aversion models of Bell (1985), Loomes and Sugden (1986) and Gul (1991).

The novel prediction of Koszegi and Rabin (2007), that the degree of uncertainty with which a decision problem is naturally framed matters for risk aversion, is yet to be tested. Two stylised examples that would ring true in the setting of our study can make clear the importance of testing this prediction. Consider two products in a rural area of a poor country in the semi-arid tropics, say somewhere in sub-Saharan Africa: index insurance and inorganic fertiliser. Index insurance is a relatively new product which pays out in the case of poor rains, compensating farmers for low yields. The reference in this case is non-purchase as farmers do not expect to purchase index insurance because they have, for many years, farmed without formal insurance against low rains. The reference of non-purchase has a wide spread, including the possibility of harvest failure due to drought, and the alternative offers a less widely distributed lottery.

By contrast, inorganic fertiliser is not a new product, but its use in sub-Saharan Africa is low relative to other continents (Morris et al., 2007). A typical semi-subsistence farmer grows enough to feed her family and sells a small surplus. The use of inorganic fertiliser offers a higher return for an initial cash outlay assuming good rains, an absence of pests and a good price at the time of harvest. If any of these shocks occur, her net profit may be lower than it would have been without inorganic fertiliser. Given the low levels of fertiliser use, the reference in this case is non-purchase of fertiliser. The key difference between the two examples is that in the insurance decision the reference lottery is more widely distributed than the lottery contemplated and vice versa in the investment (fertiliser) decision. Other things being equal, we would expect more risk aversion in the second example, if

[1] The theory's most impressive feat to date is to accommodate the within-day labour supply decisions of New York City taxi drivers that previous economic theory was unable to rationalise; see Crawford and Meng (2011) and references therein.

Koszegi and Rabin (2006, 2007) model is correct.

We test the prediction that, ceteris paribus, more widely distributed reference lotteries yield less risk aversion in an artefactual field experiment among 292 randomly selected participants from farming villages in eastern Uganda, using a between-subject design. We find that in the $(0.2, 10-x; 0.8, 10+x)$ lottery, where x is a freely chosen number of 500 shillings coins out of the set $0, 1, \ldots, 10$, a mean 5.0 instead of 6.4 coins are risked when the reference lottery is (0.2, 9; 0.8, 11) instead of (0.2, 1; 0.8, 19). The difference is statistically very highly significant. Empirical estimates from an ordered logit show that treatment effects are larger than the effect of well-known correlates of risk preferences, such as gender. Using Maximum Likelihood methods we are able to recover estimates of the model's structural parameters, which provide evidence for both loss aversion and the dual importance of consumption and gain-loss utility.

The expectational environment is not fixed; over time we would expect it to change and develop. Social effects are a plausible transmission mechanism[2] and so we test for this using a within-subject design. In a second round of the experiment, subjects were informed of the most popular choice in a parallel session (unbeknownst to them, other subjects may have faced a different reference lottery). We find social effects that are somewhat larger than the expectations-of-risk effect found in the main experiment. The propensity to change the number of coins risked in the two rounds, and the direction of any change, is strongly influenced by the social mode. A bivariate regression estimates convergence of 0.375 per unit of difference. Maximum Likelihood estimation of a structural model that includes the social mode as a competing reference point give surprisingly consistent parameter estimates.

As well as being the first experimental evidence for the influence on risk aversion of varying non-degenerate reference lotteries, our findings may help interpret low demand for products that would seem to be attractive to consumption-smoothing farmers operating in the hazardous environments of the semi-arid tropics, such as index insurance, new seeds and soil fertiliser. At the same time, our finding that social effects are larger than the expectations-of-risk effect highlights a possible avenue for policy-makers trying to encourage the uptake of such products.

The rest of the paper is structured as follows: section 2 presents the model and experimental design. Sections 3 and 4 contain the analysis of the first and second round decisions respectively. Section 5 discusses the results in the context of persistent findings of underinvestment and underinsurance in developing countries, and section 6 concludes.

2 Theory and Experimental Design

In Koszegi and Rabin (2006) a model of reference-dependent preferences is introduced in which utility consists of two parts: consumption utility and "gain-loss" utility. The former is subject to the assumption conventional in expected utility theory of diminishing marginal utility of wealth and the latter is thought of as derived from the difference between the consumption utility that corresponds with a reference and actual consumption utility. The model is applied in Koszegi and Rabin (2007) to study preferences over monetary risk and we follow the latter's notation. Utility u derived from wealth w is thus defined as

$$u(w|r) = m(w) + \mu(m(w) - m(r)), \tag{1}$$

where $m(w)$ is reference-independent consumption utility and the shape of μ, which converts deviations from a reference level of consumption utility into gain-loss utility, exhibits the loss aversion and diminishing sensitivity postulated by prospect theory. Probability weighting is abstracted from.

The reference point, given by r, is allowed to be stochastic. When a lottery F is evaluated, it is compared with reference lottery G:

$$U(F|G) = \int\int u(w|r)\mathrm{d}G(r)\mathrm{d}F(w) \tag{2}$$

[2]There is a body of evidence regarding how social networks influence people's risky choice, with several different conceptualisations (Conley and Udry, 2010; Bandiera and Rasul, 2006; Karlan et al., 2013; Cai, 2012). We have in mind the mechanism where another's choice may change one's one expecations about 'normal' behaviour in a given setting.

The evaluation of lottery F can thus be thought of as taking part in two stages. First, an outcome in F is compared with each outcome in G, and the utility discrepancies are averaged using the G outcome probabilities. This is done for each outcome in F separately. Next, the average utility discrepancies thus obtained for each outcome in F are averaged using the F outcome probabilities to yield the "expected utility" of lottery F.

The reference lottery G is conceptualised as recently held probabilistic beliefs about outcomes. For situations in which these are exogenous to the actual choice set, Koszegi and Rabin (2007) develop two propositions[3]. Proposition 1 states that a lottery evaluated relative to a deterministic reference level of wealth can be no less attractive when it is evaluated relative to a reference lottery, which implies an "endowment effect for risk" when the reference lottery is the evaluated lottery itself. Such an endowment effect, in which a lottery is valued more when owned, has been found in the lab (Knetsch and Sinden, 1984; Kachelmeier and Shehata, 1992; Sprenger, 2010).

Proposition 2 states that a sufficiently widely distributed reference lottery induces approximate risk neutrality. We test its implication that, ceteris paribus, more widely distributed reference lotteries decrease or not increase risk aversion. Formally, Koszegi and Rabin (2007, p.1057) show that for certain existing $A, \varepsilon > 0$, all constants k, any lottery F with positive expected value and any lottery H, $U(H + F|G) > U(H|G)$ when $Pr_G[r \in (k - A, k + A)] < \varepsilon$. An immediate corollary is that if lottery G' satisfies $Pr_{G'}[r \in (k - A - v, k + A + v)] < \varepsilon$ for $v > 0$ and the $A, \varepsilon > 0$ delimited in the proposition, then $U(H + F|G') \geq U(H + F|G)$. Intuitively, since the proposition states that for a sufficiently widely distributed reference lottery, F will be accepted, the decision-maker will be no less willing to accept F for an even wider distributed reference lottery.

Based on a pilot, we calibrated parameters as follows: $k = 10$ (500 shillings coins), $A = 1$, $v = 8$, and $\varepsilon > 0.8$. It follows that $G = (0.2, 9; 0.8, 11)$ and $G' = (0.2, 1; 0.8, 19)$. F is determined by subjects choosing x from the set $0, 1, \ldots, 10$ yielding lottery $F = (0.2, 10 - x; 0.8, 10 + x)$ for testing the hypothesis that a higher mean x will be chosen (subjects are less risk averse) for a more widely distributed reference lottery.

For estimating risk aversion, loss aversion and the relative weight of gain-loss utility, we proceed as follows. Our subjects evaluate two-outcome prospects $F = [x_1, q; x_2, 1-q]$ against reference lottery $G = [r_1, p; r_2, 1 - p]$ in which in our set-up (normally) $x_1 > r_1 > r_2 > x_2$ and $p = q$. If we assume "narrow bracketing" - both consumption utility and gain-loss utility refer exclusively to the experimental outcomes - then $u(x_i|r_j) = m(x_i) + \mu(m(x_i) - m(r_j))$ with $i = 1, 2$ and $j = 1, 2$. Plugging this into equation (2) yields

$$\begin{aligned} U(F|G) = \quad & q^2[m(x_1) + \mu(m(x_1) - m(r_1))] \\ & + (1-q)^2[m(x_2) + \mu(m(x_2) - m(r_2))] \\ & + q(1-q)[m(x_2) + \mu(m(x_2) - m(r_1))] \\ & + q(1-q)[m(x_1) + \mu(m(x_1) - m(r_2))] \end{aligned} \tag{3}$$

For estimation purposes, we assume that the shape of μ is identical in the gains and losses domain apart from loss aversion λ, so that $\mu(|y|) = -\lambda\mu(-|y|)$. For ready interpretation, we adopt Crawford and Meng's (2011) weight η of gain-loss utility and weight $(1 - \eta)$ of consumption utility (where $0 \leq \eta \leq 1$), which is a simple transformation of Koszegi and Rabin's (2007) weights. Noting that $x_i > 0$ and $r_j > 0$, we use $m(x) = \dfrac{x^{1-\alpha}}{1-\alpha}$, which for $\eta = 0$ (no gain-loss utility) amounts to the conventional CRRA specification of expected utility theory.

Adapting Equation (3) accordingly gives our estimation equation for $U(F|G)$.

[3]Beliefs being influenced by the decision-maker's reflection on the implications of her own strategies is at first abstracted from in Koszegi and Rabin (2007), but they are formed endogenously in the second part of their paper in recognition of their dependence on utility-maximising behaviour. This is clearly an important part of their theory but not our focus.

$$
\begin{aligned}
U(F|G) = \quad & q^2\left[(1-\eta)\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right)+\eta\left(\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right)-\left(\frac{r_1^{1-\alpha}}{1-\alpha}\right)\right)\right] \\
&+(1-q)^2\left[(1-\eta)\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right)+\eta\lambda\left(\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right)-\left(\frac{r_2^{1-\alpha}}{1-\alpha}\right)\right)\right] \\
&+q(1-q)\left[(1-\eta)\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right)+\eta\lambda\left(\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right)-\left(\frac{r_1^{1-\alpha}}{1-\alpha}\right)\right)\right] \\
&+q(1-q)\left[(1-\eta)\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right)+\eta\left(\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right)-\left(\frac{r_2^{1-\alpha}}{1-\alpha}\right)\right)\right]
\end{aligned}
\tag{4}
$$

The first line describes expecting r_1 but obtaining x_1, which will occur with a probability $q \cdot q$. In each line, there is a weighted average of consumption and gain-loss utility. Following Crawford and Meng (2011), we state that gain-loss utility receives a weight of $\eta \in [0, 1]$. The piecewise utility function states that where losses are present (i.e. the subject will receive less than expected), λ is included. We use the standard Constant Relative Risk Aversion (CRRA) formulation, which allows comparison to the existing literature.

Status Quo Bias

A word is required on the relation of reference lotteries to default/status quo bias, which has received compelling empirical support from hypothetical surveys (Samuelson and Zeckhauser, 1988; Ritov and Baron, 1992, 1995) and natural experiments (Madrian and Shea, 2001; Duflo and Saez, 2003). Status quo bias is almost exclusively discussed in a binary or ordinal choice set, where proof of its effect takes the form of a difference in the popularity of an option when it is the default option to when it is not (popular examples include organ donation and enrolment in a pension plan). We are only aware of two cases where continuous options are examined, both unincentivised atheoretical hypothetical questions reported in Samuelson and Zeckhauser (1988). Choi et al. (2003) is representative of the theoretical understanding of default bias, where it is explicitly conceptualised as a cost of switching from the default, a cost that does not depend on the distance between the default and the alternative.

Our focus on a continuous option set is also in contrast to previous investigations of KR. Koszegi and Rabin (2007) allow either a stochastic reference or a stochastic prospect, and Sprenger (2010) examines binary choices between degenerate and non-degenerate lotteries in an experimental setting. We apply the KR model to truly risky choice, i.e. where both prospect and reference are non-degenerate lotteries. This is a common domain outside of the lab, as it is rare that an option in an insurance or investment decision does not contain an element of risk.

Field Experiment

To examine the effect, if any, of changing the reference lottery we used an artefactual field experiment with 292 randomly selected participants from a predominately agricultural community in Eastern Uganda. Before the experiment began we explained that they would be asked to make two decisions, one of which (decided by a coin toss) would be played out for real. In the experiment each subject was endowed with 5,000 Ugandan Shillings in the form of ten 500 shilling coins, approximately a local daily wage. Each participant was asked to distribute the coins between two options, represented by a safe basket and a risky basket. Each coin placed in the safe basket meant a secure income, whereas each coin put in the risky basket offered a 80% chance of being doubled in value and a 20% chance of becoming worthless. Chance was determined at the end of the experiment by picking one counter out of a bag with 4 white and 1 green counters. If the subject picked the green counter all the coins in the risky option would become worthless. Table 1 shows the pay-off table of the different options, with Expected Utility Theory CRRA bounds.

To present the choice task, all subjects from a session approached a table where an experimenter explained the decision problem using the coins and two baskets, representing the safe and risky options. Subjects then returned to a waiting area before being recalled to answer control questions and make

Table 1: Pay-off table

Coins in safe option	Coins in risky option	Low Payoff p=0.2	High Payoff p=0.8	Risk Aversion Parameter (EUT) r where $v(x) = \dfrac{x^{1-r}}{1-r}$
10	0	10	10	$+\infty$ to 0.93
9	1	9	11	0.93 to 0.82
8	2	8	12	0.82 to 0.73
7	3	7	13	0.73 to 0.65
6	4	6	14	0.65 to 0.59
5	5	5	15	0.59 to 0.53
4	6	4	16	0.53 to 0.47
3	7	3	17	0.47 to 0.41
2	8	2	18	0.41 to 0.35
1	9	1	19	0.35 to 0.25
0	10	0	20	0.25 to $-\infty$

their decision in private. The same lottery game was presented in one of three ways to separate groups of participants: where inaction leads to a risky choice (insurance treatment); where inaction leads to a safe choice (investment treatment); and where inaction is not possible (neutral treatment). In particular, the three treatments differed in the starting position of the coins, both during the explanation of the game and when the subject approached the table to make their decision. In the insurance treatment 9 coins started in the risky basket and 1 in the safe basket. In the investment treatment the starting positions were reversed. In the neutral treatment 1 coin was placed in the safe and 1 in the risky basket, with 8 coins placed in between. Consequently, while the instructions given to the participants and the option set across the three treatments were the same, the only difference was the initial distribution of the coins between the baskets.

The elicitation procedure we use is closet to Gneezy and Potters (1997, hereafter GP), who frame a choice problem as an investment game where each subject is given an endowment X, which they may choose to keep $(X - x)$ or invest (x). The part invested has an uncertain yield such that the individual's total will be $X - x + kx$ with probability p and $X - x$ with probability $1 - p$ (see Charness and Gneezy, 2012, for a survey on this method). A related method used by, for example, Eckel and Grossman (2002, hereafter EG) is framed as a free choice between various binary lotteries which have a constant relation between proximate options. Conceptually the two approaches are identical; one presents the mechanism and the other the possible final options. To make this clear, we can state that in the EG lottery a subject may choose option 1 where they will receive X with certainty. Moving to the next riskiest option will result in $X - x + kx$ with probability p and $X - x$ with probability $1 - p$. We could present our specific experiment using the GP method by stating that $X = 10$, $p = 0.8$ and $k = 2$, or using the EG method as in table 1.

In the second half of the experiment, subjects were then presented with the same decision at a different table laid out in the same way. Two sessions ran in parallel with approximately 10 subjects per session, sharing a common waiting area but using a different table in the first round before switching for the second round. Before making their second round decision, they were told of the most popular decision on the new table in the previous round. This was done in a way that maximised the variation in the sample, by pairing treatments of different types. The experiment took circa two and a half hours, and average earnings were 9,000 Ugandan shillings once an unannounced show-up fee of 2,000 is included. The experiment was complemented by a questionnaire organized a few weeks before the experiment, which captured data on individual characteristics such as gender, wealth, education and risk preferences.

3 Analysis: First Round

The first round decisions are found in figure 1, with actual, summary and test statistics in tables 2-4. Table 5 provides summary statistics by treatment, which helps to check whether the randomisation resulted in a balanced groups in each treatment. There are three aspects to highlight. First, there is a substantial treatment effect despite the rather subtle difference between treatments. Those in the insurance treatment risk an average of 1.38 more coins than those in the investment treatment. A two-sample t-test for a difference in means shows that the investment treatment is significantly different from the other two treatments at the 1% level. The same test does not show a significant difference between neutral and insurance treatments[4]. It may appear remarkable that the Investment and Neutral treatments are significantly different, as one coin is risked in both cases. The difference must then stem from subjects having different expectations over the coins in the safe basket and those placed between the two baskets. Second, the observed behaviour is not consistent with inertia: only 2 people in the safe treatment and 12 people in the risky treatment make no change. The evidence suggests that the treatment affects the attractiveness of the entire choice set, and is not consistent with a binary disutility incurred from deviating from the status quo. Third, a simplifying heuristic appears to have played a role in subjects' choices. The most popular option, chosen by 27% of subjects, is to risk exactly half of the coins. It is also possible that the extreme options may have been more popular than they would have otherwise been[5].

Figure 1: Number of coins risked, by treatment

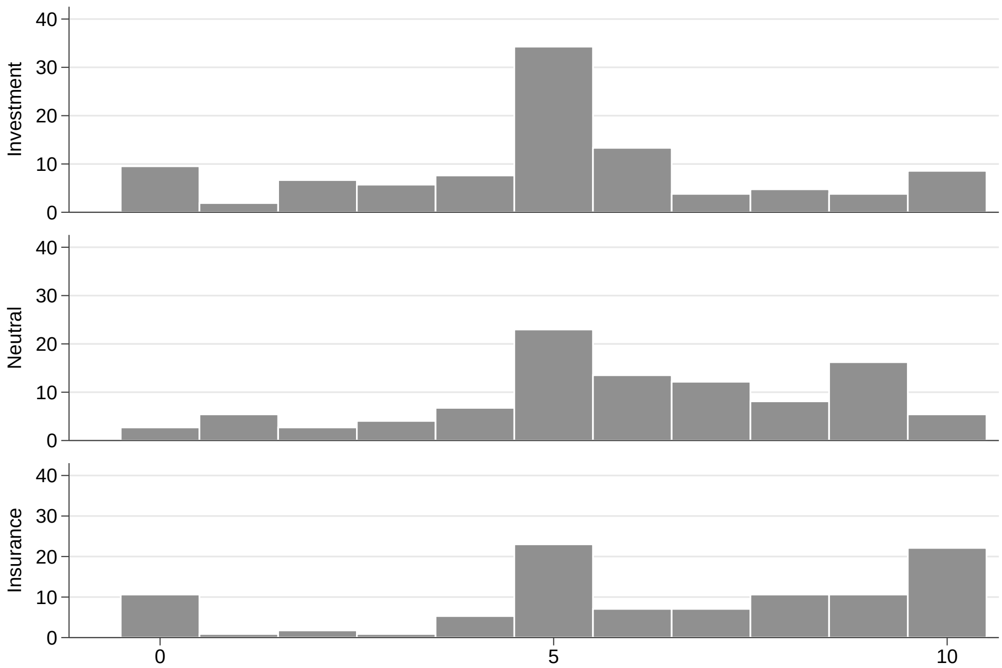

Note: Y scales are percentages. X axes are the number of coins risked in the first round

3.1 Analysis Using an Ordered Probit

Table 6 contains the results from a series of ordered logit regressions, providing parameter estimates for treatment effects with and without controls for household characteristics. We use a logit in recognition that a simplifying heuristic appears to have been at work, i.e. risking five coins was popular

[4]However, note the larger standard deviation for the risky treatment. Of the 12 subjects in the investment treatment to risk no coins in the first round, 5 come from a single session (12 sessions played the investment treatment). We are not aware of any extraneous factors and so have treated this as chance.

[5]This is the opposite of a censoring effect where the true latent preferences of subjects lay outside of the available decision space. Rather we note it is possible that the extremes gather more support that they would if there were a broader option set.

Table 2: Numbers of Coins Risked, by Treatment

Freq	Invest.	Neutral	Insur.	Total
0	10	2	12	24
1	2	4	1	7
2	7	2	2	11
3	6	3	1	10
4	8	5	6	19
5	36	17	26	79
6	14	10	8	32
7	4	9	8	21
8	5	6	12	23
9	4	12	12	28
10	9	4	25	38
Total	105	74	113	292

Table 3: Summary Statistics

Treatment	Mean	S.D.	N
Investment	4.99	2.67	105
Neutral	5.96	2.55	74
Insurance	6.37	3.13	113
Total	5.77	2.88	292

Table 4: T-Test for Difference in Means

Test	T Stat	P Value
Invest. = Insurance	3.49***	0.00
Invest. = Neutral	2.43***	0.00
Neutral= Insurance	0.94	0.17

Table 5: Summary Statistics, by Treatment

	Statistic	Safe	Neutral	Risky	Total
Female	Percent	49%	48%	42%	46%
Age	Mean (S.D)	42 (14)	41 (14)	39 (14)	41 (14)
Married	Percent	13%	18%	19%	16%
Catholic	Freq.	35	16	45	96
Anglican	Freq.	35	31	36	102
Muslim	Freq.	10	6	8	24
7th Day Adventist	Freq.	1	0	3	4
“Born Again”	Freq.	19	17	17	53
Other Protestant	Freq.	4	3	4	11
Total	Freq.	104	73	113	290

Note: Two participants took part in the experiment but did not take part in the survey. These two subjects are excluded from the Ordered Logit and Maximum Likelihood analysis, so as to ensure a constant number of observations across specifications.

partly because it was a simple choice. This decision appears borne out as the cut point estimates show the space between cut 5 and 6 (i.e. risking 5 coins) is larger than any other. In all specifications the estimated coefficients on treatment effects are significant, in that neutral and insurance effects are significantly different from investment treatment. However, in no specification are the estimated coefficients for neutral and insurance treatments significantly different from each other: Wald tests cannot reject the null hypothesis of equal coefficients, with p-values of 0.22, 0.27 and 0.29 for specifications 1, 2 and 3 respectively.

The size of the treatment effects is clear when compared to other household characteristics (specifications 2 and 3). The only significant household characteristic is that married subjects risk fewer coins. Women tend to risk fewer coins but the effect is not significant, and no parameter noting religious affiliation is significant. The coefficient on age is negative but rather small, as moving from the youngest participant to the oldest in the sample (from 18 to 70) implies a decrease of 0.45 units: smaller than the coefficient on being married.

3.2 Analysis in a KR framework

Two main types of risk parameter estimation are common. First, the bounds of an individual's risk parameter can be inferred from decisions, much as table 1 makes clear in the final column. This is normally restricted to a single parameter (though Tanaka et al. (2010) and Liu (2013) use an innovative method where a combination of choices can be used to infer the bounds of more than one

Table 6: Ordered Logit on Numbers of Coins Risked

	(1)	(2)	(3)		(1)	(2)	(3)
Neutral	0.64***	0.70***	0.69***	cut1	-1.96***	-2.39***	-2.47***
	(2.65)	(2.83)	(2.80)		(-8.51)	(-6.00)	(-5.46)
Insurance	0.96***	0.98***	0.97***	cut2	-1.71***	-2.14***	-2.22***
	(3.67)	(3.69)	(3.54)		(-8.16)	(-5.64)	(-5.11)
Female		-0.33	-0.24	cut3	-1.38***	-1.81***	-1.89***
		(-1.58)	(-1.06)		(-7.15)	(-4.91)	(-4.43)
Age		-0.0045	-0.0086	cut4	-1.12***	-1.54***	-1.61***
		(-0.65)	(-1.19)		(-6.05)	(-4.23)	(-3.83)
Married		-0.54*	-0.53*	cut5	-0.70***	-1.12***	-1.19***
		(-1.78)	(-1.69)		(-4.06)	(-3.17)	(-2.89)
Anglican			0.31	cut6	0.55***	0.15	0.095
			(1.18)		(3.22)	(0.44)	(0.24)
Muslim			0.034	cut7	1.03***	0.64*	0.59
			(0.10)		(5.83)	(1.89)	(1.47)
7th Day Adventist			0.84	cut8	1.37***	0.99***	0.95**
			(0.95)		(7.27)	(2.93)	(2.34)
"Born Again"			-0.38	cut9	1.79***	1.42***	1.38***
			(-1.05)		(8.83)	(4.19)	(3.37)
Other Protestant			0.34	cut10	2.48***	2.12***	2.08***
			(0.88)		(10.4)	(6.01)	(4.87)

Note: The parameter estimates are in relation to a Catholic subject in the Investment treatment. The dependent variable is simply the number of coins risked in the first round.

parameter, in their case loss aversion and value function parameters). An alternative method uses maximum likelihood to estimate a structural model of risk preferences (Harrison and Rutström, 2008), typically applied to data where subjects have made a series of choices between two lotteries. However, Crosetto and Filippin (2013) and Dave et al. (2010) show that other kinds of data can be transformed to resemble a set of binary lotteries. We follow this second method by representing a subject's choice not as 1 choice between 11 options, but rather 10 choices between 2 options: the one they chose was preferred to the 10 alternative options. This transformation means that the standard maximum likelihood method developed for binary lottery data can be used, and allows results to be compared to the existing literature.

Table 7 contains the parameter estimates obtained using ML. Column 1 reports a parsimonious model with only the three parameters used in the model included. The estimate of α is consistent with expectations, and is similar to other findings (e.g. Harrison et al., 2010, find a population wide estimate of α for their subjects from India, Ethiopia and Uganda of 0.46 using the same CRRA formulation). The estimate of η means that, in crude terms, consumption utility has a weight of 0.7 and gain-loss utility a weight of 0.3. The estimate of λ shows that there is loss aversion, though it is smaller than typically found. Other experiments tend to find a lambda of around 2, but these relate to quite different settings where the endowment and alternatives are different goods in a riskless setting. In column 2 demographic characteristics are also included in the estimation of alpha, such that the population wide estimates of men's alpha is estimated to be 0.5, while women's alpha is 0.58. The inclusion of gender, age and marital status characteristics do not have a great influence on the model's parameters, and the household characteristics themselves are not significant. However, the well known finding that women tend to be more risk averse than men (Eckel and Grossman, 2008) is replicated.

Columns 3 and 5 add a single household characteristic to the estimation of λ, and implies that women are more loss averse then men, significantly so in specification 5. Men's λ is estimated at 1.46 and 1.54 in specifications 3 and 5, with women's at 2.19 and 2.26. While differences in risk aversion by gender have been discussed at length (Eckel and Grossman, 2008), gender differences in

Table 7: KR Parameters, First Round Decision

	(1)	(2)	(3)	(4)	(5)
α	0.55***	0.50***	0.49***	0.45***	0.43***
	(14.75)	(6.00)	(5.84)	(5.35)	(5.22)
Female		0.08	0.09	0.05	0.07
		(1.37)	(1.27)	(1.00)	(1.11)
Age		-0.0001	-0.00002	0.001	0.001
		(0.03)	(0.01)	(0.64)	(0.70)
Married		0.11	0.11	0.13*	0.14*
		(1.35)	(1.39)	(1.77)	(1.81)
Anglican				-0.047	-0.05
				(0.81)	(0.85)
Muslim				0.27	0.02
				(0.26)	(0.24)
7th Day Adventist				-0.22*	-0.24*
				(1.77)	(1.84)
"Born Again"				0.16**	0.16**
				(2.03)	(2.06)
Other Protestant				-0.13	0.12
				(1.39)	(1.23)
η	0.30	0.33	0.49	0.35**	0.54***
	(0.66)	(1.19)	(1.30)	(2.12)	(2.80)
λ	1.30***	1.45***	1.46***	1.57***	1.54***
	(4.73)	(5.37)	(3.19)	(5.03)	(5.13)
Female			0.73		0.72*
			(1.20)		(1.78)

Note: Absolute Z statistics are shown in parenthesis below coefficient estimates. Standard errors and associated Z statistics are clustered at the individual level. Where religious affiliation is included, parameters should be interpreted relative to being Catholic.

loss aversion have not been found. For example, Tanaka et al. (2010) find men are less loss averse, but this is difference is never significant. However, Croson and Gneezy (2009) argue that women are more sensitive to social cues than men, accounting for the variability in gender differences across experiments. This would be consistent with the finding here, as women are apparently more sensitive to the initial frame.

The coefficient estimates in columns 3 and 5 couple higher average loss aversion for the population with a higher η. The estimates where gender differences in loss aversion are allowed see η jump from around 0.3 to around 0.5, implying a higher weight on the gain-loss element of the utility function.

Columns 4 and 5 include religious affiliation, where parameter estimates should be understood in relation to Catholic subjects. Two groups differ significantly from the others: Seventh Day Adventists are significantly less risk averse and Born Again protestants are significantly more risk averse. These extra parameters seem to allow a more precise identification of η, as they allow more variation to be controlled for.

4 Analysis: Second Round

Once subjects had made their first round choice they returned to a central waiting area. They were then told that they were to go to the other table to face the same decision problem, where they would be told of the most popular decision on that table in the previous round. This allows us to measure the effect of two references - one from the starting position of coins and the other from announcing the social mode.

What characterises people's behaviour in the second round? First, the majority risk a different number of coins in the two rounds; 115 people (39%) choose to risk the same number of coins in both rounds. Figure 2 plots these changes in coins risked on the x axis. The top panel is for people who initially risked more coins than the socially most popular option (the magnitude of this difference is not shown here), the middle panel for those that had risked the same number of coins as the most popular option and the bottom panel for people who risked fewer coins than the most popular option. The figure illustrates several points. First, while we cannot control completely for the effect of seeing the same decision problem a second time, the middle panel relates to the 51 subjects who agreed with the social mode. Of these, exactly two-thirds made the same choice in both rounds. By contrast, subjects that received a social mode that differed from their first choice were only consistent in around 30% of cases. This is clearly evidence of a social effect, as the propensity to change is influenced by one's peers. Second, figure 2 clearly shows that the direction of change is overwhelmingly towards the social signal as those in the top panel risk fewer coins, and those in the bottom panel risk more.

Figure 2 doesn't include information on the distance between the social signal and an individual's choice. For this analysis, figure 3 plots the difference between rounds on the x-axis and the difference between a subject's choice and the most popular choice in the first round on the y-axis. The line of best fit corresponds to a coefficient of -0.375 in a bivariate regression (see table 8), meaning that for each unit of difference between a subject's own choice and the social mode, we should expect convergence of 0.375. Thus for 8 units of difference (e.g. a subject risked one coin but the most popular option was 9), we should expect average convergence of 3 coins.

Figure 2: The difference between coins risked in the first and second round, by social signal

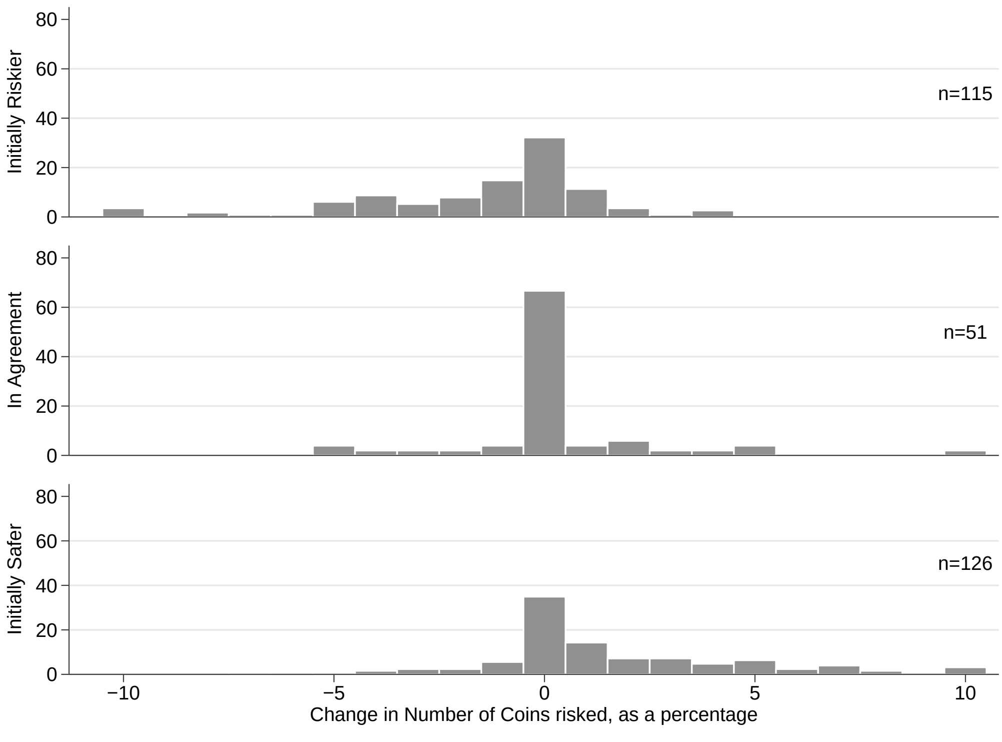

How should we interpret the size of the social mode effects? A simple comparison to the effect of the default lottery is informative, where the average difference in the number of coins risked between risky and safe treatments is 1.7. These two treatments differ by 8 coins. The social mode sees convergence of 0.375 per unit of difference, which over 8 units of difference translates to around 3 units of convergence. This is *prima facie* evidence that the social mode is stronger than the reference implied by the starting position of the coins.

Figure 3: Convergence to the Social Mode

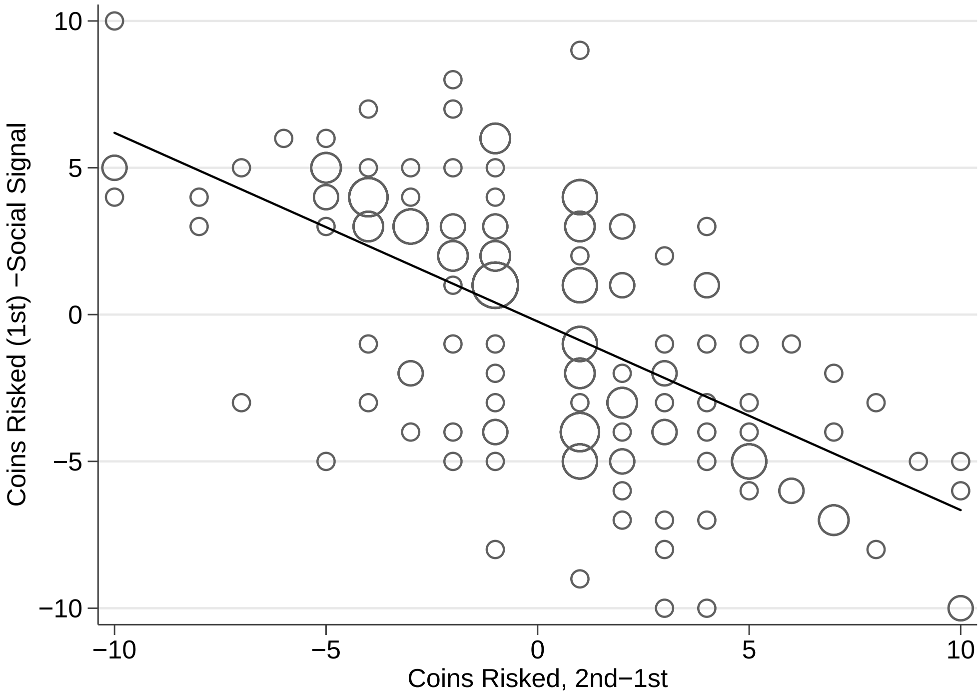

Note: those at zero on either axis are excluded, to aid readability. The size of the circle indicates the frequency of a particular data point.

Table 8: Dependent Variable: Change in coins risked

Variable	Coefficient (Std. Err.)
1st Decision - Social Signal	-0.375**
	(0.039)
Constant	0.058
	(0.163)

4.1 Analysis of second round, KR framework

More formally, we can conceptualise the social mode as a second reference point. Imagine a prospect $[x_1, q; x_2, 1-q]$, a first reference lottery $[r_1, q; r_2, 1-q]$ and a second reference lottery $[s_1, q; s_2, 1-q]$ drawn from S. We abstract from cases where the probabilities between the different reference points do not coincide, as this does not relate to our setting. For simplicity of exposition, assume $x_1 > r_1 > r_2 > x_2$ and $x_1 > s_1 > s_2 > x_2$ such that both reference lotteries have lower spread than the prospect. We introduce $\omega \in [0, 1]$ to describe is the weight attached to the (first) reference lottery G.

We then get to the following reduced form, where the first line describes getting x_1 when one was expecting r_1 and/or s_1.

$$
\begin{aligned}
U(F|G,S) = q^2 & \left[(1-\eta)\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right) + \eta\left(\omega\left(\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right) - \left(\frac{r_1^{1-\alpha}}{1-\alpha}\right)\right) + (1-\omega)\left(\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right) - \left(\frac{s_1^{1-\alpha}}{1-\alpha}\right)\right)\right)\right] \\
+(1-q)^2 & \left[(1-\eta)\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right) + \eta\lambda\left(\omega\left(\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right) - \left(\frac{r_2^{1-\alpha}}{1-\alpha}\right)\right) + (1-\omega)\left(\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right) - \left(\frac{s_2^{1-\alpha}}{1-\alpha}\right)\right)\right)\right] \\
+q(1-q) & \left[(1-\eta)\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right) + \eta\lambda\left(\omega\left(\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right) - \left(\frac{r_1^{1-\alpha}}{1-\alpha}\right)\right) + (1-\omega)\left(\left(\frac{x_2^{1-\alpha}}{1-\alpha}\right) - \left(\frac{s_1^{1-\alpha}}{1-\alpha}\right)\right)\right)\right] \\
+q(1-q) & \left[(1-\eta)\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right) + \eta\left(\omega\left(\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right) - \left(\frac{r_2^{1-\alpha}}{1-\alpha}\right)\right) + (1-\omega)\left(\left(\frac{x_1^{1-\alpha}}{1-\alpha}\right) - \left(\frac{s_2^{1-\alpha}}{1-\alpha}\right)\right)\right)\right]
\end{aligned}
\tag{5}
$$

Table 9: KR Parameters for Second Round Decision

	(1)	(2)	(3)
α	0.52***	0.49***	0.52***
	(15.84)	(5.76***)	(5.28)
Female		-0.001	0.02
		(0.08)	(0.27)
Age		0.002	0.0001
		(0.08)	(0.06)
Married		0.10	0.08
		(1.22)	(0.93)
Anglican			-0.05
			(0.76)
Muslim			0.83
			(0.65)
7th Day Adventist			-0.21
			(1.41)
"Born Again"			0.01
			(0.04)
Other Protestant			0.23***
			(2.77)
η	0.19	0.24	0.14
	(0.74)	(1.00)	(0.20)
λ	1.48**	1.55	0.87***
	(2.46)	(1.38)	(6.26)
Female			1.58**
			(2.40)
ω	0.17	0.26	0.00002
	(0.37)	(0.97)	(0.02)

Note: Absolute Z statistics are shown in parenthesis below coefficient estimates. Standard errors and associated Z statistics are clustered at the individual level. Where religious affiliation is included, parameters should be interpreted relative to being catholic.

Table 9 reports the parameter estimates for second round decisions, using the structural model outlined above. Results should be treated with some caution, as estimates are more fragile for the second round than they are for the first. In essence there are more parameters to recover but the same number of observations, and so this increased instability is unsurprising. This being said, columns 1 and 2 provide parameter estimates that are remarkably similar to these found in table 7 relating to first round decisions. The parameter on risk aversion is around 0.5, and lambda around 1.5. In both cases the weight given to the gain-loss utility has fallen in this second round. This hints that, in the face of competing references, consumption utility becomes more important.

Estimates of the new parameter in the model, ω, show that the social mode acts as a stronger reference than the starting position of the coins. This is particularly striking in column (3), where the weight given to the first reference tends to zero. In the estimation of ω it is constrained to lie between 0 and 1, and clearly the social mode dominates the gain-loss utility entirely in the third specification. The external validity of this result should not be overstated, as the relative strength of nature's framing and social signals will be highly context dependent. A further worry is that experimenter demand effects may lead us to overstate the importance of the social mode. However, it is clear that in this particular setting, the social mode has a very strong influence. A further reason for the size of this effect may be the complexity and unfamiliar decision environment. Apesteguia et al. (2007, p.217) states that "[i]mitation is prevalent in much of everyday decision making, in particular when the environment is complex or largely unknown": in rural Uganda such decision environments are neither well-known nor considered straight-forward.

Column 3 reports estimates when the estimate of λ is allowed to vary by gender. As in the previous round's decision the coefficient estimate finds that women are more loss averse than men. In this case, however, men's λ actually implies slight loss attraction (0.87). With women's λ estimated at 2.45, this means a much greater gender difference. The apparent importance of social cue's and framing for women's behaviour discussed elsewhere (Croson and Gneezy, 2009) chimes with the results here.

5 Discussion

In the first round of the experiment, we provide the first evidence that we are aware of that reference lotteries affect risky choice in a continuous option set. Changing the reference lottery has a significant and sizeable effect, despite the subtle difference in framing. There is, perhaps surprisingly, even a significant effect between the neutral treatment and the investment treatment, despite the reference lottery in each treatment risking only one coin. This difference must then come from different expectations over coins in the safe basket and those placed on the table between the two baskets. Ordered logit results show that the treatment effects are larger than well known correlates of risk preferences (e.g. gender). We also provide parameter estimates from a structural model, estimated using Maximum Likelihood. In order to recover these parameter estimates we extend the KR model to a new domain, truly risky choice, where both the reference and prospect(s) can be stochastic lotteries. These parameter estimates are surprisingly consistent across specification, with α (a CRRA risk aversion parameter) consistently around 0.5. The coefficient of loss aversion is also very consistent, at around 1.5. Estimates of the weight given to the gain-loss utility are between 0.3 and 0.54. Together, these consistent parameter estimates provide strong evidence of the importance of reference lotteries.

In the experimental literature, our paper is most closely related to Sprenger (2010), who showed that in cases where either the reference *or* the prospect is a stochastic lottery there is an endowment effect for risk. This was established through finding a significant gap in the willingness to pay for a lottery when endowed with a certain amount, and vice versa. By using a different experimental set-up, we are able to explore a continuous option set where risk is included in both reference lotteries. We show that changing the reference lottery affects the entire distribution of choices, finding evidence that the spread of the reference affects one's evaluation of the entire option set.

In the second round of the experiment, the social mode has a large effect of subject's choices. It affects the propensity to deviate from one's first round decision, and the direction of that movement. On average, subjects move 0.375 units towards to social mode for each unit of difference. In order to recover structural parameter estimates, we extend the KR model to incorporate two reference lotteries and model the social mode and default lottery as two competing reference points. It is striking that, despite the number of subjects that change the number of coins risked, parameter estimates from the second round are remarkably consistent with those in the first round. Loss aversion and risk aversion parameters are very similar at around 0.5 and 1.5 respectively. The weight given to the gain-loss utility is lower in the second round, dropping from a range of 0.3-0.5 in the first round to around 0.2 in the second. This is in line with expectation, as competing reference points dilute the signal of gains and losses, and increase the focus on consumption utility. The estimate on the relative weight of the two references shows that the social mode is much more important.

Second round estimates should be treated with some caution, as a greater list of caveats are needed. One of these caveats relates to estimation - a greater number of parameters are estimated using a similar amount of data. Another is that we cannot rule out an experimenter demand effect in the second round, as the social signal is explicitly announced to the subject, perhaps implying that they should respond to this new information. First round decisions do not suffer from this concern, as treatment effects are very subtle. Given this caveat, it is remarkable that consistent parameter estimates are recovered from apparently inconsistent behaviour.

A further innovation in this paper is to play the experimental games in rural Uganda, whereas previous related research was conducted in western labs. This decision was in part motivated by two common puzzles in development economics. The first puzzle is that of underinvestment, found in many developing countries. Duflo et al. (2008) offer a compelling empirical example from Western Kenya: investment in fertilizer has an expected annualised return of 69.5%, but only 37% of the sample report ever having used fertilizer. The second puzzle is that of underinsurance, typically index insurance products for rural farmers that pay-out in the case of poor rains (for a discussion see Karlan et al., 2013, p.37). It has been estimated that index insurance products are equivalent to an increase in consumption of almost 17%, yet take-up remains low (de Nicola, 2012). The most commonly repeated explanation is a lack of trust in or understanding of insurance products (Cai, 2012; Giné et al., 2008; Karlan et al., 2013). These studies typically refer to index insurance which pays local farmers when rains are below a certain threshold; a relatively new insurance product which requires a high level of numeracy to fully comprehend.

One factor has been mentioned in both insurance and investment contexts: risk preferences. Low levels of real-world investment has been linked to observed risk aversion in the lab: Liu (2013) shows that measures of risk aversion obtained from artefactual field experiments are a good predictor of adoption of Bt cotton amongst Chinese farmers. This new seed required a higher initial outlay but cut overall costs to the farmer, and it is easy to see why more risk-averse farmers would be slower to adopt a promising new technology. Low levels of insurance can conversely be explained by low levels of risk aversion. Zant (2008) simulates the demand for index insurance, and shows that demand for the product ranges between 5% and 30%, depending on the degree of risk aversion. However, it cannot simultaneously be the case that risk aversion is low enough to explain low levels of insurance behaviour, and high enough to explain low levels of investment behaviour.

We enter these debates with a simple observation: while insurance and investment decisions are conceptually identical (a choice between a higher paying but riskier option and a lower paying but safer alternative) they differ in their framing. In insurance decisions inaction leads to taking a riskier option, whereas non-investment leads to a safer alternative. This is not a superficial framing effect but rather a fundamental and innate aspect of investment and insurance decisions; nature's framing of investment and insurance decisions implies different reference lotteries even for decisions that are conceptually identical. Our extended KR model confirms the importance of reference lotteries on risky choice, with experimental evidence that is consistent with the model prediction that a reference lottery with a greater spread invokes more risk neutral behaviour. If our observation regarding nature's framing of investment and insurance decisions is valid, we would then expect underinvestment and underinsurance.

Recently, a number of papers have gone beyond offering explanations of underinsurance or underinvestment to offering or testing solutions. On the insurance side, Giné et al. (2008) hint that in India social effects may hold the key to combating low levels of trust and understanding, as a large proportion of people ask their friends before deciding whether or not to purchase the insurance product. Likewise, Cai (2012) provides evidence that an insurance education program in China is subject to substantial positive spillover effects from participants to their friends. Karlan et al. (2013) built on some of these findings in designing a policy intervention to increase demand for index insurance, focusing on increasing trust for insurance products. It included providing free insurance for a randomly selected subgroup of Ghanaian farmers and found a significant increase in the demand for insurance products once the intervention had finished, especially amongst those who knew someone who had received a pay-out from the insurance company. On the investment side, Duflo et al. (2011) argue that while many farmers plan to use fertilizer, procrastination means that more pressing concerns take precedence. They report evidence that ‘nudging’ farmers by marketing fertilizer with a 15% discount in the period directly after the sale of crops increases fertilizer purchase substantially[6].

The model and experimental results presented here offer a new framework through which to interpret these successes, as each intervention discussed above can be seen as introducing a new reference. Duflo et al. (2011) argue the increase in fertiliser purchase is related to combating procrastination. However, as discount vouchers were disseminated in public gatherings such as churches and schools, this may have changed people’s expectations of others’ behaviour and challenged the local norm of low fertiliser purchase. Karlan et al. (2013) argue that the endowment of insurance increases insurance take-up by increasing trust, as insurance take-up is increasing in payouts to the farmer themselves and their social network. However, it is also possible that by endowing insurance a new less-risky (lower spread) reference is created, invoking less-risky behaviour.

Our results do not rule out procrastination or low trust as reasons for underinsurance and underinvestment, rather they offer a complementary explanation. Furthermore, our results provide further evidence that social information information is a key channel for policy-makers who wish to combat underinsurance and underinvestment. Our results also show that these two puzzles are not unrelated.

6 Conclusion

We extend the KR model to a new domain, truly risky choice, where both reference and prospect(s) may be stochastic lotteries. In this setting we test KR’s prediction that a reference lottery with a higher spread invokes less risk averse (or at least not more risk loving) behaviour. Using several methods of analysis, our experimental results find strong supporting evidence for this prediction. Estimates of structural parameters are surprisingly stable across specifications and even the two rounds. Both consumption utility and gain-loss utility play their part in decision making, but in the case of competing reference points the weight on gain-loss utility diminishes. Consistent with previous research, we find evidence for loss aversion. However, we find that women are more loss averse than men, which may be related to the finding that women are more influenced by framing and social signals.

We discuss these results in the context of underinvestment and underinsurance in developing countries. We argue that while investment and insurance decisions are conceptually identical (both are choices between a riskier but higher paying option and a less risky but lower paying alternative) nature’s framing of them means that we should expect underinvestment and underinsurance. The experiment mimics nature’s framing of a risky choice problem, and finds that a subtle suggestion of a different reference lottery has a large influence.

Second round decisions, which allow for social influences, are analysed by extending the KR model further, to incorporate two references. In our specific experiment, the social mode exerts much larger influence than the subtle suggestion of different reference lotteries. This underlines the importance of using social information to combat low levels of investment and insurance, suggesting an avenue for policy-makers to explore.

References

Apesteguia, J., Huck, S., and Oechssler, J. (2007). Imitation - theory and experimental evidence. *Journal of Economic Theory*, 136(1):217–235.

Bandiera, O. and Rasul, I. (2006). Social networks and technology adoption in northern mozambique. *The Economic Journal*, 116(514):869–902.

Bell, D. E. (1985). Disappointment in Decision Making under Uncertainty. *Operations Research*, 33(1):1–27.

[6]Unfortunately they do not have evidence on fertilizer use, so it is possible that some fertilizer is bought for resale. The authors argue that the relatively small discount makes this unlikely.

Cai, J. (2012). Social Networks and the Decision to Insure: Evidence from Randomized Experiments in China. *University of California, Berkeley Mimeo.*

Charness, G. and Gneezy, U. (2012). Strong Evidence for Gender Differences in Risk Taking. *Journal of Economic Behavior & Organization*, 83(1):50–58.

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2003). Optimal defaults. *The American Economic Review*, 93(2):180–185.

Conley, T. and Udry, C. (2010). Learning about a new technology: Pineapple in Ghana. *The American Economic Review*, 100(1):35–69.

Crawford, V. P. and Meng, J. (2011). New York City Cab Drivers' Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income. *American Economic Review*, 101(5):1912–1932.

Crosetto, P. and Filippin, A. (2013). A Theoretical and Experimental Appraisal of Five Risk Elicitation Methods. *Jena Economic Research Papers*, 9.

Croson, R. and Gneezy, U. (2009). Gender Differences in Preferences. *Journal of Economic Literature*, 47(2):448–474.

Dave, C., Eckel, C. C., Johnson, C. A., and Rojas, C. (2010). Eliciting risk preferences: When is simple better? *Journal of Risk and Uncertainty*, 41(3):219–243.

de Nicola, F. (2012). The impact of weather insurance on consumption, investment, and welfare. *IFPRI Mimeo.*

Duflo, E., Kremer, M., and Robinson, J. (2008). How High Are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya. *American Economic Review: Papers & Proceedings*, 98(2):482–488.

Duflo, E., Kremer, M., and Robinson, J. (2011). Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. *American Economic Review*, 101(6):2350–2390.

Duflo, E. and Saez, E. (2003). The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment. *The Quarterly Journal of Economics*, 118(3):815–842.

Eckel, C. C. and Grossman, P. J. (2002). Sex differences and statistical stereotyping in attitudes toward financial risk. *Evolution and Human Behavior*, 23(4):281–295.

Eckel, C. C. and Grossman, P. J. (2008). Men, women and risk aversion: Experimental evidence. *Handbook of Experimental Economics Results*, 1:1061–1073.

Giné, X., Townsend, R., and Vickery, J. (2008). Patterns of rainfall insurance participation in rural India. *The World Bank Economic Review*, 22(3):539–566.

Gneezy, U. and Potters, J. (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, 112(2):631–645.

Gul, F. (1991). A Theory of Disappointment Aversion. *Econometrica*, 59(3):667–686.

Harrison, G. and Rutström, E. (2008). *Risk Aversion in Experiments*, chapter Risk Aversion in the Laboratory. Bingley, UK: Emerald, Research in Experimental Economics Volume 12.

Harrison, G. W., Humphrey, S. J., and Verschoor, A. (2010). Choice under uncertainty: evidence from Ethiopia, India and Uganda. *The Economic Journal*, 120(543):80–104.

Kachelmeier, S. J. and Shehata, M. (1992). Examining Risk Preferences under High Monetary Incentives: Experimental Evidence from the People's Republic of China. *American Economic Review*, 82(2):1120–1141.

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263–291.

Karlan, D., Osei, R., Osei-Akoto, I., and Udry, C. (2013). Agricultural Decisions after Relaxing Credit and Risk Constraints. *Mimeo, Yale University.*

Knetsch, J. L. and Sinden, J. A. (1984). Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value. *Quarterly Journal of Economics*, 99(3):507–521.

Koszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165.

Koszegi, B. and Rabin, M. (2007). Reference-dependent risk attitudes. *The American Economic*

Review, 97(4):1047–1073.

Liu, E. M. (2013). Time to Change What to Sow: Risk Preferences and Technology Adoption Decisions of Cotton Farmers in China. *The Review of Economics and Statistics*, forthcoming.

Loomes, G. and Sugden, R. (1986). Disappointment and Dynamic Consistency in Choice under Uncertainty. *Review of Economic Studies*, 53(2):271–282.

Madrian, B. and Shea, D. (2001). The Power of Suggestion: Inertia in 401 (k) Participation and Savings Behavior. *Quarterly Journal of Economics*, pages 1149–1187.

Morris, M., Kelly, V. A., Kopicki, R. J., and Byerlee, D. (2007). *Fertilizer Use in African Agriculture : Lessons Learned and Good Practice Guidelines*. Washington, DC: World Bank.

Ritov, I. and Baron, J. (1992). Status-quo and omission biases. *Journal of Risk and Uncertainty*, 5(1):49–61.

Ritov, I. and Baron, J. (1995). Outcome knowledge, regret, and omission bias. *Organizational Behavior and Human Decision Processes*, 64(2):119–127.

Samuelson, W. and Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of risk and uncertainty*, 1(1):7–59.

Sprenger, C. (2010). An Endowment Effect for Risk: Experimental Tests of Stochastic Reference Points. *Mimeo, San Diego: University of California*.

Tanaka, T., Camerer, C. F., and Nguyen, Q. (2010). Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam. *The American Economic Review*, 100(1):557–571.

Thaler, R. H. and Johnson, E. J. (1990). Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice. *Management Science*, 36(6):643–660.

Zant, W. (2008). Hot Stuff: Index Insurance for Indian Smallholder Pepper Growers. *World Development*, 36(9):1585–1606.