The London School of Economics and Political Science

Rainfall Index Insurance in India

Daniel Stein

Dissertation submitted to the Department of Economics for the degree of
Doctor of Philosophy in Economics at
The London School of Economics and Political Science

September 2011

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others.

The first chapter draws on field research that was carried out jointly with equal share by Shawn Cole, Jeremy Tobacman, and me. The data analysis and writing of the first chapter was performed primarily by myself, with contributions from Shawn Cole and Jeremy Tobacman. The second chapter was conceived and written solely by myself. The third chapter draws from a project designed with equal share by Jeremy Tobacman and myself. The project execution, data analysis, and writing of the third chapter were performed primarily by myself, with contributions from Jeremy Tobacman.

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

Abstract

This thesis provides three works which each contribute to understanding of the promising yet struggling market for rainfall index insurance in India. The first chapter contains an analysis of the willingness-to-pay (WTP) for rainfall insurance by poor farmers in Gujarat, India. It develops a theoretical model to predict individual WTP and contrasts it with emprical estimates of WTP using the Becker-DeGroot-Marshalk (BDM) mechanism. We find that BDM works well as a predictor of WTP, but that our model significantly overestimates WTP. The second chapter seeks to provide a possible explanation for demand being lower than theoretical predictions by looking at the dynamics of insurance demand. Using a panel dataset of insurance purchasers in India, it shows that people who receive an insurance payout are 9-22% more likely to purchase insurance the following year. The results are consistent with a dynamic model of insurance demand featuring loss aversion, in which receiving an insurance payout shifts the reference point such that people become more risk averse the following season. I provide evidence against other possible explanations, such as increased trust and learning about insurance, and direct effects of bad weather. The final chapter explores the possibility that combining rainfall insurance with savings may result in a more attractive financial product than insurance on its own. We conduct a laboratory experiment with Indian farmers that uses the BDM mechanism to assess the valuation of various insurance/savings combinations, which we title WISAs (Weather Insured Savings Accounts). We find that, contrary to theoretical predictions, most people prefer both pure savings and pure insurance to any combination of the two. This paper hopefully provides valuable contibutions to solving the puzzle of how to shield poor farmers from uncertain rainfall.

Acknowledgements

This work would not have been possible without the assistance of a vast array of people. I'd first like to thank my PhD supervisors at LSE, Tim Besley and Greg Fischer, who have provided invaluable advice and guidance throughout the dissertation process. Similarly, I owe a debt to my research collaborators Shawn Cole and Jeremy Tobacman, who invited me to join their umbrella of research projects in Ahmedabad, and shared valuable ideas, advice, and resources. During the 2010-2011 academic year I was a guest of the Center for Research in Economic Development at the University of Namur, and I thank the faculty (particularly Jean-Philippe Platteau, Gani Aldashev, and Catherine Gurikinger) for their helpful advice on my research.

I consider myself lucky to have received funding for all five years of my PhD. I thank LSE and Frederic and Elizabeth Sipiere for endowing me with a scholarship which allowed me to fund my studies at LSE for the first four years of the PhD. I also thank the AMID program of the European Union for the fellowship which funded the 5th year of my studies.

All three chapters of this thesis draw on data gathered from the field, which I certainly could not have managed all by myself. For Chapter 1, I first have to thank SEWA, the Ahmedabadi NGO whose partnership made our entire rainfall insurance study possible. They (especially Reema Nanavaty, Chhaya Bhavsar, and Nisha Shah) have provided incredible collaboration, direction, and resources for the project over the years. For the survey and marketing efforts that provided data for Chapter 1, I thank the entire team of the Centre for Microfinance in Ahmedabad, who run the project year round. Specifically, Nilesh Fernando, William Talbott, Sangita Vyas, and Dhruv Sood have been invaluable managers of the project over the years.

For the data in Chapter 2, I thank the microinsurance team at BASIX in Hyderabad, who were generous enough to share both their customer data and details of their insurance operation with me. Specifically, Sridhar Reddy and Venu Pratapani have been extremely helpful.

For Chapter 3, I thank my talented and hard-working team of facilitators who completed our laboratory sessions incredibly quickly and efficiently. I especially thank Maulik Chauhan, who made valuable contributions to all aspects of the field effort including translation, recruitment of subjects, recruitment of staff, and training. I also thank Neha Agarwal and Taroonkumar Madiwal from ICICI-Lombard who worked with me to create a custom insurance product for our study. The project was made possible thanks to a generous grant by the Microinsurance Innovation Facility of the International Labour Organization.

Many of my colleagues have helped with ideas for my research or by reviewing papers. While I am certain to forget people, I can at least thank Oliver Vanden Eynde, Michael Best, Kara Contreary, Tara Mitchell, Vincent Somville, Anders Fredriksson, and Petros Sekeris for their help.

Finally, I would like to thank my mother (Erica Karp) and stepfather (Wylie Crawford) for patiently reading each chapter of the thesis for spelling, grammar, and usage errors.

Contents

List of Figures

List of Tables

Preface

Roughly 60% of India's population is employed in agriculture, and over 50% of agricultural land is dependent on rainfall to nurture the crops.[1] But the Indian monsoon is notoriously unpredictable, prone to droughts and floods that can have devastating effects on the livelihood of rural Indians. While Townsend (1994) argues that Indian villages do an effective job of providing informal consumption insurance against idiosyncratic shocks, a poor monsoon will hit whole villages and districts at once, likely rendering intra-village transfers ineffective. Beginning in the early 2000s, rainfall index insurance was introduced in India as a potentially important tool to help poor farmers deal with rainfall risk (Hess, 2004; Skees et al., 2001), but insurance providers have struggled to reach a critical mass of customers, especially when offering unsubsidized policies.

This thesis attempts to contribute to the understanding of rainfall index insurance through three chapters that each offer a unique contribution. The first chapter adds to the measurement of demand by testing and implementing the Becker-Degroot-Marshak (BDM) mechanism in the field to empirically estimate willingness-to-pay (WTP) for rainfall insurance. The following two chapters look at two possible solutions for the low demand measured in the first chapter.

Previous studies have shown low take-up of rainfall index insurance in India, despite high theoretical benefit and widespread interest (Hill and Robles, 2010; Cole et al., 2010; Giné et al., 2008). Chapter 1 seeks to understand the structure of demand, specifically by estimating households' willingness to pay (WTP) for rainfall index insurance. We develop two approaches to estimating WTP, and evaluate them against an experiment in which fixed prices are randomly assigned. We believe that fixed prices provide a reasonable benchmark for "true" demand, as customers purchasing insurance in the marketplace would generally be presented with fixed price options.

Our first approach uses a simple structural model of index insurance demand that includes basis risk– the possibility that policy-holders may suffer a negative shock yet receive little or no payout. We use insurance policies, historical rainfall data, and survey data from members of an insurance pilot in Gujarat, India to fit the model and estimate each household's WTP for rainfall insurance coverage. Relative to the demand we observe at randomly assigned fixed prices, the structural model significantly overestimates demand. Our second approach uses a Becker-Degroot-Marschak (BDM) methodology to empirically elicit WTP from potential insurance customers at the time of marketing. We find that BDM does a better job of predicting fixed price purchasing behavior, but the distribution of stated willingness to pay has large mass at focal points. Finally, we directly compare the two approaches, and find the theoretical model has weak predictive power for WTP as elicited by BDM. We explore which household characteristics are correlated with WTP, and determine that recent experiences with rainfall and insurance are important factors not captured in our static model, suggesting that purchasing dynamics may be a promising

[1] CIA World Factbook: India (https://www.cia.gov/library/publications/the-world-factbook/geos/in.html); Indiastat.com

area for future analysis.

Following this lead, Chapter 2 looks at the dynamic nature of rainfall insurance purchasing decisions, specifically studying whether and why receiving an insurance payout induces a greater chance of purchasing insurance again the next year. I first develop a model that illuminates loss aversion as a possible mechanism for why receiving an insurance payout could spur repurchases in the following year. I theorize that receiving an insurance payout shifts the reference point in such a way that risk aversion, and therefore insurance demand, increase in the following period. This logic follows from Thaler and Johnson's (1990) work on "Gambling with House Money," where they find that gamblers who had recently won money become more risk loving.

I then test predictions of the theory using a uniquely-constructed panel data set of insurance customers from the Indian micro-finance institution BASIX. I find that receiving an insurance payout is associated with a 9-22% increased probability of purchasing insurance the following year, which corresponds to predictions of the model. I then test for other mechanisms not covered by the model that may explain insurance repurchasing. I first check whether direct effects of weather could be driving the results, and find that places that had experienced adverse weather in the year before insurance was offered actually had lower insurance demand than normal. This suggests that adverse weather itself is not driving increased insurance purchases. I next test for whether increased trust or learning could be driving insurance repurchasing, hypothesizing that if these channels were active they would result in spillover effects in the community. I do not find evidence that insurance payouts in a village drive new insurance purchasers, casting doubt on trust and learning as a valid explanation for the results.

In Chapter 3 we propose a new type of financial product that combines savings and rainfall insurance, called a Weather Insured Savings Account (WISA). Defining a WISA's type as the proportion of insurance versus savings it offers, we develop a model which shows that there should be an ideal WISA type for each person, and that the utility provided by the WISA will always decline as one moves away from this ideal type. We also show that (under reasonable conditions) people who are more risk averse or value the future more will have an ideal WISA type with more insurance.

We then conduct a laboratory experiment in Gujarat, India to measure farmers' risk aversion, time preference, and ideal WISA type. To measure the ideal WISA type we assess the willingness-to-accept (WTA) for various WISA types using a BDM mechanism. Contrary to predictions of our model, we find that a plurality of participants prefers both pure insurance and pure savings to any mixture of the two, and that this preference is most pronounced among those who are more risk averse. We present a number of possible explanations to try to square this result with our model, including that these results were driven by confusion over the insurance/saving mixtures or uncertainty over whether future payment would actually be received. We do not find behavior consistent with these explanations. One possible explanation for our results is that subjects exhibited diminishing sensitivity to losses as proposed by prospect theory. These findings suggest that the introduction of a WISA is unlikely to be successful.

This work exposes the struggles of the rainfall insurance market, and does not provide optimism for various suggested solutions. Overall, it suggests that rainfall index insurance is a struggling product that is not showing promising signs of improvement.

Chapter 1

What Is Rainfall Index Insurance Worth? A Comparison of Valuation Techniques[1]

1.1 Introduction

Rainfall index insurance is a microinsurance product designed to help farmers cope with the risk of uncertain rainfall. Its payouts are based not on individual outcomes of its customers, but instead on rainfall measured at a nearby "reference" weather station. This contract structure eliminates moral hazard, adverse selection, and costly claims adjustment, facilitating sale to small-scale farmers. Despite vast theoretical promise and extensive policy development, demand for rainfall index insurance has been low, especially when offered at market rates. Several years of field work with the NGO SEWA in Gujarat, along with a parallel study in Andhra Pradesh, have shown take-up of around 16% for market-priced insurance in India, despite intensive door-to-door marketing by trusted representatives (Cole et al., 2010; Giné et al., 2008). Giné et al. (2010) provide greater detail on the Indian rainfall insurance market.

This paper seeks to reconcile empirical findings of limited demand with an individually-calibrated structural model. Specifically, we develop a static model of index insurance demand that predicts willingness-to-pay (WTP) for a fixed amount of insurance coverage, given an individual's risk aversion and distance from the reference weather station. Our model contains a key insight highlighted by Clarke (2011), which is that the chance that the farmer could experience a shock but not receive a payout may reduce rainfall insurance demand by the most risk averse. We then perform three sets of tests.

First, we examine how well the model predicts observed insurance purchases at experimentally-manipulated fixed prices. Customers' decisions when presented with random fixed price offers provide useful benchmarks because they most closely reflect the real-world sales en-

[1] This chapter based on joint work with Shawn Cole and Jeremy Tobacman.

vironment. This test compares the percentage of the population that the model implies would have bought at a given fixed price to the percentage of people offered this price that actually purchased. If these percentages are the same, it indicates that the model is performing well at predicting WTP at the given price. In fact, we find that at each fixed price the model predicts a greater percentage of purchasers than we observe, indicating that the model is overestimating WTP.

Only two fixed prices were offered, limiting the resolution available for the first test. Consequently, second, we introduce and evaluate a methodology for obtaining higher-resolution empirical measures of WTP: a Becker-DeGroot-Marschak incentive-compatible mechanism (BDM). We implemented the BDM mechanism with 2,165 farmers for the opportunity to purchase real insurance policies. Using the same procedure we used to evaluate the model, we analyze the decisions made by people who received fixed prices to test the effectiveness of BDM in estimating WTP. (Subjects were randomly assigned BDM or an opportunity to buy at a fixed price.) We find that participants in the BDM exercise are very likely to express willingness to pay equal to a focal point (Rs. 50 or 100), making the comparison with fixed discounts difficult to interpret. However, when the fixed price corresponds to a focal point of the distribution, the WTP distribution elicited via BDM is consistent with purchasing behavior at fixed prices.

In our third test, we directly compare WTP predicted by the structural model to that measured using BDM. At the individual level we regress the WTP estimated using BDM (BDM bids) on the WTP predicted by the model (calculated WTP). A positive coefficient on calculated WTP would suggest that our model has predictive power in determining the BDM bids. In our full sample we find a positive correlation, showing that a one rupee increase in the calculated WTP is associated with an increase of Rs .27 in BDM bids, but this correlation is only significant at the 12% level. This indicates that the model has relatively weak power in predicting the BDM bids.

The remainder of the paper analyzes the strengths and limitations of the model and the BDM procedure in order to resolve the discrepancy between their implied WTP's. In response to recent papers exploring the risk aversion and insurance demand (Cole et al., 2010; Clarke, 2011; Bryan, 2010), we test how the relationship between risk aversion and insurance demand has evolved over time throughout our study. We find that while risk-averse people were less likely to purchase insurance at the beginning of the study (in 2006), by 2010 risk aversion was positively correlated with insurance demand, which corresponds with predictions of our model.

Finally, we examine other household characteristics that may be correlated with the BDM bid, hoping to gain insight into what other factors may influence WTP. We find that recent experiences with rainfall and insurance have significant correlations, suggesting that adding dynamic components of demand to our neoclassical model may be important.

This paper also makes a number of methodological contributions to the implementation of BDM in the field. First, it highlights the potential for focal points around round numbers in the distribution of WTP estimated by BDM. This suggests that researchers looking to test the effectiveness of BDM should make sure that their fixed price comparisons correspond to focal points of the BDM bid distribution. Next, we show that the outcomes of a "practice" BDM game, which teaches subjects how the game works, can affect their decisions in the 'real' game for insurance. As experiences in the practice BDM game (for a napkin) had strong effects on BDM bid for insurance, this suggests that researchers should use caution when teaching subjects about BDM.

This paper draws on a line of theoretical papers that attempt to explain low insurance takeup in the field. deNicola (2011) calibrates a dynamic infinite-horizon model, showing that basis risk, premium loading, and uninsurable background risk can lead to low insurance adoption. Cole et al. (2010) calibrate a simple neoclassical model and predict significant insurance demand for people with high risk aversion. On the other hand, Bryan (2010) uses a model of ambiguity aversion to show that people who are ambiguity averse will have demand for insurance decreasing in risk aversion. Clarke (2011) develops a model highlighting basis risk, showing that the possibility of not receiving a payout in the bad state of the world can reduce demand among the most risk averse individuals. Our model is closest in spirit to that of Clarke (2011).

As far as we know, ours is the first study to use BDM to study WTP for rainfall insurance. Perhaps the most closely related paper is Cole et al. (2010), which estimates demand elasticity for rainfall insurance using discount coupons, finding an elasticity between -.66 and -.88. The demand curve we estimate using BDM gives shows how the elasticity varies over a wider range of possible prices.

There have been relatively few field tests of the effectiveness of BDM as a methodology to assess WTP. While it is easy to show that the true statement of WTP is a dominant strategy for models of expected utility maximization (Becker et al., 1964), others have shown that BDM can give biased results if the expected utility framework does not hold (Horowitz, 2006a; Karni and Safra, 1987). Horowitz (2006b) provides a good overview of previous tests of BDM, along with some reasons for skepticism. Berry et al. (2011) test the effectiveness of BDM in the field by comparing WTP from BDM to demand elicited by fixed price offers for a water filter in Ghana, and find that BDM systematically underestimates WTP.

This paper will proceed as follows. In Section 1.2 we give an overview of the insurance products and data used in the experiment. In Section 1.3 we develop our model of insurance demand, and Section 1.4 presents benchmark tests of its predictions against insurance decisions at fixed prices. In Section 1.5 we discuss the implementation of BDM in the field, and test the predictions of BDM against insurance decisions at fixed prices. In Section 1.6 we directly compare WTP estimates from our model to those of BDM. Section 1.7 assesses reasons for the discrepancies between the model and the empirical measures of WTP. Section 1.8 concludes, and offers policy prescriptions based on the results.

TABLE 1.1: POLICY FOR ANAND TEHSIL IN ANAND DISTRICT (PAYOUTS DOUBLED TO REFLECT NABARD SUBSIDY)

WEATHER BASED CROP INSURANCE SCHEME (KHARIF 2010)
TERM SHEET

State GUJ Distrcit: Anand Tehsil: Anand

Crop:	Generic Crop	Reference Weather Station:	Anand	Unit:	PER ACRE

1. DEFICIT RAINFALL

			PHASE - I	PHASE - II	PHASE - III
1 A.	RAINFALL VOLUME	PERIOD	16-Jun to 15-Jul	16-Jul to 20-Aug	21-Aug to 30-Sep
		TRIGGER I (<)	80 mm	160 mm	60 mm
		TRIGGER II (<)	30 mm	75 mm	20 mm
		EXIT	0	0	0
		RATE I (Rs./ mm)	3	1.5	2.5
		RATE II (Rs./ mm)	24.16	9.96	20.00
		Max. Payout (Rs.)	875.00	875.00	500
		TOTAL PAYOUT (Rs.)	2250		

			PHASE - I	PHASE - II
2.	EXCESS RAINFALL (Multiple events)	PERIOD	1-Sep to 20-Sep	21-Sep to 10-Oct
		DAILY RAINFALL TRIGGER (>)	80 mm	60 mm
		EXIT (mm)	160 mm	120 mm
		Payout (Rs. / mm)	3.02	8.34
		Max. Payout	250	500
		TOTAL PAYOUT (Rs.)	750	

TOTAL SUM INSURED (Rs.)	3000
PREMIUM with S Tax (Rs.)	150
PREMIUM %	10.00%

1.2 Product and Data Description

1.2.1 Policy Explanation

Our local partner in this project is SEWA, an NGO based in Ahmedabad, India, that describes itself as "an organization of poor, self-employed women." Responding to concerns about rainfall risk from its rural members, SEWA piloted a rainfall insurance product in Patan in 2005, and began a broader offering of rainfall insurance to households in three districts (Ahmedabad, Anand, and Patan) during the summer (kharif) growing season in 2006.

This study uses data from 2010, when SEWA offered a five-phase rainfall insurance policy underwritten by the Agricultural Insurance Company of India (AICI) to its members. The first three phases of the policy provide coverage against deficit rainfall, while the final two phases provide coverage against excess rainfall as heavy rainfall or storms can damage crops near harvest time. The policy terms as provided in AICI's termsheet are included here as Table 1.1.

SEWA offered policies linked to 14 different rainfall stations. The policies were all priced the same (Rs 150), but gave slightly different terms due to different historical rainfall. They all followed the same general structure as the example given above. The deficit phases of coverage offer piecewise-linear payouts based on the cumulative amount of rainfall within the specified timeframe. If this cumulative amount is below Trigger I (II), the policy pays out the difference between actual rainfall and Trigger I (II) times Rate I (II). (Note that when rainfall is below Trigger II, the customer is also paid [TriggerI-TriggerII]*Rate I.)

The two excess phases pay out if rainfall on any single day within the coverage period exceeds the trigger threshold. Figure 1.1 shows the payout structure for the first phase of

Figure 1.1: Payout Scheme for Phase 1, Anand Tehsil, Anand District

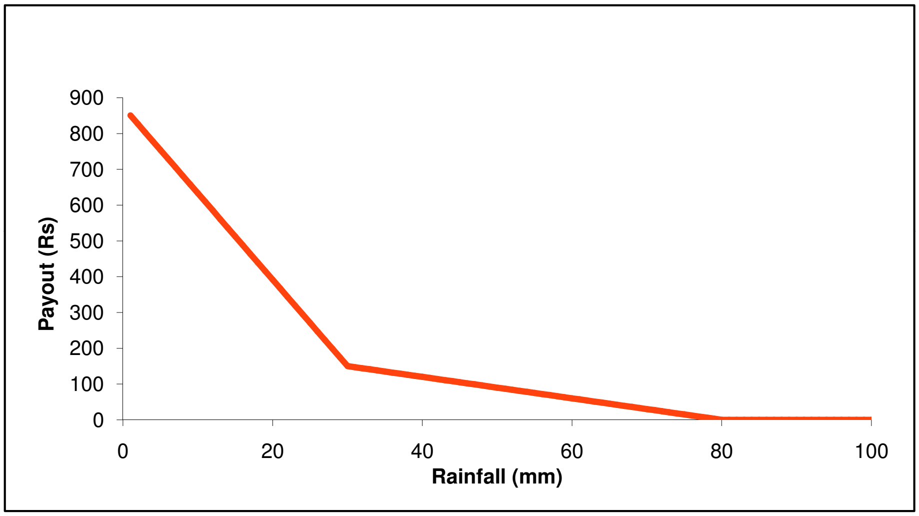

the insurance policy for Anand Tehsil.

While they vary somewhat based on the weather station, the policies offer coverage that is roughly actuarially fair, meaning the expected value of insurance payouts equals the premium paid. This favorable pricing was due to a subsidy from the government of India's National Bank for Agriculture and Rural Development (NABARD). NABARD offered to match premiums paid by farmers, which effectively doubled payouts from the original policies offered by AICI. When selling the policies, SEWA chose to market the subsidy as a "Buy One Get One Free" promotion to its members. SEWA explained that anyone who purchased a policy (either at full price or as a result of the BDM game) would instead be awarded two policies, effectively doubling coverage.

Insurance policies are written for a certain policy holder only, and are not transferrable. While an informal secondary market for the insurance policies could technically exist, we have never witnessed any evidence of this.

1.2.2 Data

The data in this study comes from household surveys conducted with SEWA members from 2006-2010, and also from data collected during insurance marketing efforts in 2010. In 2010, SEWA marketed insurance to around 3,351 households in 60 villages. We can divide this sample into two groups: the sample of people who received household surveys, and non-surveyed households.

Since 2006, we have conducted annual household surveys with 750 of these households. One third of the surveyed households were selected randomly from SEWA's membership rolls, while the other two thirds were identified by SEWA as people who may be interested in rainfall insurance. In 2009, we added an additional 8 villages to the study, surveying and visiting 50 households per village (all of whom were suggested by SEWA.) Survey data is used to calculate risk aversion parameters for participants, calibrate constants in the

Table 1.2: Summary Statistics

BDM bid / Total price	0.593	Irrigation spending (Rs '0000)†	0.039
	(0.24)		(0.142)
Bought insurance	0.62	Uses HYV seeds	0.335
	(0.486)		(0.472)
Total Monthly expenditure (Rs '0000)†	6.101	Experienced Drought in Previous Yr	0.268
	(4.588)		(0.443)
Experience with SEWA insurance	0.192	Food adequacy	0.058
	(0.394)		(0.234)
Experience with Gov't crop insurance	0.069	Rainfall last year	3.640
	(0.254)		(1.43)
Outstanding credit (Rs '0000)†	3.48	Basis risk	10.865
	(5.006)		(5.073)
Main income own agriculture	0.165	Financial literacy	0.617
	(0.371)		(0.233)
Main income agricultural labor	0.203	Risk Aversion	1.331
	(0.402)		(2.362)
Input spending (Rs '0000)†	0.212	Discount factor	0.772
	(0.455)		(0.165)

Standard Deviations in Parentheses † Windsorized at 1% upper tail

theoretical model, and to correlate BDM bids with household characteristics. All surveyed households were given the opportunity to play the BDM game.

Most household data used in this paper is taken from the survey conducted in early 2010. One exception is the measure of risk aversion, as questions pertaining to these subjects were only asked in the first year customers were surveyed (which is either 2006 or 2009). Table 1.2 presents summary statistics for our surveyed population.

The non-surveyed households were additional households suggested by SEWA that would be good candidates for rainfall insurance. As the surveyed and non-surveyed populations were selected differently and also have received different marketing efforts in the past, the two populations potentially have different underlying insurance demand. Both surveyed and non-surveyed households were used to populate a marketing list, which directed SEWA's marketing efforts.

1.2.3 Insurance Marketing Strategy

Insurance policies were marketed to 60 villages from May-June 2010 by SEWA. The marketing began with a village meeting to which all SEWA members were invited, which explained the concept of rainfall insurance and the policies that would be offered. In the meetings attendees were given the opportunity to discuss the policies and ask questions of the SEWA representatives.

Following completion of the village meetings, the SEWA marketing team returned to each village to conduct household-level marketing visits. They focused on reaching people on the pre-specified marketing list. Each person on the marketing list received a household visit by a member of SEWA's marketing team, during which they received an explanation of the insurance policy, viewed a video about rainfall insurance on a handheld player, and received additional marketing flyers with randomly assigned marketing messages.

In addition, each household was given a preprinted scratch card that enabled the client to either receive fixed policy discounts or play the BDM game. The participant's name was printed on the scratch card, and only they could use it. Participants first scratched off the top panel of the scratch card, which revealed whether they received an offer of a fixed discount or an offer to play the BDM game.

If the household was selected to play the BDM game, they were then asked to state the maximum amount of money they would be willing to pay for insurance (their bid). They then scratched off another panel on the card which revealed their random offer price. If the offer price was below their bid they purchased the policy for the offer price. If the offer price was higher than their bid, there was no sale. Further explanation of the BDM procedure is given in Section 1.5.

Surveyed and non-surveyed households were treated with discount arms in different proportions. To maximize power for tests involving household characteristics, only BDM games (for 1 and 4 policies) were assigned to surveyed households. To maximize power for evaluating the BDM methodology, either BDM games (for 1 and 4 policies) or fixed-price discounts were randomly assigned to 1,035 non-surveyed households identified by SEWA as potentially interested in rainfall insurance. The fixed discounts resulted in final prices for a single insurance policy of Rs 130 or Rs 100. Note that we use purchasing data from people given fixed discounts to validate estimates of WTP elicited via BDM and our model. An image of the scratch card used to conduct the randomization is given in Appendix Figure A1.4.

Of the 3,351 people visited in 2010, 2,165 filled out the scratch cards. Table 1.3 outlines the various discounts and games offered.

1.3 Structural Approach

1.3.1 Models of insurance Demand

In this section, we construct and calibrate a model of demand for rainfall insurance that captures the key features of the farmer's problem.

Classic theories of insurance demand (Schlesinger, 2000; Borch, 1990) generally focus on traditional indemnity insurance, in which insurance payouts are a function of financial loss. These models predict full insurance coverage for risk-averse individuals when insurance is priced at actuarially fair rates, and at least some coverage when insurance is more expensive. These models do not match the observed low take-up rates of index insurance.

Standard models of indemnity insurance omit a key feature of index insurance: basis risk. Basis risk is the possibility that the insurance may not pay out even though the customer has experienced a loss (or if the insurance pays out even though no loss occurs.) This happens if the weather on the farmer's land differs from that at the reference weather station or if a farmer experiences crop failure for any other reason (e.g., pest). Basis risk is an important limitation of index insurance as compared to traditional indemnity insurance, and may be an important aspect of a model of index insurance demand.

Table 1.3: Discounts Offered

	Surveyed	Nonsurveyed
All Discounts		
Visited	1158	2193
Played Scratch Card	865	1300
BDM Game for 4 Policies		
Number Played	410	295
Number Won	326	190
Number Bought	294	167
Average Bid (Rs.)	297.6	280
	(137.1)	(139.8)
BDM Game for 1 Policy		
Number Played	448	345
Number Won	420	310
Number Bought	387	286
Average Bid (Rs.)	102.8	105.2
	(31.0)	(38.1)
Fixed Price Rs 100		
Number Scratched	0	327
Number Bought	0	232
Fixed Price Rs 130		
Number Scratched	0	314
Number Bought	0	182
Standard Deviations in Parentheses		

In the following model we allow for basis risk by assuming that the rainfall which produces crop input is not the same as the rainfall used to calculate insurance payouts; they are instead related by a bivariate lognormal distribution, where the correlation between the two variables determines the basis risk.

We allow each village to have its own measure of basis risk, which increases with the distance to the reference weather station. We also allow the coefficient of partial risk aversion to vary for each individual, as we have estimates of risk aversion from experimental lotteries conducted during the survey. These two factors allow us to generate individual-level estimates of WTP for insurance.

1.3.2 Basic Model Structure

We construct a simple model of insurance demand to determine how much an individual would be willing to pay for a fixed amount of insurance coverage. An individual has fixed income Y, but is also subject to a random income shock S. The individual can purchase an insurance policy at price P which gives a payout M as a function of the shock. The premium P^* that satisfies Equation 1.1 sets the expected utility from purchasing insurance equal to the expected utility from not purchasing insurance, representing the maximum WTP.

$$E[u(Y - P^* - S + M)] = E[u(Y - S)] \tag{1.1}$$

Timing is as follows:

1. Customer makes insurance purchase decision.
2. Income shock S and payout M are realized.
3. Final wealth is consumed.

We make a significant assumption that liquidity constraints do not affect WTP. Although (Cole et al., 2010) indicates that liquidity constraints may play a role in demand for rainfall index insurance, a parallel study to this one showed a limited role. In this study we randomly offered some customers the ability to pay their insurance premium after the harvest (when liquidity constraints are lower) as opposed to the normal time of before the growing season. Very few customers offered this option elected to take it, and the WTP for people offered these premium loans was not significantly different from those who had to pay right away. Based on these results, we do not include liquidity constraints in this model.

In the next section we calibrate this model by developing a structure for the shocks, payouts, and utility function.

1.3.3 Calibration

The main challenge in adapting the simple model above to our situation is to develop a structure for both income shocks and insurance payouts. We use historical rainfall, crop models, and the actual insurance policies used in Gujarat to develop such a framework.

SEWA offered insurance contracts for 14 different rainfall stations in 2010, but we have varying amounts of historical data for each station. We have a particularly long data series (44 years) from the weather station in Anand city due to data collection by the Anand Agricultural University. We therefore calibrate the model using Anand's historic rainfall data and its corresponding insurance policy.

We estimate crop losses using an adaptation of the Food and Agriculture Organization's crop water satisfaction model (Cole and Tufano, 2007; Bentvelsen and Branscheid, 1986). In this model, crop losses are proportional to the percentage evapotranspiration[2] deficit from the maximum evapotranspiration by the crop-specific yield response factor K_y. We proxy for evapotranspiration with rainfall, and define the shock S as follows:

$$S = \mathbf{1}(R < R_{max})K_y(1 - \frac{R}{R_{max}})Y_n \tag{1.2}$$

[2]Evapotranspiration is the sum of water evaporating from a surface (evaporation) and water vapor being released by a plant (transpiration). Transpiration is directly related to the amount of water absorbed by a plant, but in practice, it is generally difficult to measure the two effects separately. Therefore evapotranspiration is used as a proxy to measure water intake by the plant. When crops receive all their water from rainfall (as is the case with most of our sample population), evapotranspiration will be closely related to rainfall.

Figure 1.2: Actual versus Fitted Rainfall over first three phases of the Monsoon. Anand District 1965-2003, fitted to lognormal distribution.

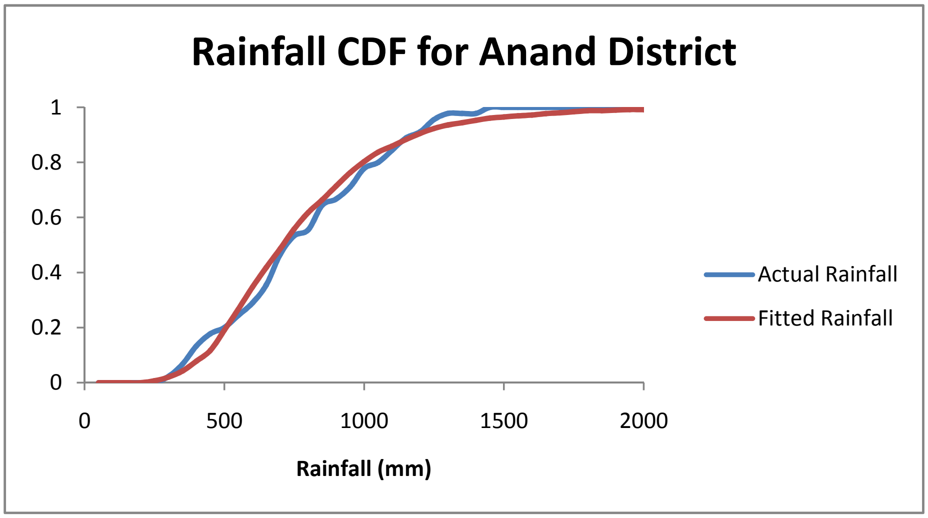

Where R is rainfall and R_{max} corresponds to the 90th percentile of the rainfall distribution, which we assume is the rainfall threshold below which crop losses begin to occur.[3] We assume that the shock is proportional to Y_n, which represents the maximum level of income that can be lost due to a shock. We set Y_n at 10,500, which is equal to the average yearly difference between income in a good rainfall year versus a bad rainfall year as self-reported by our farmers.

We assume that the relevant R used to calculate income shock due to drought is the cumulative rainfall over the period of time when our insurance policies offer drought coverage. Following Cole and Tufano (2007), we assume that this rainfall follows a lognormal distribution. The parameters of both variables are set to fit the historical rainfall distribution in Anand district over the beginning of the monsoon (when drought coverage was offered), giving a location parameter of 6.57 and a scale parameter of .41. A Kolmogorov-Smirnov test cannot reject the equality of distribution between actual rainfall and our fitted lognormal distribution. Figure 1.2 plots the cumulative distribution function for both historical rainfall and our lognormal approximation.

The main crops grown by farmers in our sample are millet and sorghum. While we do not have yield response data for millet, the FAO estimates the K_y coefficient of crop loss for sorghum to be around .9 over the entire growing season (Bentvelsen and Branscheid, 1986). We therefore use a value of .9 for K_y.

The rainfall insurance policies sold in Gujarat in 2010 were quite complicated, consisting of three phases of drought coverage and two phases of coverage against single days of particularly heavy rains. We estimate insurance payouts using a simpler scheme with one phase of drought coverage covering the time period of drought coverage on the actual policies. This simplification costs us the opportunity to correctly analyze situations where

[3]Our results are not sensitive to this assumption.

overall rainfall in a season is normal, but the distribution is heavily skewed, affecting crops and also triggering insurance payouts. However, we feel the gain in simplicity from a one-phase policy is worth this sacrifice.

While the actual insurance policies sold in Gujarat in 2010 varied based on location, they all had roughly the same structure. For drought coverage, linear payouts are based on the difference between cumulative rainfall over a phase and two defined "triggers". When rainfall falls below the first trigger, the policy pays out a small payment for each millimeter of rainfall below the trigger. When rainfall falls below the second trigger, recipients receive payouts per millimeter that are much higher than deficits below the first trigger. In Anand Tehsil, rainfall has historically fallen below the first trigger 41% of the time and hit the second trigger 15% of the time. We use these thresholds of the estimated rainfall distribution to create the payout structure, with payment per millimeter below the second trigger being seven times the payment per millimeter below the first trigger. As the policies in 2010 were roughly actuarially fair, we set the payout amounts such that in expectation the payout equals the premium of Rs 150.[4] This corresponds to a payout of Rs 1.11 for each millimeter below the first threshold and Rs 7.77 for each millimeter below the second threshold.

Insurance payouts are based on rainfall at a local rainfall station, which may be different from rainfall R that farmers experience in their fields. We denote the rainfall used to calculate the insurance payout as R_s, and to provide for basis risk we draw R and R_s from a bivariate lognormal distribution. It is worth noting that this choice deviates from the structure of basis risk used by Cole et al. (2010) and Clarke (2011). Cole et al. (2010) assume that the two shocks (the equivalent of R and R_s) are different due to an additive, independent, mean-zero normal error term. Importantly, this structure creates very few situations where there is a bad shock yet no payout, minimizing the importance of basis risk. The model in Clarke (2011) has a constant probability that the insurance will not give a payout even when there has been an income shock. This creates many situations where there is a bad shock yet there is no payout. The difference in this structure determines why Cole et al. find insurance demand increasing with risk aversion, while Clarke finds demand to be either uniformly decreasing or increasing then decreasing in risk aversion. Our bivariate normal structure presents a strategy that is somewhat in between in terms of the number of times where it creates a bad shock but low insurance payout. But overall, it creates a structure of basis risk closer in spirit to that of Clarke (2011). We think that this functional form provides a flexible and plausible method for introducing basis risk into our model.[5]

Empirically, we have examined the relationship between correlations of daily rainfall realizations at the 15 GSDMA weather stations in our study area, and the distance between

[4] As is standard in the insurance literature, this definition of "actuarially fair" does not take into account time preference.

[5] We have also run the model with different functional assumptions of basis risk. If we assume basis risk similar to that of Cole et al. (2010), demand for insurance is much higher than in our model, especially for people with high levels of risk aversion. If we assume a binary form of basis risk as in Clarke (2011), demand is lower among the most risk averse, but the results are not fundumentally different.

Table 1.4: Model Calibration

Risk Exposure	
Fixed Income (Y)	41800
Maximium Loss (Yn)	10500
Crop Factor (Ky)	0.9
Rainfall Distribution (Lognormal)	
Location Parameter	6.568
Scale Parameter	0.405
Shock Distribution	
Average shock	5224
Stdev of shock	1694
Policy Characteristics	
Probability Payout > 0	41%
Payout Rs/mm after First Trigger	1.1
Probability of Reaching Second Trigger	15%
Payout Rs/mm after Second Trigger	7.7
Average payout	150
Basis Risk	
Correlation between money shock and payout shock	
Average	0.76
Standard Deviation	0.11

those weather stations. The correlations fall with distance, as expected, and the linear fit between the correlations and distance has an R-squared of 0.62. In our model, we adopt this linear structure of rainfall correlation, and assign a level of correlation to each village based on its distance from its reference weather station. For the typical distance of 10 km between a study village and its weather station, the predicted correlation in daily rainfall is 0.65.

We assume that people have CRRA utility with coefficient of relative risk aversion ϕ. Utility as a function of consumption c is given as:

$$U(c) = \frac{c^{1-\phi}}{1-\phi} \tag{1.3}$$

Survey enumerators played Binswanger (1981) lotteries with subjects for real money, which allows us to estimate ϕ for each respondent. Given that the amount of money in the games is relatively low compared to subjects' total wealth, a simple calculation of the CRRA parameter would give unreasonably high values (the no-risk value of the lottery is Rs 25, or around \$.50.) Therefore, we follow Binswanger (1981) and estimate the partial risk coefficient, and use this as an estimate of the coefficient of relative risk aversion. This assumption gives a range of values consistent with empirical estimates of risk aversion (Halek and Eisenhauer [2001] have a good summary.) More detail about estimation of these risk coefficients and comparisons to other estimates are given in Appendix Table A1.1.

Certain income Y is set equal to the average level of yearly nonfarm income according to our survey, which is Rs. 41800. Table 1.4 outlines the calibration of various constants in the model.

To calculate the willingness to pay for insurance, we numerically solve Equation 1.1 for P^* for each person who played the BDM game. In the next section we present the results of this model.

1.4 Test of Model Predictions against Fixed Prices

1.4.1 Benchmark Tests of the Structural Model

In this section we compare predictions of the model against purchasing decisions made by people who faced fixed prices for a single policy, which we consider a reasonably reliable indicator of true WTP.[6] This can only be thought of as an illustrative test, however, since the population which received fixed-price offers does not overlap with the population of those for whom we have survey data. Most importantly, the households for which we have survey data are more likely to have received insurance marketing visits over the past years, which could influence their WTP. However, we believe this is still a useful exercise, as the sampling frame for the surveyed households and the households receiving the marketing price are roughly similar.

Since everyone for whom we have survey information was offered the BDM game, we cannot use the model to calibrate demand for those offered fixed prices. But we can still use aggregate statistics to provide a rough test of how well the model mirrors true purchasing decisions. We compare the percentage of people who purchased insurance at a fixed price to the percentage of people whom the model predicts would have a WTP above the fixed price. If these two percentages are equal, it is an indication that the model is accurately predicting WTP, at least around that fixed price.

When we visited households in the field, we delivered the opportunity to purchase insurance for reduced prices (or play the BDM game) via a scratch card. Some households (around 1/3) refused to scratch off their card, generally due to complete lack of interest in insurance. A reasonable assumption is that these households' true WTP was below any of our fixed prices and therefore they would not have bought even if they had scratched off the card. We report results for both the sample of just people who scratched off a card and also the full visited sample, assuming that these people would not have purchased at either of our fixed prices.

The comparison between model predictions and fixed price purchasing is presented graphically in Figure 1.3. The solid line in the graph is the demand curve predicted by the model, showing the percentage of people we expect to purchase at each price.[7] The

[6]While we can use our model to predict WTP for any amount of insurance, we only have fixed price data for purchases of single policies, so we use predicted WTP for one policy for the comparison.

[7]Note that the demand curve is generated for the entire surveyed population that was visited to market insurance. We could generate a second demand curve just for people who agreed to play the scratch card to provide clearer comparison to the "Played Card" results for fixed prices, but this demand curve is virtually

Figure 1.3: Comparison of Model Predictions to Purchasing at Fixed Prices

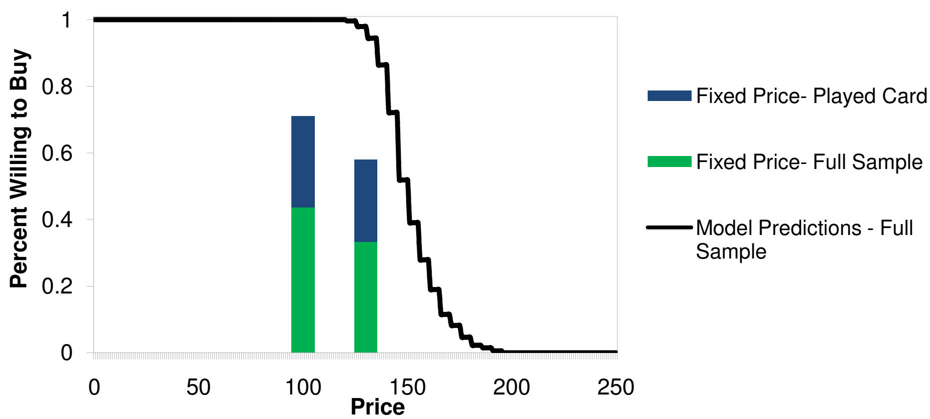

Table 1.5: Model Predictions Compared to Fixed Prices

	Played Card Only	All Visited
	(1)	(2)
Bought at Price of Rs 100	71.00%	43.61%
Number offered Fixed Price	327	532
Model Predicts WTP >= Rs 100	100.00%	100.00%
Number Modeled	396	538
Bought at Price of Rs 130	57.96%	33.21%
Number of Customers	314	551
Model Predicts WTP >= Rs130	98.48%	98.30%
Number of Customers	396	538
Note That the Policy Price is Rs 150		

columns represent the actual proportion of purchasers at fixed prices of Rs 100 and Rs 130. The dark columns include the whole sample, while the light columns restrict the analysis to only people who filled out the cards. The graph clearly shows that the model overestimates the amount of purchasers at all fixed prices.

In Table 1.5 we present numerically the comparisons between the model predictions and purchasing at fixed prices. In Column 1 our sample of fixed price purchasers is everyone who scratched off their scratch card to reveal a fixed price (which would either be Rs 100 or Rs 130). We see that the model predicts near-universal takeup- 100% at a price of Rs 100 and 98.48% at a price of Rs 130, while actual takeup was 71% and 58% respectively. In Column 2 we include all people who were visited, even if they refused to fill out a scratch card. Here we see that the model still predicts near-universal takeup, while the actual takeup is 43% and 33%. These results verify the fact that our model is overestimating WTP. While the surveyed and fixed-price populations are different, it is unlikely these

indistinguishable from that of the full sample.

differences are driving a wide gap in WTP between the model and observed behavior. We therefore conclude that the model severely overestimates WTP.

1.4.2 Structural Model Sensitivity Tests

In this section, we test the sensitivity of the model to shed light on which parameters may be responsible for the inaccurately high estimates of insurance demand. Charts and figures related to these tests can be found in Appendix Section 1.9.2.

We consider three key factors that could affect model predictions: expectations about payouts, risk exposure, and basis risk.

We used all available data, 44 years, to characterize the rainfall distribution. However, much of the policy value derives from extreme events, which are by definition rare. It is quite possible that people had varying beliefs about the probability of the payout. A farmer who believes the expected payout to be significantly lower will have a lower WTP for the product. Given our 44 years of data, we can put bounds on our estimate of the expected value of the insurance, and see how we would expect WTP to change for different beliefs withing these boundaries. If a farmer believes the insurance is not actuarially fair, and instead has a loading factor in the range of 21-42% (which correspond to one and two standards deviations below our estimate of expected payouts), predicted WTP lines up more closely with fixed price behavior.

A second factor affecting demand is the degree of risk exposure. We model this as the ratio of wealth susceptible to loss due to a rainfall shock, set to roughly .2, based on self-reported loss exposure by farmers. Since this may be a noisy measure, we consider alternative ratios, from 1 (which means a farmer risks losing all wealth) to .1. Ratios below .2 do not have much effect on insurance demand. Increasing this ratio does increase demand, but at high levels of risk aversion the prospect of a total loss not covered by insurance can decrease insurance demand (this is one of the central conclusions of Clarke [2011]).

Finally, the amount of basis risk present can affect insurance demand. In the Appendix we present results from the model with a range of correlation between shocks and payouts, including the endpoints of 0 and 1. While lower basis risk does lead to higher insurance demand, this effect plays out mostly for those with the highest risk aversion.

The factor that seems to have the most potential to square our predicted WTP with observed behavior is the belief about average payouts; possibly the customers did not believe that this insurance was in fact actuarially fair. We will discuss other ways to possibly improve the model's predictions in Section 7.

1.5 BDM and Fixed Discounts in the Field

1.5.1 Explanation of BDM Implementation

As mentioned before, the opportunity to play the BDM game was determined using a scratch card, which was given to all households visited for insurance marketing. If the participant received the chance to play the BDM game, the SEWA team then explained how the game worked. The steps of the BDM game (using the game for 1 policy as an example) are as follows:

1. Participant states the maximum amount they are willing to pay for the insurance policy. This "bid" is recorded by the facilitator.

2. Participant scratches off the random "offer" price from the scratch card.

3. If the offer is less than the bid, then the participant purchases the insurance at the offer price. If the offer is greater than the bid, the participant cannot purchase policies during that marketing visit, though she or he is free to purchase the insurance at full price through an agent or SEWA sales team member at another time.

The participant first practiced by playing the BDM for a SEWA napkin, which had a market value of Rs 10. The napkin game was resolved on the spot to show exactly how the game worked. Then they played for insurance. After stating their bid, participants were reminded that bid above the offer price was an agreement to purchase insurance, and that if the bid was below the offer price there would be no sale. In order to make sure that the BDM bid did not capture short-term liquidity fluctuations, participants were told that if they didn't have the money to purchase insurance on the same day, a SEWA representative would return in two weeks to complete the sale if they won the game. Before scratching off the offer, participants had a chance to adjust their bid, but once the offer was scratched off they could no longer change their bids.

The distribution of BDM offers was skewed towards low prices, as we wanted many people to win the game and end up with rainfall insurance.[8] The range of the offers was between 0 and 150 (the market price for insurance), and the probability density function of the offer prices was: $Density = 2 - 2 * (\frac{OfferPrice}{Premium})$. We told participants the range of the offer prices, but not its distribution.[9]

[8] This is because the BDM game acts as an instrument for take-up for an impact evaluation of insurance, to be described in a future paper.

[9] We make the top of the distribution equal to the offer price due to the fact that regulations prevent us from selling insurance above the market price, and evidence in Bohm et al. (1997), which shows that BDM performs better when the upper bound on the offer distribution is the market price. While we didn't address the distribution of offers with the subjects, experiments in Mazar et al. (2010) suggest that exposure to different price distributions will change a subject's stated WTP.

1.5.2 Theoretical Concerns about BDM

As mentioned earlier, various authors have put forth concerns about the validity of BDM, especially if participants have preferences that cannot be expressed by expected utility (Horowitz, 2006a; Karni and Safra, 1987). Karni and Safra emphasize that BDM can give incorrect valuations for lotteries, which would include insurance contracts. While different classes of preferences cause the WTP estimated by BDM to be either upward or downward biased, it is instructive to think how a specific deviation from expected utility could affect BDM bids.

We take Karni and Safra's example of using probability weights as in prospect theory, and consider how this class of preferences could affect a BDM bid for index insurance. If participants play BDM for a lottery with a particularly important low probability event (we will call it a "catastrophe"), the probability of this event will be lower when playing the BDM game because there is a chance that BDM will result in the lottery not being offered at all. If a participant overweighs low probability events, playing BDM for this lottery makes the subjective probability of the catastrophe higher relative to its actual probability. One catastrophic event that may weigh on a participant's mind in the case of rainfall index insurance is basis risk, or more specifically that there would be a bad rainfall shock yet would receive a small or nonexistent payout. With overweighting of small probabilities, the BDM game would magnify the effect of this negative event on decision making, tending to cause BDM to underestimate WTP.

While this is a legitimate concern, we will be able to look for evidence of this effect by comparing WTP as measured by BDM to behavior at fixed prices. If there was systematic underestimation of WTP using BDM, the above criticism might be playing a role. However, this does not correspond with the patterns we observe.

1.5.3 Tests of BDM Implementation

In April and May 2010, SEWA visited 3,351 households in 60 villages, of which 2,268 were assigned to play the BDM game. In a large implementation like this, there are likely to be some errors in the field. These worries are magnified by the fact that explaining and implementing the BDM game is somewhat complicated, and there may be opportunity for collusion between the facilitator and player of the BDM game. Fortunately, we can use the data collected to test the validity of the BDM implementation. In this section we present an overview of possible concerns and data either supporting or rejecting these worries.

There are a handful of specific things that we thought could have affected our field implementation. First, some scratch cards may have been lost, and if this was correlated with the outcome of the BDM game it could potentially bias our results. Next, people may have scratched the cards before they recorded their final offer price. There is also the worry the people may "win" the game by scratching off a bid lower than their offer but then decline to actually purchase the policy. Finally, people may be influenced by the test BDM game for the napkin.

We take each of these concerns in turn. Data analysis relating to these issues is available in Appendix Section 1.9.3.

- **Censoring of Cards**: In order to check for censoring of cards, we can check whether the distribution of the BDM offers from people who played the cards in the field is the same as the distribution generated on all the cards. This does seem to be the case, as the equivalence of the distributions cannot be rejected by a number of statistical tests. We therefore think that censoring of cards was not a large issue.

- **Scratching Cards Before Stating WTP**: If people saw the offer price before they made their bid (and it affected the bid), we should see a correlation between BDM bids and offers. Unfortunately we do see this, indicating some lapses in implementation in the field. However, the result is a bit puzzling, as we see a positive correlation in two districts and a negative correlation in the third. (Each of the three districts in our sample had different marketing teams.) The positive correlation could make sense for a few reasons. First, people who had offers less than their bids but then decided they didn't want to purchase the policies may have lowered their bids after the fact. Next, people who had offers higher than their bids may have decided that they did want to purchase the policy at that offer price, and therefore raised their bids. Finally, it is possible that people simply viewed the offer price before they made their bid, making it a type of price anchor. The negative correlation in the single district is difficult to justify.

- **Refusal to Purchase Policy**: When someone scratches off an offer price below their bid, they are technically required to purchase the insurance at the offer price. However, around 10% of the people who won the game refused to purchase the insurance. This most likely arose due to the fact that the ability to purchase insurance is affected by liquidity constraints, which may not be well known at the time of making their bid. Respondents had two weeks to come up with the money, and some may not have been able to collect sufficient funds to purchase the policy. Most respondents who refused to purchase the insurance after winning the game claimed they did not have the money available.

- **Insurance Bid Influenced By Napkin Game**: We played the BDM game with each respondent first for a napkin to show how the BDM game works. In theory this should have no effect on WTP for rainfall insurance, but we do find it affects the BDM bid. A 1 Rs increase in the price offered to purchase the napkin (which is revealed by scratching the card) correlates with an increase in the BDM Bid (expressed as percentage of premium) by 1.5 percentage points. (The standard error of this estimate is .42%) The mere fact of winning the napkin game may also have a strong negative effect on the BDM bid. This result seems to indicate a misunderstanding of how BDM works, as maybe people thought that they could achieve a better outcome in the insurance game by taking the results of the napkin game into account.

While there were clearly some irregularities in our implementation of BDM, it is difficult to understand what it means for our interpretation of the BDM bids. There is no rational

Figure 1.4: Comparing BDM Bids to Fixed Price Decisions for Non-Surveyed Population

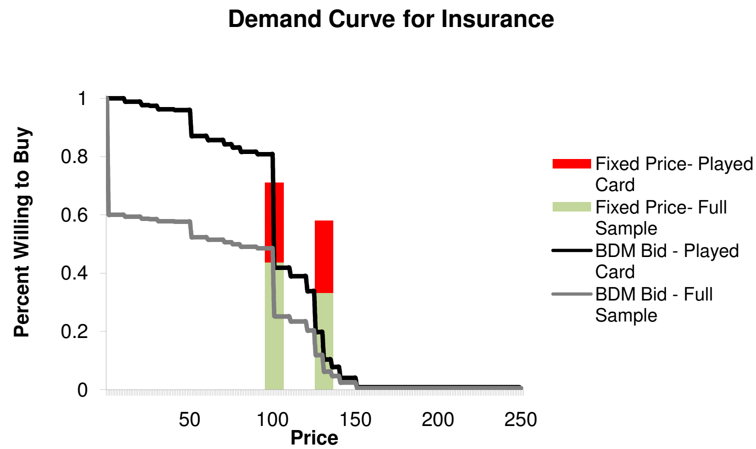

reason that someone should change their BDM bid after viewing the BDM offer unless viewing the offer somehow changes their preferences. Similarly, experience with the unrelated napkin game should not affect preferences over rainfall insurance. This behavior likely reveals more subtle clues as to the nature of WTP. Perhaps it is somewhat misleading to think that people have an intrinsic, unchanging WTP, and simply seeing the price that they could have paid changes their demand for the product.

Despite these doubts, we consider testing against fixed price demand to be the best test of BDM validity, which we do in the next section. While our results are somewhat mixed, they indicate overall that BDM gives an accurate measure of WTP.

1.5.4 Test of BDM against Fixed Prices

If BDM is eliciting the true WTP, then the percentage of people who have a WTP over a certain threshold should be the same as the percentage who purchased when offered a corresponding fixed price. Participants who were on our list but were not previously surveyed randomly received either the opportunity to play the BDM game or fixed prices of Rs 100 or Rs 130 for one policy. We can therefore compare the decisions among these two groups to assess the validity of BDM. Note that although we played the BDM game for one or four policies, we only offered varying fixed prices for single policies. Therefore, we just use the single policy results for this comparison.

We offer a graphical comparison of the two demand measures in Figure 1.4, which plots the demand curve for insurance as predicted by BDM and also demand as observed at the fixed price points. The dark demand curve reflects the demand as a percentage of everyone who filled out a scratch card. The lighter curve assumes that everyone who did not fill out a card had a WTP of zero, and includes the full sample. The two columns for fixed price purchasing have analogous definitions. The dark bars restrict the sample only to those who filled our cards, while the light bars include the full sample.

Figure 1.4 shows that demand predicted by BDM is close to fixed price demand at a

Table 1.6: Comparison of BDM to Fixed Discounts

Sample is Non-Surveyed Customers in 2010		
	Played Game Only	Full Sample
	(1)	(2)
Bought at Price of Rs 100	71.00%	43.61%
Number offered Fixed Discount	327	532
Bid >= Rs100	81.75%	48.54%
Number Played BDM Game	345	581
Bought at Price of Rs 130	57.96%	33.21%
Number offered Fixed Discount	314	551
Bid >= Rs130	20.00%	11.88%
Number Played BDM Game	345	581

price of RS 100, but is much lower at a fixed price of Rs 130. In Table 1.6 we directly compare the demand at these prices.

In Column 1 we consider the entire population of people who filled out scratch cards. Here we see that while 81.75% of people who played the BDM game bid greater than or equal to Rs 100, only 71% of people purchased at a price of Rs 100. This suggests that BDM is overestimating the true WTP. Column 1 also makes the same comparison with people who received a fixed price of Rs 130. In this comparison, BDM performs far more poorly, with only 20% of people bidding Rs 130 or more, while 58% of people purchased when offered a fixed price of Rs 130. This suggests that BDM is underestimating the true WTP.

In Column 2 we assume that households who refused to play the scratch card game would not have purchased if offered any discount, and also would have bid less than Rs 100 if they had agreed to play the BDM game. While adding this group to the analysis mechanically makes the BDM correspond more closely to fixed discounts, omitting the group arguably improperly censors people with low insurance demand. In this analysis we see that 43.6% of people bought at a fixed price of Rs 100, while 48.5% gave BDM bids greater than or equal to Rs 100. This comparison is much closer than in Column 1, but still suggests that BDM is overestimating WTP. The comparison with prices of Rs 130 also improves compared to Column 1, but still suggests that BDM underestimates WTP.

To get a more quantitative comparison of BDM and purchasing behavior from fixed discounts we can adopt a regression framework akin to that of Berry et al. (2011). To do this we create a dummy variable that takes a value of 1 if the participant was assigned the BDM game and their bid was greater than or equal to the fixed discount threshold or they were assigned a fixed discount and purchased insurance. We regress this dummy on a variable that takes a value of 1 if the participant was assigned the BDM game and zero if they were assigned a fixed price. A positive coefficient means that BDM gives a higher value of WTP than you would expect from looking at the behavior of people assigned fixed discounts. Results are presented in Table 1.7.

The results in Table 1.7 confirm the comparisons outlined in Table 1.6. We see that

Table 1.7: Regression Comparing BDM to Fixed Discounts

Sample is Non-Surveyed Customers in 2010				
	Fixed Price of Rs 100		Fixed Price of Rs 130	
	Played Card	Full Sample	Played Card	Full Sample
	(1)	(2)	(3)	(4)
BDM	0.0939**	0.0379	-0.367***	-0.216***
	(0.0445)	(0.0376)	(0.0636)	(0.0312)
Constant	0.718***	0.442***	0.580***	0.334***
	(0.0230)	(0.0196)	(0.0336)	(0.0160)
Observations	678	1113	665	1132
R-squared	0.237	0.157	0.277	0.195
Robust standard errors in parentheses			Village Fixed Effects	
*** p<0.01	** p<0.05	* p<0.1	Village Level Clustering	

Figure 1.5: Histogram of BDM bids for 1 Policy

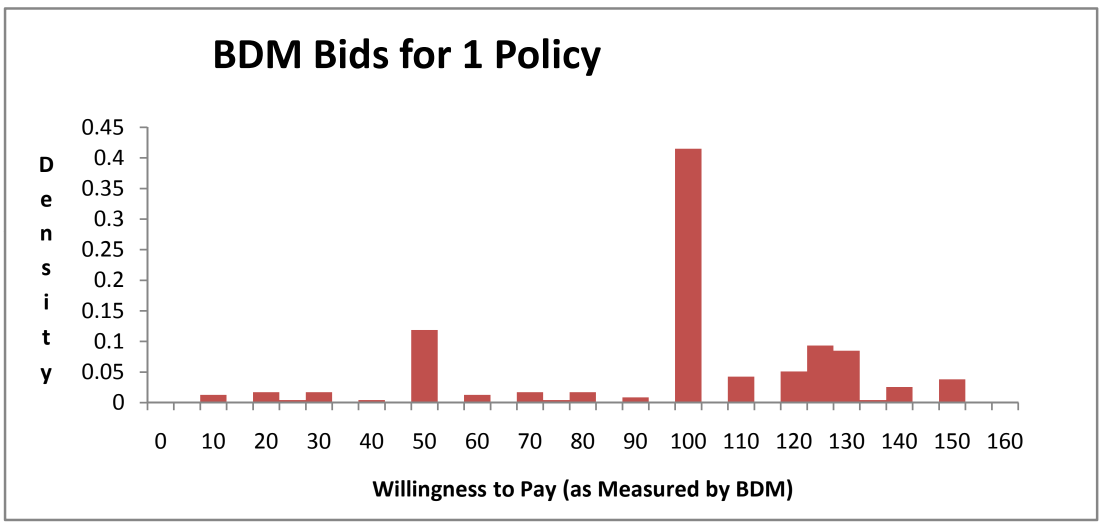

BDM generally gives inflated values of WTP compared with a fixed price of Rs 100 but decreased values compared with a fixed price of Rs 130. However, when we include the full sample in Column 2, we see that there is no significant difference in measured WTP when compared to a fixed price of Rs 100.

This analysis is clouded by the existence of focal points in the BDM data. From viewing the histogram of BDM bids for 1 policy in Figure 5, we can see that the majority of bids are the "round" numbers of 50 and 100. As a price of Rs 130 failed to encompass even the largest of these focal points, BDM appears to drastically underestimate WTP. The fixed price of Rs 100 probably provides a more realistic comparison, as this discount corresponds exactly to a focal point of the BDM bid distribution.

We argue that the fixed price of Rs 100 gives the most reliable comparison, and that the correct population to consider is the sample of all people who were visited to market

Figure 1.6: Demand Curves for 1 Policy

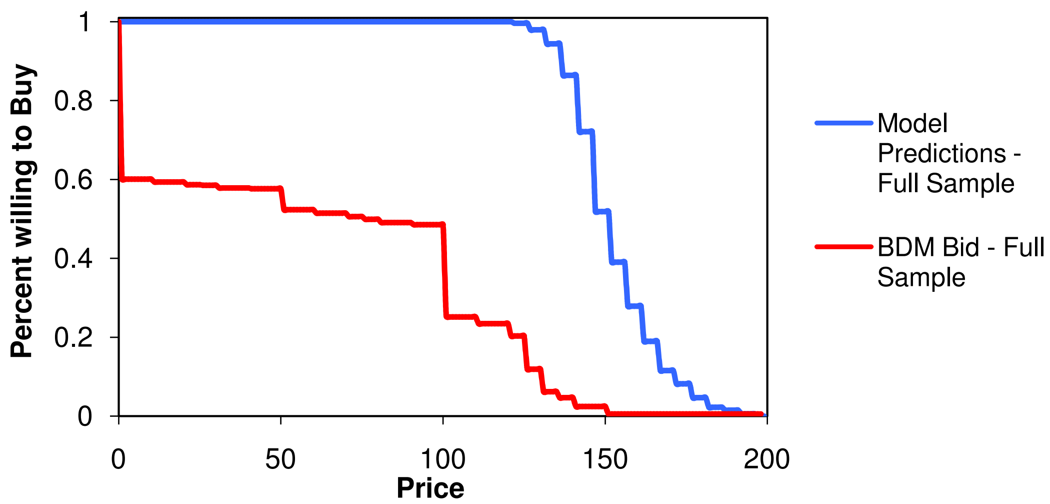

insurance, even if those people refused to play the scratch card game. Using this benchmark, BDM performs very well, as there is no significant difference in buying behavior for those assigned BDM versus those offered fixed prices.

These results are notably different than those of Berry et al. (2011), who find that BDM consistently undervalues WTP (compared to fixed price offers) through a number of frames and sub-treatments.

1.6 Test of BDM against the Theoretical Model

In this section we compare the WTP results from our theoretical model and the BDM procedure. We start by examining the demand curves as predicted by the model and BDM, and then look at whether the estimated WTP from the model has predictive power for BDM bids.

1.6.1 Demand Curves

In Figure 1.6, we plot the predicted demand curves for 1 insurance policy from both the theoretical model and from BDM. We include the full sample of those visited in the BDM plot, assigning a WTP of zero to people who refused to play the scratch card game. As expected from previous analysis, the theoretical model predicts a higher WTP at all price levels.

There are a couple of caveats to keep in mind when comparing the two demand curves, especially when looking at the lowest or highest prices. While people who refused to play

Figure 1.7: Demand Curves for 4 Policies

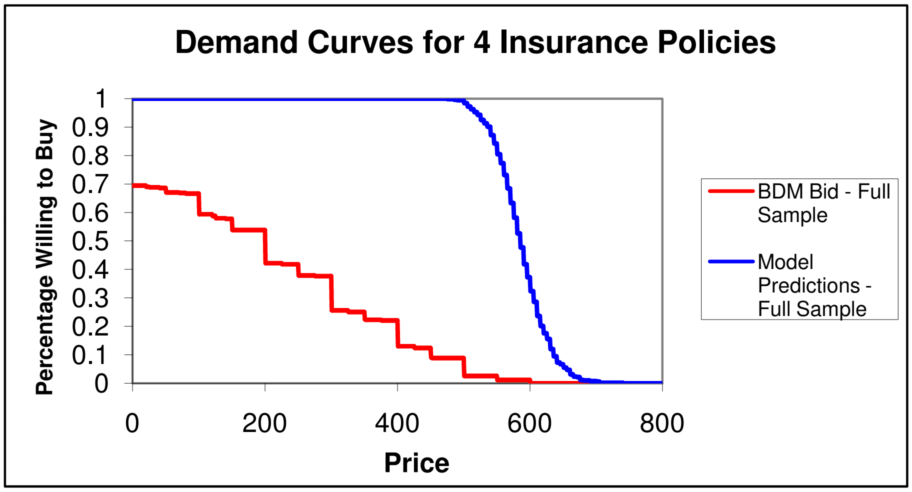

the scratch cards likely had low WTP, we simply assigned a zero WTP to this population, potentially underestimating WTP at low prices.

Similarly, people playing BDM had no incentive to ever bid more than the maximum price of the offer distribution, which was Rs 150. Accordingly, there were very few BDM bids above the maximum offer price of Rs 150. Perhaps this is due to people not exactly understanding the game, as they may have thought that by bidding less they were more likely to get a good deal. While we didn't offer anyone a fixed price of Rs 150, it is unlikely that no one would have bought at this price, as previous years' experience with the same population tells us that roughly 10% of people are willing to purchase insurance at market price. Therefore, BDM bids near 150 may actually reflect people with WTP greater than 150. Even with these caveats taken into account, the model clearly predicts higher WTP than the BDM game.

As we also played the BDM game for 4 policies, we can generate similar demand curves for a package of four policies, which is presented in Figure 7. Once again, we see that the theoretical model predicts higher WTP at all levels compared to BDM.

1.6.2 Can the model predict BDM bids?

While the previous analysis showed that the model predicts higher WTP than BDM, it is still possible that the estimated WTP at an individual level will be correlated with the BDM bids. To explore this we regress the BDM bids on the estimated WTP from the model, presenting the results in Table 1.8. In Column 1 we include only people who filled out scratch cards, and regress the BDM bid on a dummy that takes a value of 1 if the participant played the BDM game for 1 policy (as opposed to 4 policies) along with the model's predicted WTP. In this specification we see a positive yet insignificant point estimate on the model's predicted WTP.

In Column 2 we include the full sample, assigning a BDM bid of zero to people who

Table 1.8: Individual Comparison of BDM Bid and Predicted WTP

Dependent Variable is BDM Bid		
	Filled Card Only	Full Sample
	(1)	(2)
WTP Predicted by Model	0.224	0.268
	(0.163)	(0.170)
Game for 1 Policy	-87.61	-8.061
	(72.29)	(75.04)
Constant	155.2	41.21
	(97.11)	(100.9)
Observations	744	1045
R-squared	0.473	0.198

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

were assigned the BDM game yet refused to play. In this specification the point estimate is higher, indicating that an increase of Rs 100 of the predicted WTP is associated with an increase in the BDM bid of Rs 27. However, the estimate is only statistically significant at the 12% level.

1.7 Discussion

Our results have shown that, with some caveats, BDM provides an accurate measure of WTP. In contrast, a simple model overestimates WTP in the aggregate, and provides only weak predictions of WTP at the individual level. In this section we discuss some limitations of the model and ways it can be improved.

1.7.1 Insurance Demand and Risk Aversion

As described in the introduction, we do not include some features that others argue are important in determining insurance demand, such as ambiguity aversion (Bryan, 2010). This model, along with that of Clarke (2011) suggests that there may not be a simple relationship between risk aversion and demand for insurance. Since one main source of heterogeneity in our model is the coefficient of relative risk aversion, this relationship is important for the functioning of the model.

In Table 1.9 we regress a dummy which takes the value of 1 if the individual purchased insurance on risk aversion and risk aversion squared for each year of our study. In Column 1 we reproduce the results in Cole et al. (2010), showing that people with higher risk aversion had lower insurance demand in 2006, though the positive squared coefficient shows that these results weaken for high levels of risk aversion. However, these results disappear in subsequent years, with all significant correlation between risk aversion and insurance demand disappearing between 2007 and 2009. In 2010 there is a positive yet diminishing relationship between risk aversion and insurance demand. In a way, these results are consistent with a story of ambiguity aversion, as one may expect ambiguity towards a

Table 1.9: Risk Aversion and Insurance Purchasing

Year	2006	2007	2008	2009	2010
	(1)	(2)	(3)	(4)	(5)
Risk Aversion	-0.0895**	0.0115	-0.0223	0.0194	0.0661*
	(0.0341)	(0.0329)	(0.0236)	(0.0299)	(0.0337)
Risk Aversion Squared	0.00895**	-0.0000841	0.00159	-0.00197	-0.00921**
	(0.00438)	(0.00439)	(0.00303)	(0.00395)	(0.00443)
Constant	0.305***	0.317***	0.203***	0.149***	0.537***
	(0.0473)	(0.0412)	(0.0359)	(0.0338)	(0.0444)
Observations	315	824	765	725	756
R-squared	0.034	0.003	0.006	0.002	0.008
Robust standard errors in parentheses	Errors Clustered at Village Level				
*** p<0.01, ** p<0.05, * p<0.1					

new product to decrease over time. However, the positive coefficient in 2010 suggests that ambiguity aversion is no longer much of a factor for our sample.

Clarke (2011) focuses on the possibility that basis risk can make an insured individual worse off than an uninsured individual in a bad state of the world. Under certain circumstances, the demand for insurance can be increasing then decreasing in risk aversion, which is supported by the results in Column 5 of Table 1.9. People with low levels of risk aversion are uninterested in insurance (assuming there is premium loading), while people with high risk aversion will not want to take the risk of paying for a policy and subsequently suffering a loss that is not covered by the insurance policy.

While our model contains many of the same mechanisms as those of Clarke (2011), over the range of risk aversion in our sample our model always predicts WTP to be increasing in risk aversion. The main reason for this difference is in the structure of shocks and basis risk. In Clarke's model (and numerical example), there are many situations where people experience a heavy rainfall shock yet receive no payout, which makes insurance especially unattractive for the risk averse. In our model this situation still exists, but is less common. It is possible that our structure of basis risk does not adequately expose this possibility, causing our model to overestimate WTP.

1.7.2 Household Characteristics and WTP

While our theoretical model only weakly predicts BDM bids, it admits risk aversion and basis risk as the only sources of individual heterogeneity, and may miss other important individual factors that affect demand for insurance. In Table 10 we take a look at correlations between a number of household characteristics and BDM bids, which may give insight into other drivers of WTP.

Table 10 contains two types of outcome variables. The first labeled "BDM Bid/Total Price - Filled Card" is the BDM bid divided by the premium. We scale it this way so that we can easily pool together analysis for people who received the BDM game for 1 or 4 policies. This sample contains only surveyed households who filled out their scratch card. The second outcome, labeled "BDM Bid/Total Price – Full Sample" contains the full sample of surveyed participants, with their BDM bid being set to zero if they refused to play the game.

In Panel A we report the correlation of BDM bids with a number of household characteristics. Columns 1-2 report the coefficients obtained from regressing the outcome variables on each covariate individually. In Columns 3-6 we repeat the regressions with all right hand side variables included at once, and also run them with or without fixed effects.

The results give some evidence that people who have experienced drought recently, store goods, or have a loan from SEWA have lower insurance demand. People who have used other forms of SEWA insurance or have higher risk aversion are more likely to purchase insurance. But none of these results is robust across all specifications.

One interesting result comes out of our financial literacy variable, which measures the respondent's ability to answer a few simple questions about savings and credit. Just taking into account the people who filled out scratch cards, people with higher financial literacy tended to give a lower WTP. But if we include the full sample, then there is a positive correlation between financial literacy and WTP. This seems to indicate that people with higher financial literacy were more willing to play the BDM game (maybe due to the fact that they were more open to purchasing insurance), but had lower valuations conditional on playing that game.

Some other results are a bit counterintuitive. We would expect people who had lower risk exposure to have lower demand for insurance. This prediction is borne out somewhat in the data, as people who store goods have lower demand. But the amount spent on farm inputs, which we would think would be positively correlated to risk exposure, was negatively correlated to insurance demand. One theory is that this might reflect stronger liquidity constraints, but if liquidity constraints were a driving factor then access to credit would increase purchases. However, people who have loans from SEWA (which is an indication of access to credit) have lower demand.

Panel B restricts the sample to people who purchased rainfall insurance in 2009, and looks at the correlation between their experiences with insurance and their BDM bid the next year. There is some evidence that people who received a payout or reported higher satisfaction with insurance have greater insurance demand. We also include a variable called "Understanding of Product" which is the percentage of simple questions about rainfall insurance that they answered correctly. This variable is not significant in any of the specifications.

Overall, the results in Table 10 suggest that the most important factors related to insurance demand that are omitted from our model are dynamic considerations. This suggests that a model which takes into account previous experience with insurance may have better predictive power.

Table 1.10: Correlates of WTP

	Univariate		Multivariate			
Dependent Variable:	BDM bid / Total price	BDM bid (zeroed)	BDM bid / Total price	BDM bid / Total price	BDM bid (zeroed)	BDM bid (zeroed)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All Survey Respondents						
Total expenditure (Rs '0000)†	-0.004	-0.003	-0.003	-0.002	-0.003	-0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Total savings (Rs '0000)†	0.001	0.010	0.003	0.005	0.007	0.005
	(0.008)	(0.012)	(0.008)	(0.009)	(0.014)	(0.010)
Experienced drought	-0.092	-0.118	-0.075	0.002	-0.104	0.033
	(0.028)***	(0.033)***	(0.027)***	(0.023)	(0.030)***	(0.028)
Experience SEWA insurance	0.076	0.064	0.088	0.056	0.039	0.048
	(0.024)**	(0.033)*	(0.027)***	(0.026)**	(0.034)	(0.033)
Financial literacy	-0.111	0.072	-0.092	-0.024	0.113	0.045
	(0.050)**	(0.052)	(0.052)*	(0.046)	(0.056)*	(0.059)
Input spending (Rs '0000)†	-0.044	-0.037	-0.022	-0.016	0.007	0.023
	(0.017)**	(0.024)	(0.022)	(0.029)	(0.026)	(0.031)
Basis risk	-0.003	-0.004	-0.001	0.000	-0.004	0.000
	(0.003)	(0.004)	(0.003)	(0.000)	(0.004)	(0.000)
Uses HYV seeds	0.003	-0.020	0.019	0.010	0.007	-0.000
	(0.018)	(0.025)	(0.019)	(0.019)	(0.026)	(0.026)
Stores goods	-0.066	-0.115	-0.026	-0.010	-0.075	-0.058
	(0.034)*	(0.063)*	(0.033)	(0.041)	(0.059)	(0.058)
Risk Aversion	0.004	-0.004	0.027	0.005	0.041	0.021
	(0.004)	(0.006)	(0.014)*	(0.015)	(0.020)**	(0.018)
Risk Aversion Squared	0.000	-0.001	-0.003	-0.001	-0.006	-0.003
	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)**	(0.002)
Discount factor	-0.074	0.013	-0.095	-0.031	-0.025	-0.068
	(0.072)	(0.068)	(0.062)	(0.066)	(0.065)	(0.076)
Holder of loan from SEWA	-0.032	-0.014	-0.004	0.009	0.012	0.042
	(0.017)*	(0.024)	(0.017)	(0.020)	(0.023)	(0.022)*
Game for 4 policies	-0.199	-0.155	-0.198	-0.195	-0.160	-0.168
	(0.019)***	(0.022)***	(0.019)***	(0.018)***	(0.021)***	(0.019)***
Constant			0.830	0.710	0.527	0.501
			(0.057)***	(0.057)***	(0.074)***	(0.069)***
FE	NO	NO	NO	YES	NO	YES
Observations	745	1018	745	745	1018	1018
R-Squared			0.239	0.386	0.104	0.292
Panel B: People who Purchased Insurance in 2009						
Satisfaction with rainfall insurance	0.01	0.019	0.009	0.020	0.016	0.040
	(0.007)	(0.010)*	(0.007)	(0.014)	(0.010)	(0.012)***
Understanding of product	0.021	0.017	0.034	0.050	0.062	0.016
	(0.043)	(.069)	(0.040)	(0.064)	(0.065)	(0.071)
Payout last year (survey)	0.04	0.175	0.042	0.006	0.176	0.096
	-0.048	(0.044)***	(0.047)	(0.132)	(0.044)***	(0.096)
Game for 4 policies	(0.154)	-0.181	-0.158	-0.189	-0.188	-0.222
	(0.043)***	(0.053)***	(0.045)***	(0.060)***	(0.055)***	(0.060)***
Constant			0.627	0.609	0.409	0.387
			(0.044)***	(0.056)***	(0.068)***	(0.063)***
Village Fixed Effects	NO	NO	NO	YES	NO	YES
Observations	154	203	154	154	203	203
R-squared			0.122	0.363	0.151	0.437

Robust Standard errors in parentheses † Windsorized at top 1%
* significant at 10%; ** significant at 5%; *** significant at 1% Standard errors clustered at the village level

1.8 Conclusion

This paper outlined two approaches for measuring willingness-to-pay (WTP) for rainfall index insurance, and evaluates their effectiveness by comparing their predictions to decisions made my people facing fixed priced. The first approach a structural static model that generated predictions of WTP based on an individual's risk aversion and basis risk. We found that this model significantly overestimated WTP. We also implemented the BDM mechanism in the field, and found that it performed much better. While there were some problems with the BDM procedure that make the results hard to interpret (correlation of BDM bids and offers, focal points in BDM bid distribution), BDM gave predictions that were consistent with buying behavior of people who faced a fixed price of Rs 100. Finally, we found that our model's did have some predictive power over the BDM bids, but that the correlation between the two was weak and only significant at the 12% level.

One main shortcoming of this paper is that we were unable to conclusively determine the cause of our model's failure. While it is possible to tweak the parameters of the model such that its predictions correspond more closely with observed behavior (for instance, by assuming greater risk exposure and lower beliefs about expected payouts), we don't have any evidence that these calibrations are actually what is driving the shortcomings of our model. Most likely, a richer modeling framework will be necessary to generate trustworthy predictions of WTP.

The results from this paper have a number of policy implications. First, the distribution of WTP (as measured using BDM) shows us that in order to have high adoption of rainfall insurance, the policies must be heavily subsidized to above actuarially fair levels. Our data shows that in order to get 50% take-up of a single insurance policy, it needs to cost around Rs 100, which is roughly 2/3 the actuarially fair price. In order to get 50% takeup of a bundle of four policies the price needs to drop to around Rs 250, which is less than 50% of the actuarially fair value. For policy makers looking to promote risk mitigation among poor farmers, this suggests that very heavy subsidies will be necessary to convince farmers to purchase index insurance.

But one question that is still open is why demand for insurance is so low. Our model, which focused heavily on how basis risk can make insurance less attractive for the risk averse, still predicts much higher WTP than we see in practice. Generating the correct policy response depends on figuring out which mechanism is at play that we have not accounted for in the model. If people actually have other risk coping mechanisms so that rainfall shocks are not as damaging as we assume, then perhaps subsidizing rainfall insurance is a foolhardy effort. But if people are not buying because they don't understand its value, then perhaps WTP will increase over time as people become more familiar with insurance.

This suggests that it may be appropriate to take a dynamic approach to insurance demand, seeing how previous experience with rainfall insurance affects future demand. In the first three years of our study (2006-2008) we didn't see any insurance payouts, and the insurance payouts in 2009 were very modest, making any type of dynamic analysis difficult. In Chapter 2, I will explore the issue of dynamic demand using data from another microinsurance provider in India, BASIX.

Table A1.1: Risk Factors

				Approximate Partial Risk Aversion Coefficient			
Gamble	Head Payoff	Tails Payoff	Risk Level	Upper Bound	Lower Bound	Coeff Used for Regressions	Percentage of Respondents
1	25	25	Extreme	∞	7.51	7.5	9.31
2	23	45	Severe	7.51	1.74	3.615	9.06
3	20	60	Intermediate	1.74	0.812	1.189	14.1
4	18	63	Inefficient			0.8475	17.08
5	15	75	Moderate	0.812	0.316	0.506	10.48
6	10	80	Inefficient			0.337	17.98
7	5	95	Slight-to-Neutral	0.316	0	0.168	7.76
8	0	100	Neutral-to-Negative	0	-∞	0.01	14.23

1.9 Appendix

1.9.1 Risk Factors

In order to calculate risk factors, we use answers from games played during the household surveys. Each participant was asked to choose a list of lotteries that would be settled with a coin toss. The lotteries increased in both risk and expected payout, and the participants received the payout in real money at the end of the survey. We estimate CRRA risk aversion coefficients using the partial risk aversion coefficient, as calculated by Binswanger (1981). Table A1.1 shows the various gambles offered and calculated coefficients of risk aversion.

These estimates are in line with previous estimates of relative risk aversion, as explained in the introduction to Halek and Eisenhauer (2001). They list a number of studies that estimate CRRA coefficients using different methodologies, which generally find coefficients between zero and two. This is consistent with our estimates, as around 82% of the subjects fall into this range. The other 82% are higher, but finding outliers on the high end of the CRRA distribution is also supported in the literature. For instance, Hansen and Singleton (1983) find outliers as high as 58.25.

1.9.2 Model Sensitivity Tests

In this section we look at how varying some of the parameters of the structural model affect its predictions. In Figure A1.1 we first look at how varying the subject's expectation of premium loading changes insurance demand. To do this we change the amount of payouts in our model until they correspond to lower or higher payouts on average. The standard deviation of the estimate of expected payouts is 21% of the premium, so we present the results of the model with a loading factor of -42%, -21%, 21%, and 42%.

As expected, this has a large effect on insurance demand, and if customers had lower expectations of insurance payouts, this could partially explain the gap between the model and observed demand.

Next, we consider how the ratio of potential losses to wealth affects insurance demand.

Figure A1.1: Sensitivity to Premium Loading

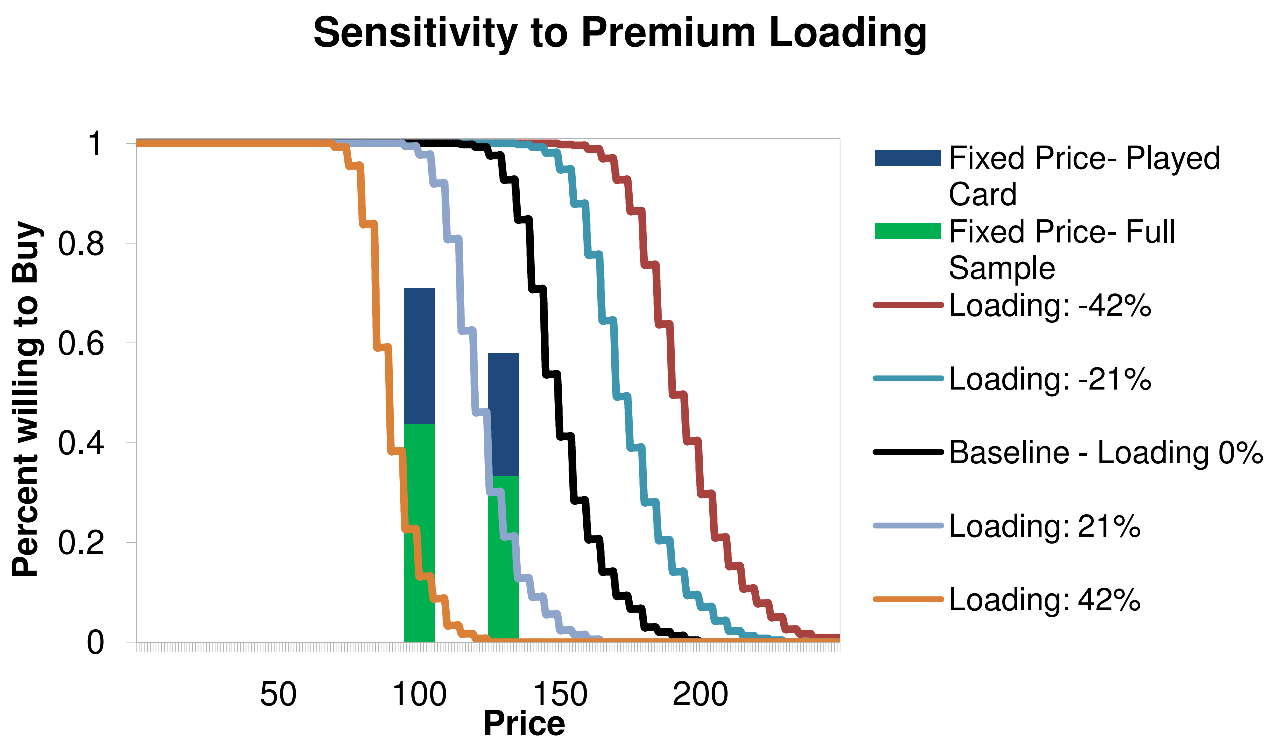

We would expect that people with larger wealth (keeping potential losses constant) would have lower insurance demand since with CRRA utility risk aversion decreases with wealth. This is exactly what we see in Figure A1.2. Doubling wealth from the baseline of 41800 decreases insurance demand but not by much. Decreasing wealth increases demand, but most of this effect comes at high prices. Although not shown on the graph, further analysis shows that this demand for insurance at high prices is driven by people with high levels of risk aversion.

Finally, insurance demand can be sensitive to basis risk. In our model the correlation between rainfall used to calculate shocks and rainfall used to calculate payouts varies based on the distance between someone's village and the reference rainfall station, which varies from roughly .63 to .67. In Figure A1.3 we present the demand curve for different levels of constant basis risk, ranging from no correlation between the income shock and payouts to perfect correlation of the rainfall used to calculate income shocks and payouts. As expected, higher basis risk leads to lower demand, but the effect is not extremely strong. Even at zero correlation, predicted demand is above observed demand.

Figure A1.2: Sensitivity to Wealth

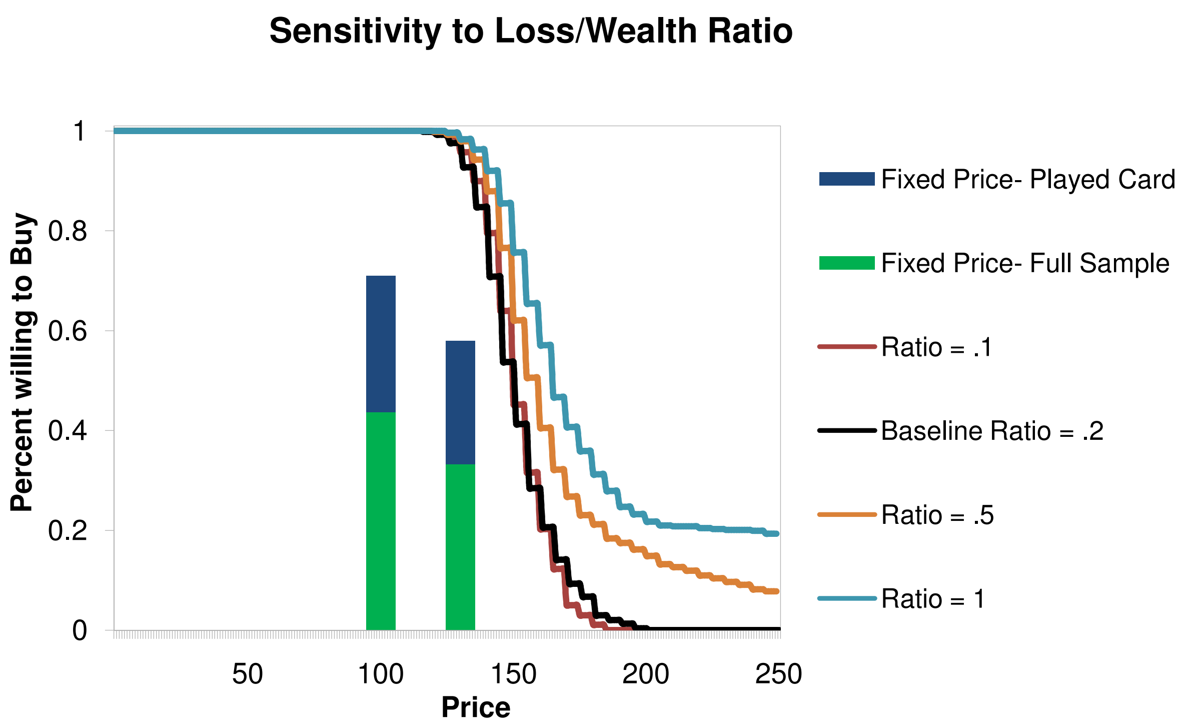

1.9.3 BDM Implementation

Scratch Cards

An example scratch card used to deliver the BDM game is shown in Figure A1.4. The text in the top panel translates to "Scratch Here to Reveal Discount". This panel was scratched first, revealing whether the customer was going to play the BDM game or receive a fixed discount. If they were supposed to play the game, they next played a practice game for a napkin. Top left boxed on the back of the card provided a space to write the bid for the napkins, and then the second scratch panel revealed the napkin offer price. The bid for insurance went into the bottom left boxes on the back of the card, and the offer for insurance was under the bottom right scratch panel.

Censoring of Cards

When participants ended up purchasing a policy as a result of the BDM game, the enumerators had high incentives to carefully record the participants' bid, as this would be proof that they won the game and therefore were allowed to purchase at a discounted price. But when participants "lost" the game, meaning they did not purchase the policy, we worry that the BDM bid may have not been reported. If this was true, then we would expect the cards that were filled out in the field to have the distribution of their offer prices skewed downward.

Figure A1.3: Sensitivity to Basis Risk

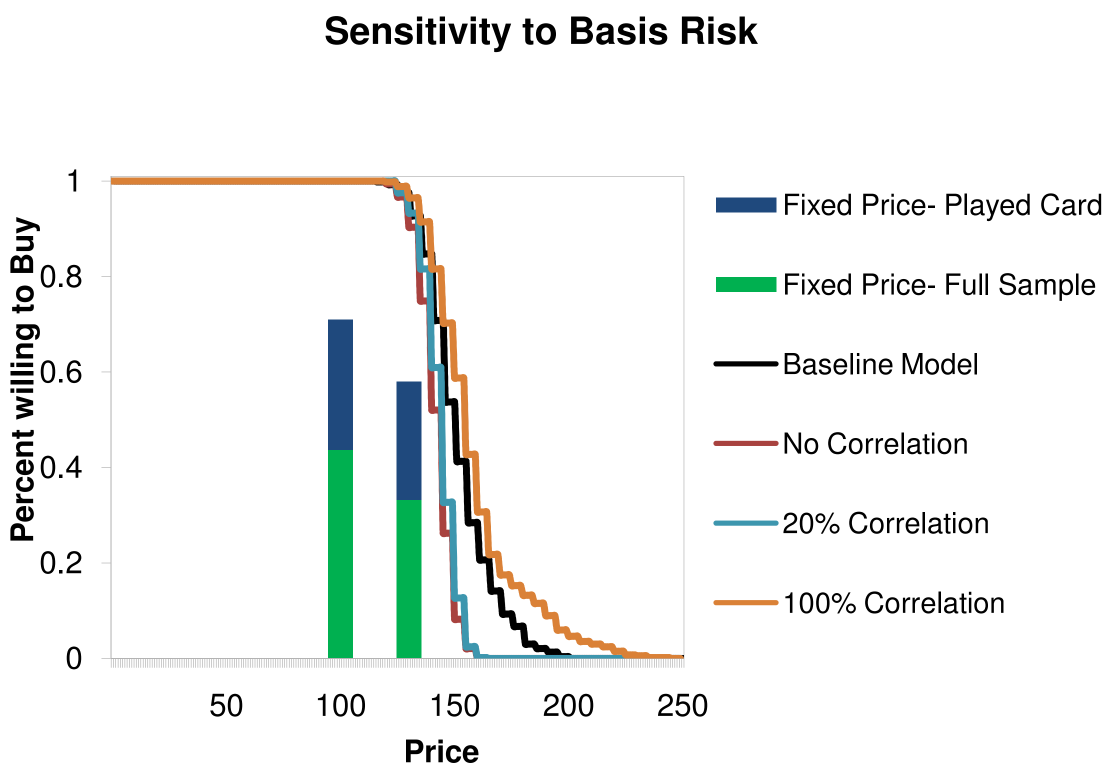

Figure A1.4: Example Scratch Card. Left is front, Right is Back.

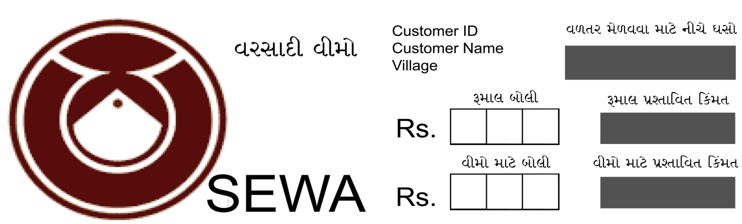

FIGURE A1.5: DISTRIBUTION OF BDM OFFERS. LEFT IS OFFERS SEEN IN THE FIELD. RIGHT IS OFFERS GENERATED ON ALL THE CARDS.

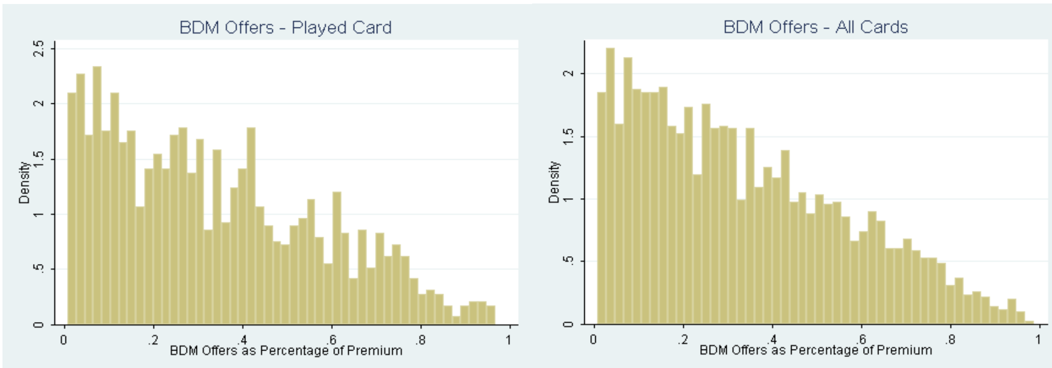

TABLE A1.2: CORRELATION OF BDM BIDS AND BDM OFFERS

Dependent Variable is BDM Bid as % of premium	All Districts (1)	Patan (2)	Anand (3)	Ahmedabad (4)
BDM Offer (as % of Premium)	0.166***	0.170***	-0.0932**	0.428***
	(0.0489)	(0.0397)	(0.0448)	(0.0724)
Game for 1 Policy	0.208***	0.268***	0.202***	0.144***
	(0.0160)	(0.0221)	(0.0257)	(0.0226)
Constant	0.427***	0.304***	0.626***	0.327***
	(0.0308)	(0.0245)	(0.0220)	(0.0495)
Observations	1506	419	557	530
R-squared	0.192	0.366	0.207	0.246
Robust standard errors in parentheses	Standard Errors Clustered at Village Level			
*** p<0.01, ** p<0.05, * p<0.1				

Figure A1.5 graphs the histograms of these two distributions side by side, and they seem to be quite similar. A Kolmogorov-Smirnov test is unable to reject the null hypothesis that the two distributions are the same. Similarly, a Fisher exact test (using five bins) cannot reject the hypothesis that there were different distributions on cards that were scratched off in the field versus those that were not.

Correlation of BDM Price to BDM Offer

In Table A1.2 we observe a correlation between the BDM Offer and BDM bid, indicating that some people saw the BDM offer before they committed to their bid.

Column 1 shows that there is a positive correlation overall between bids and offers, suggesting that people changed their bids after seeing the offer price or had their WTP anchored by viewing the offer price. Columns 2-4 perform the regression separately for each district (since each district had a different marketing team), and reveals some puzzling heterogeneity. Patan and Ahmedabad districts show the same pattern as the overall

Table A1.3: Napkin Game

Dependent Variable is BDM Bid as % of premium	OLS	OLS	IV
	(1)	(2)	(3)
Napkin BDM Offer	0.0150***		
	(0.00417)		
Won Napkin Game		0.0435	-0.316***
		(0.0321)	(0.0765)
Constant	0.553***	0.564***	0.891***
	(0.0276)	(0.0277)	(0.0642)
Observations	1450	1450	1450
R-squared	0.021	0.003	.
Robust standard errors in parentheses	Standard Errors Clustered at Village Level		
*** p<0.01, ** p<0.05, * p<0.1			

data, with an especially strong correlation in Ahmedabad. However, there is a negative correlation between bids and offer in Anand, which is difficult to explain.

Napkin Correlation

As shown in Table A1.3, the outcome of the napkin game seems to have influenced the BDM bids of the participants. Column 1 regresses the BDM bid on the napkin offer, finding that an increase of 1 Rs in the napkin offer is correlated with a 1.5 percentage point increase in the BDM bid. It is possible that the mechanism for this effect is whether or not the participant won the napkin game. In Column 2 we regress the BDM bid on a dummy which takes a value of 1 if the participant won the napkin game. This yields an insignificant coefficient, but this specification is of dubious quality due to the likelihood of unobserved variables that would drive both the napkin bid and BDM bid. If one is willing to believe that the only channel in which the napkin offer can affect the BDM bid is through winning the napkin game, we can instrument for winning the napkin game with the napkin offer. We do this in Column 3, and see that winning the napkin game causes BDM bids to decrease by an astonishing 32 percentage points.

Chapter 2

Paying Premiums with the Insurer's Money: How Loss Aversion Drives Dynamic Insurance Decisions

2.1 Introduction

Unlike physical goods (or even credit), it is difficult for customers to evaluate the benefits of insurance since its main benfits only occur when a payment is received. If customers are unfamiliar with how insurance works, they may be influenced by their recent experiences with insurance and also by the experiences of their friends and neighbors. Providing evidence from the developed world, Kunreuther et al. (1985) observe that purchases of flood and earthquake insurance in the US are greatly influenced by recent experiences with disasters and insurance payouts, and also argue that peoples' insurance decisions are influenced by their friends and neighbors' experiences with insurance. Reacting to low rates of rainfall insurance uptake in Andhra Pradesh, India, Giné et al. (2008) suggest that "over time, lessons learned by insurance 'early adopters' will filter through to other households, generating higher penetration rates among poor households."

This paper seeks to understand how previous insurance payouts can affect future insurance purchasing decisions, and what mechanisms can explain this behavior. Using data on three years of insurance purchasers from the Indian micro-finance institution BASIX, I find that customers who received an insurance payout are 9-22% more likely to repurchase in the following year as compared to customers who did not receive any insurance payments. Customers who received larger insurance payouts are more likely to repurchase than those who received small payouts. However, the paper finds no evidence of positive spillover effects to other people in the village even at large levels of payouts, casting doubt on the hypothesis that witnessing insurance payouts will spur new buyers.

I introduce a model in which customers exhibit loss aversion and repurchases of insurance are driven by the psychological effects of a reference point shift. Purchasers of insurance who receive insurance payouts will be more likely to purchase insurance in the future because their previous insurance "profits" make future premiums seem like less of a loss. Following the logic of Thaler and Johnson (1990), I propose that customers who receive an insurance payout will not regard future premium payouts as a true loss, which can be modeled as a shift of the reference point that increases risk aversion. I define the conditions under which risk aversion will increase after a payout, and argue that these conditions are likely to hold.

I also explore some alternative hypotheses as to why receiving payouts could increase insurance demand the following year. First, it is possible that weather shocks themselves could have an effect on insurance demand, such as Kunreuther et al. (1985) observe in the US. This could happen because weather shocks change customers' beliefs about future shocks, change their wealth, or simply increase the salience of shocks. I look for direct effects of weather by testing how rainfall in the year before insurance was introduced affected insurance purchases, and find evidence that previous rainfall shocks tend to *decrease* insurance purchasing. This provides evidence against the argument that it is weather shocks as opposed to insurance payouts that are driving insurance purchases.

Next I test whether receiving insurance payouts could induce trust in the insurance company or learning about the insurance product. Bryan (2010) suggests that index insurance take-up is low due to ambiguity aversion, which should decrease as customers learn more about the product. We assume that if trust and learning are driving purchases, we should be able to witness spillover effects on other people in the village. This is because people in a village who witnessed payouts but did not receive them should also have been able to learn about insurance and gain trust in the insurance company. I do not find convincing evidence that these spillover effects are present, and argue that this is evidence that insurance repurchasing is not being driven by trust or learning. Overall, the empirical analysis suggests that it is the physical reception of payout money that drives future purchases, which is consistent with the proposed model of loss aversion.

This paper is related to a few separate lines of research. First, it contributes to a growing list of empirical studies that attempt to determine demand for weather index insurance (Hill and Robles, 2010; Cole et al., 2010; Giné et al., 2008). One overarching conclusion from previous studies on index insurance is that demand for index products is low when provided at market rates[1], and only increases when prices are slashed significantly. However, most of these studies look at insurance as a static purchasing decision, seeing what factors lead people to become first time customers.

One exception is Hill and Robles (2010), who provide rainfall insurance for free as part of an experimental game in Ethiopia, and then return the next year to sell the same insurance. Despite the fact that two-thirds of the people who were granted insurance during the experiment received payouts, this group had low take-up rate of 11% the next season, which was a lower rate than those who had not participated in the experiment.

Our study differs from these in that it uses a much larger dataset, allowing comparisons across weather stations to identify the dynamic effects, which I argue allows a clearer identification of the effects of payouts. Additionally, this paper studies a real world insurance

[1] Market rates tend to be around 2-6 times actuarially fair rates.

implementation at market rates, making its results potentially more relevant for policy making.

Another related work is an observational study on mutual insurance among fisherman in the Ivory Coast by Platteau (1997). Platteau observes malfunctioning mutual insurance cooperatives and theorizes that they are failing because members view insurance as a system of balanced reciprocity, meaning that they expect to break even over the lifetime of the scheme. When members have not received the services (in this case sea rescue) of the mutual in a long time, they start to view the insurance as a bad deal and ask for their contributions back. This paper's loss aversion model provides a different theoretical framework that can also generate predictions of a desire for balanced reciprocity among insurance customers.

This work also contributes to the literature on choice under uncertainty by providing a theoretical framework to understand how shifting reference points can help explain dynamic insurance choices. The basic ideas behind the model are not new: in their original paper on prospect theory Kahneman and Tversky (1979) note that if customers "expect" to purchase insurance such that their payment of the insurance premium is not counted as a loss, this makes insurance more attractive. Similarly, Köszegi and Rabin (2007) argue that if a risk and the ability to insure it are "anticipated," then payment of the insurance premium should be incorporated into the reference point and therefore not counted as a loss.

But neither of these papers are specific about what may cause a customer to move from one reference point to another. Kaheneman and Tverky muse that someone may expect to purchase insurance (and therefore have a different reference point) "perhaps because he has owned it in the past or because his friends do." This paper builds on these ideas but adds the crucial assumption that receiving insurance payouts is the key event that can trigger a change in the reference point. One way to justify this assumption is to invoke the concept of the "lagged status quo," which means that after experiencing a gain (or loss), the new reference point may not immediately update to the current level of wealth. If the reference point lags after an insurance payout, loss aversion dictates that people will become more risk averse and therefore more likely to purchase insurance in the next period.

Behavior consistent with the lagged status quo is demonstrated by Thaler and Johnson (1990), who show that gamblers who had experienced recent gains were less sensitive to subsequent losses. (This is the "gambling with house money" effect.) Their framework offers the equivalent of a reference point shift such that after winning money subjects do not psychologically consider subsequent negative outcomes as true "losses." When the same logic is applied to insurance purchases, it means that people who receive an insurance payout will not regard subsequent premium payouts as a loss, making them more risk averse and therefore more likely to purchase insurance.[2] This is the mechanism utilized in the theoretical framework of this paper.

Another relevant work is Köszegi and Rabin (2007), who adapt their theory of stochastic

[2]The subjects in Thaler and Johnson's experiments play gambles involving gains, so shifting the reference point makes them more risk loving, as potential losses are now less painful while gains are relatively unchanged. However, when the same principle is applied to insurance, a gamble which involves only losses, the same reference shift will make customers more risk averse.

reference points (Köszegi and Rabin, 2006) to insurance decisions and explicitly outline how different reference points will affect insurance decisions. In Köszegi and Rabin's model the "correct" reference point to consider when analyzing insurance decisions depends on whether the gamble is anticipated, and how long a time there will be between when the gamble is chosen and the payouts are realized. Their key insight as relates to this study is that if insurance purchases are anticipated and made far in advance of their outcome, they will not be counted as a loss as the reference point adjusts to account for the insurance premium. If the transition from a "surprise" gamble to an "anticipated" expense happens after customers receive insurance payouts, Köszegi and Rabin's model would generate predictions similar to this paper. However, this is difficult to justify because insurance decisions are made over the same time frame each year. The model presented here instead relies on a simpler mechanism, assuming that reference points are fixed within a decision period but can shift based on the outcome of that period.

The paper will proceed as follows. Section 2.2 outlines the model and determine the conditions under which an insurance payout will lead to greater insurance purchasing. The model predicts that insurance customers who receive a payout will be more likely to purchase insurance the following year. Section 2.3 explains the insurance policies and data that will be studied in the empirical section. Section 2.4 provides the main empirical evidence, which confirms the main predictions of the theory. Section 2.5 looks at alternative explanations for the increased insurance purchases, such as the possibility that they are driven by increased trust or pure weather effects. I look for evidence that these mechanisms are driving insurance decisions and fail to find them. Section 2.6 concludes and offers policy recommendations based on the results.

2.2 Theory

In this section I introduce a theoretical framework that seeks to explain how experience with insurance could affect the decision to purchase insurance during the following season. I present a model where agents exhibit loss aversion, but their reference point can shift based on previous experiences with insurance. The key component of the model is the reference point shift, which can alter the risk aversion of the agents and therefore change their demand for insurance. As the empirics of this paper show that receiving a payout correlates with increased insurance purchases, I focus on payout reception as the key moment when reference points are likely to shift. After setting up the basic framework of the model, I then determine the restrictions on the evolution of the reference point that will yield the aforementioned empirical prediction.

Subjects have a piecewise-linear utility function that exhibits loss aversion around a reference point. Given a reference point r, utility is:

$$u(c,r) = \begin{cases} \alpha c & if \quad c > r \\ \beta(c-r) + \alpha r & if \quad c < r \end{cases} \tag{2.1}$$

Figure 2.1: Utility Functions

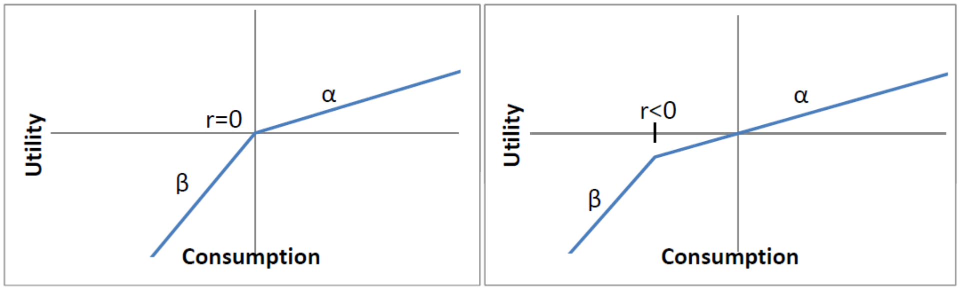

The function on the left has reference point $r = 0$, and is risk neutral for losses. The figure on the right has $r < 0$, making the subject risk averse for losses.

Equation 2.1 defines the utility funtion for $r < 0$, as this is where all the interesting dynamics of the model take place. Figure 2.1 shows the utility function for two values of r.

The model lasts two periods, and in each period there are two possible states of the world $S = \{0, 1\}$. If $S = 1$, which happens with probability p, agents suffer a consumption shock of $-X$. If $S = 0$ there is no shock. Agents also have the opportunity to purchase insurance $I = \{1, 0\}$ against the shock. If an agent purchases insurance he is completely protected from the consumption shock should $S = 1$ occur. Insurance costs a constant multiple $(1 + \lambda)$ of the expected payout of insurance, resulting in an insurance premium of $(1 + \lambda)pX$. At the end of period 1 the agent's reference point can move, as defined by the function $r_2 = f(r_1, I_1, S_1)$. This means that the reference point has a chance to shift from the first to the second period based on a customer's experience with insurance (determined by I_1, S_1) in the first period. The form of the function $f(r, I, S)$ will determine the interesting dynamics of the model, as it is the change in reference point that will generate changes in the insurance decision.

Timing is as follows:

1. Agent starts with reference point r_1
2. Agent chooses insurance decision $I_1 = \{0, 1\}$.
3. State of the world $S_1 = \{0, 1\}$ is realized. Agent receives period 1 utility.
4. Agent's reference point moves to $r_2 = f(r_1, I_1, S_1)$
5. Agent chooses $I_2 = \{0, 1\}$
6. State of the world $S_2 = \{0, 1\}$ is realized, and agent receives period 2 utility.

There is no discounting. Each period, the agent decides whether or not to purchase insurance by considering his expected utility with or without insurance. I assume that $0 \geq r \geq -X$, which restricts analysis to the case where some loss aversion is present over the possible range of consumption.[3] I also assume that agents are naive, which means that they don't take possible shifts in their reference point into account when making their first period insurance decision.[4]

Applying the utility function defined in Equation 2.1, expected utility U in each period is defined in Equation 2.2.

$$\begin{cases} U(r,I) = \beta I(((-(1+\lambda)pX) - r) + \alpha r) + p(1-I)(\beta(-X-r) + \alpha r) & if \quad -(1+\lambda)pX < r \\ U(r,I) = \alpha I(-(1+\lambda)pX) + p(1-I)(\beta(-X-r) + \alpha r) & if \quad -(1+\lambda)pX > r \end{cases} \tag{2.2}$$

In this simple decision model, an agent would choose to purchase insurance if expected utility from purchasing insurance was greater than expected utility from forgoing insurance. However, in a richer model there may be other factors that would influence a potential buyer's purchasing decision. I therefore analyze the benefit of insurance, defined as the difference in expected utility from purchasing insurance versus forgoing insurance. I assume that if people have greater benefits from purchasing insurance, they will be more likely to purchase.

$$B(r) \equiv U(r,1) - U(r,0) \tag{2.3}$$

One can interpret the benefits of insurance $B(r)$ to be a measure of risk aversion, as higher benefits from insurance imply higher risk aversion.[5]

The Period 2 reference point r_2 is determined by the initial reference point, insurance decision, and state of the world in period 1 according to $r_2 = f(r_1, I_1, S_1)$. We can use the function f to connect benefits of insurance in Period 2 to insurance experiences in Period 1. People who received a payout in period 1 will be more likely to re-purchase insurance than insurance customers who did not receive a payout if

$$B_2(f(r_1,1,1)) > B_2(f(r_1,1,0)) \tag{2.4}$$

The key is to discover the conditions on $f(r,I,S)$ such that Inequality 2.4 holds. Due to the piecewise definition of the utility function, the complete set of conditions that guarantee Inequality 2.4 holds are somewhat complex and are explained in detail in the Appendix. In order to illustrate the basic mechanism we can examine the specific and highly plausible

[3] $r = 0$ and $r = -X$ give risk neutral preferences, but I include these endpoints as limits of the interesting range of r.

[4] A sophisticated agent would make different choices in period 1 since they would anticipate how those choices would change their reference point in period 2. However, the difference in choices in the second period between an insurance customer who received a payout in period 1 versus a customer who did not receive a payout will not change if agents are sophisticated. As this is the main result we care about, considering only naive agents lends considerable mathematical simplicity to the problem.

[5] $B(r)$ as defined here will be monotonically related to the risk premium as defined in Pratt (1964).

case where $r_1 = 0$ and $f(r_1, 1, 0) = r_1 = 0$. The assumption of $r_1 = 0$ can be interpreted as the agent having no previous experience with insurance, and therefore exhibiting standard loss aversion with the reference point being his initial level of consumption.[6] $f(r_1, 1, 0) = r_1$ means that insurance purchasers who do not receive an insurance payout do not experience a shift in their reference point.

Proposition 2.1. In the case of $r_1 = 0$, $f(r_1, 1, 0) = r_1$, $\quad B_2(f(0, 1, 1)) > B_2(f(0, 1, 0))$ iff $0 > f(0, 1, 1) > -(1 + \lambda)X$.

In words, this proposition states that if an insurance payout causes a decrease in the reference point (up to $-(1 + \lambda)X$), this increases the agent's risk aversion, and therefore his benefits from purchasing insurance in period 2. Intuitively, a decrease in the reference point corresponds to a lagged status quo, and makes premium payments in the next period not seem like a real loss.

Proof. I start by looking at people who purchased insurance in period 1 but did not receive a payout. Since we have assumed that for this group the reference point remains at zero, the benefits of insurance in period two are straightforward to calculate from Equations 2.2 and 2.3.

$$B_2(0, 1, 0) = -\lambda\beta pX \tag{2.5}$$

For those who did receive a payout, the reference point will move to $r_2 = f(0, 1, 1)$. Due to the piecewise nature of the utility function, the benefits are also going to depend on whether or not the reference point remains above the insurance premium or moves below it. From Equations 2.2 and 2.3 we see:

$$\begin{cases} B_2(0, 1, 1) = -\lambda\beta pX + (1 - p)(\alpha - \beta)f(0, 1, 1) & if \quad 0 > f(0, 1, 1) > -(1 + \lambda)pX \\ B_2(0, 1, 1) = -(1 + \lambda)\alpha pX + p\beta X - p(\alpha - \beta)f(0, 1, 1) & if \quad f(0, 1, 1) < -(1 + \lambda)pX \end{cases} \tag{2.6}$$

In the case that $0 > f(0, 1, 1) > -(1 + \lambda)pX$, $B_2(0, 1, 1) > B_2(0, 1, 0)$ iff $f(0, 1, 1) < 0$. When $f(0, 1, 1) < -(1 + \lambda)pX$, we can prove the inequality by combining Equations 2.6 and 2.7.

$$B_2(0, 1, 1) - B_2(0, 1, 0) = p(\beta - \alpha)(f(0, 1, 1) + (1 + \lambda)X) \quad if \quad f(0, 1, 1) < -(1 + \lambda)pX \tag{2.7}$$

[6] If one follows Köszegi and Rabin (2006), the initial reference point would be the expectation of consumption, making $r_1 = -pX$. The main intuition that follows is still valid, and the solution to Inequality 2.4 for a general $r_1$is given in Proposition 2.2 in the Appendix. One point to notice is that the maximum benefit of insurance is obtained when $r = -\lambda pX$ (Proposition 2.3 in the Appendix). Therefore, if $\lambda \leq 1$ and the intial reference point is the expectation of consumption, the reference point yielding maximum benefit from insurance lies above the starting reference point, and the dynamics of the model change considerably. In the real world this is an unlikely scenario as market prices for insurance are generally above their expected payout, making $\lambda > 1$. In our empirical sample of BASIX insurance products, λ ranges from 1.1 to 15.7, with an average of 5.8.

So for this case $B_2(0,1,1) > B_2(0,1,0)$ iff $f(0,1,1) > -(1+\lambda)X$. Combining these two cases proves Proposition 2.1. ■

In words, this means that buyers who receive insurance payouts in period 1 will be more likely to purchase insurance in Period 2 if their reference point decreases as a result of receiving a payout. This basic result should be apparent from analyzing Figure 2.1. In the left panel of Figure 2.1, the reference point is at zero and the subject is risk neutral for losses in relation to the reference point. In the right panel, the reference point has moved to below zero and the customer is now risk averse, therefore showing increased demand for insurance.

The second condition $f(0,1,1) > -(1+\lambda)X$ will only bind if $\lambda < 0$, as we have already assumed that $r > -X$. This condition comes about from the fact that when $\lambda < 0$ there is potentially a benefit from purchasing insurance even for risk lovers, as the expected return from purchasing insurance is higher than the expected returns without. I'll call this the expectation benefit. The expectation benefit is greater if an agent's utility function is steeper between the expected returns with and without insurance. When the reference point is below the price of insurance, the utility function flattens in this range, decreasing the expectation benefit. As the reference decreases from $-(1+\lambda)pX$ towards $-X$, risk aversion decreases until the utility gains from insurance approach the utility losses from the reduced expectation benefit. These forces are equal when $f(0,1,1) = -(1+\lambda)X$.

From Proposition 2.1, we see that if the reference point moves such that it is below the level of current wealth, customers become more risk averse and therefore more likely to purchase insurance. Such a shift is likely to occur after an insurance payout assuming that insurance customers behave in a manner consistent with the participants in Thaler and Johnson's (1990) lab experiments. In these experiments people were more likely to gamble with money they had recently won as they did not regard losing this money as a true loss, which is consistent with having a reference point below the level of current wealth. I hypothesize that people who receive an insurance payout regard this payout akin to a gambling victory, and therefore this payout shifts their reference point below the level of current wealth. According to Proposition 2.1, this shift makes recipients of insurance payouts more likely to purchase insurance than other customers who did not receive payouts.

While this model is similar to models of the lagged status quo (Thaler and Johnson, 1990; Gomes, 2005), it is formulated somewhat differently to allow for greater flexibility in the movement of the reference point. In standard models of the lagged status quo, the reference point does not move after a gamble even though wealth has changed, causing a wedge between current wealth and the reference point. In this model, at the beginning of each period utility is normalized to zero, while instead the reference point is allowed to shift. Despite this change in modeling form, this model also creates a wedge between current wealth and the reference point, allowing it to give similar predictions as a lagged status quo model. For instance, if the reference point shifts down after experiencing a gain then the reference point is below the current level of wealth, just as it would be in a lagged status quo model. Allowing the reference point to shift arbitrarily, however, gives this model the ability to both reproduce the results of lagged status quo models while also allowing insights gained from other possible reference point changes.

This theory provides a framework that we can use to analyze how insurance experiences can change future insurance decisions. Specifically, it offers conditions under which customers who receive payouts would be more likely to purchase insurance in the following period. Assuming these conditions hold, this model predicts that people who receive a payout will be more likely to purchase insurance in the following year. I turn next to the empirical section to show that the behavior of rainfall insurance purchasers in India is indeed consistent with this prediction. I then show that alternative mechanisms which would fall outside of this theory, such as purchases being driven by increased trust in the insurance company or direct effects of weather, are not supported by the data.

2.3 Index Insurance and Customer Data

2.3.1 Context: BASIX Policies

In this analysis I study monsoon rainfall index insurance policies underwritten by ICICI-LOMBARD and sold by BASIX, a microfinance institution based in Hyderabad. The policies insure against excess or deficit rainfall, and are calculated based on rainfall measured at a stated weather station. By basing payoffs on just rainfall, the policies should have low monitoring and verification costs, and also should be free of adverse selection and moral hazard. These attributes make policies inexpensive to create and administer, which allows them to be sold in small quantities and priced at levels affordable for poor farmers. BASIX's policies are designed to pay out in situations where adverse rainfall would cause a farmer to experience crop loss, and are therefore calibrated to the water needs of local crops.

BASIX policies are divided into three phases, which are meant to roughly capture the three phases of the growing season: planting, budding/flowering, and harvesting. If cumulative rainfall is too low or high in any of these phases, the crop's output is potentially damaged and the farmer could suffer a loss. The policies are designed to start when farmers first start planting, which depends itself on rainfall. Therefore, the policies have a dynamic start date which means that Phase 1 begins on the day that cumulative rainfall since June 1 reaches 50mm or on July 1, whichever comes first. Each phase generally lasts 35-40 days. During this time, rainfall data is collected daily at a designated weather station, and payouts are calculated using the cumulative rainfall over the phase.

A phase of coverage is defined by three parameters: "Strike", "Exit", and "Notional". Deficit policies begin to pay out when the rainfall drops below the level of the Strike, and gives its full payout when it falls below the Exit. In between, it pays the Notional amount of rupees for each millimeter below the Strike.

In 2006 and 2007, all rainfall insurance contracts sold by BASIX included three phases, with the first two protecting against deficit rainfall, and the third protecting against excess rainfall. In 2005 the policies all had three phases, but each phase protected only against

Table 2.1: Example Insurance Policy

Phase	I	II	III
Duration (Days)	35	35	35
Type	Deficit	Deficit	Excess
Strike (mm)	135	125	730
Exit (mm)	40	40	820
Notional (Rs/mm)	10	10	10
Policy Limit (Rs)	1000	1000	1000
Premium (Rs)	110	110	90

Figure 2.2: Example Payout Schedule

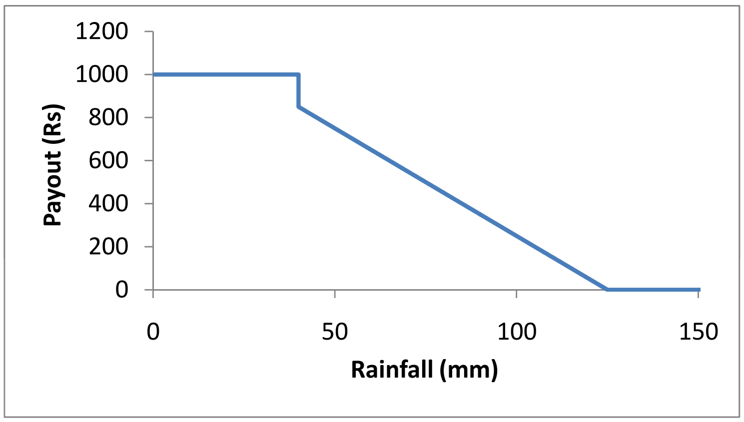

deficit rainfall. Table 2.1 presents a sample contract, from Nizamabad district in the state of Andhra Pradesh.

Given the policy parameters we can see how the payouts will evolve according to rainfall. Figure 2.2 shows the payout schedule for phase II of the above policy. There is no payout when rainfall is above the strike, which is 125mm. Then as rainfall decreases the payout increases linearly until rainfall reaches the exit of 40mm, then jumps to the policy limit of Rs 1000 once rainfall falls below 40mm.

BASIX insurance policies are sold in April and May, which are the months that precede the monsoon in India. Insurance policies cover only one season, so customers must purchase insurance again if they want coverage for the following year.

Table 2.2 presents summary statistics for the insurance policies studied.

2.3.2 Data

The data set consists of the entire set of BASIX's purchasers of rainfall index insurance from 2005-2007, which covers six states.[7] Though it ran small pilots in 2003 and 2004,

[7]The states are, in descending order of number of buyers: Andhra Pradesh, Maharashtra, Jharkand, Karnataka, Madhya Pradesh, Orissa.

Table 2.2: Policy Summary Statistics

Year	2005	2006	2007
Number of Policies	34	42	28
Average Premium For Three Phases (Rs)	283	295	287
Expected Payout (Using rainfall from 1961-2004)	76	73	80
Mean Ratio of Premium to Expected Payout	5.76	5.94	5.6
Mean Percentage of Times Policy Paid out From 1961-2004	10.2	6.8	7

BASIX began to mass-market rainfall insurance starting in 2005. The data contains limited personal information about each customer including their location, how many policies they purchased, and what payouts they received during that season. The BASIX data covers 42 weather stations, and includes a total of 19,882 customers from 2005-2007.[8] After numerous rainfall shocks in 2006, BASIX realized that many customers who had purchased only a small amount of insurance were disappointed that they received small payouts. In response, BASIX instituted a rule in 2007 that required all customers to purchase insurance coverage with a maximum payout of at least Rs 3000. This was meant to encourage people to buy a level of coverage that would actually provide meaningful payouts in the event of a shock, but resulted in a sharp decrease in the number of customers in 2007. A summary of characteristics of BASIX customers is given in Table 2.3.

For rainfall data, I use a historical daily grid of rainfall, which is interpolated based on readings from thousands of rainfall stations throughout India. This data is provided by the Asian Precipitation Highly Resolved Observational Data Integration Towards Evaluation of water resources.[9] This data set has daily readings of rainfall from 1961-2004, at a precision of .25°[10]. For each .25° x .25° block, the data contains the amount of rainfall in millimeters and the number of stations within the grid that contributed to the data. This data is used to evaluate how the insurance policies would have paid out historically, which can be used as a proxy for past rainfall shocks.

The initial challenge in processing the BASIX administrative data was to turn it into a panel. Although BASIX had three years of data, there was no way to identify unique individuals that purchased in multiple years. In order to solve this problem I matched names from year to year, taking into account the customer name, father/husband name, location, and age in order to identify which households were the same.[11]

[8] Note that BASIX also sold many policies in the district of Deogarh in Jarkhand, and those buyers are omitted from this analysis. The reason for this is that the policy for Deogarh is heavily subsidized, resulting in a policy that is completely different from all the others. For instance, the Deogarh policy for 2005 has an expected payout of Rs 1140 compared to an average of Rs 149, although the policy does not cost more than average. Because of its incredibly generous terms, the Deogarh policy has huge payouts for all years of the study, and therefore does not seem to be 'normal' enough to warrant inclusion in the main dataset. All the analysis below is performed excluding all buyers in Deogarh, though most results do not change substantially when it is included.

[9] APHRODITE's water resources project; http://www.chikyu.ac.jp/precip.

[10] .25° Latitude equals about 27.5km, while .25° longitude varies by latitude. Over the range of latitudes in this survey it equals roughly 26km.

[11] As district, village, and customer names had highly variant spelling, it wasn't possible to match

Table 2.3: Customer Summary Statistics

Year	2005	2006	2007
Number of Villages	954	1426	432
Number of Weather Stations	34	42	28
Number of Buyers	6428	10077	3377
Average Sum Insured (Rs)	3055	1612	3547
Buyers Receiving Payouts	351	1346	529
Average Payout	10.66	60.28	87.74
Average Payout (if Payout>0)	195.19	360.13	553.23
Buyers Who bought the Following Year	453	364	

The potential for matching errors causes a serious concern about the validity of the data set. Since a crucial part of my analysis revolves around determining what causes buyers to re-purchase insurance, determining who does so is extremely important. Despite a comprehensive effort which combined automated and manual matching methods, there are certain to be some errors in the dependent variable. While there is no reason to believe that this measurement error is correlated with any independent variables in the regression, since the dependent variable in some regressions is a dummy variable this can lead to downward bias on the estimated coefficients. In Section 4 I will explore the possible consequences of this problem.

A more serious problem with the data set is that it is not possible to observe the level of marketing that each person received, making "marketing intensity" an important omitted variable. When BASIX markets rainfall insurance, it first calls a group meeting in a village, and shows the villagers a video about rainfall insurance (and other BASIX products). It then speaks with visitors and answers questions. The BASIX team then makes a follow-up visit where it goes door to door, trying to sell BASIX products including rainfall insurance. Unfortunately, I have no data on on the specific marketing practices of each village and don't even know for sure in which villages BASIX actively sold rainfall insurance each year. As marketing intensity is potentially correlated with previous insurance outcomes, this may bias the estimates. This needs to be taken into account when performing the analysis and interpreting the results.

2.4 Results: The Effect of Payouts on Take-up

In this section I address the central question: is receiving an insurance payout correlated with repurchasing insurance the following year? To do this I examine BASIX's customers in 2005 and 2006, and regress repurchasing on payout reception and a year dummy. The basic econometric specification is as follows:

$$y_{i,t+1} = \alpha + \beta_1 P_{i,t} + B_2 D_{2006} + \epsilon_{t,i} \tag{2.8}$$

customers through the years using automated means.

Here $y_{i,t+1}$represents whether subject i purchases insurance at time $t+1$, and $P_{i,t}$ is a dummy variable that takes a value of 1 if person i receives an insurance payout at time t.[12] The sample is all buyers of insurance from 2005 and 2006, and I include a dummy (D_{2006}) that takes a value of 1 for purchasers in the year 2006 to control for time effects. Also, I only include purchasers who have weather insurance contracts available in their area in the following year.[13] These results are presented in Table 2.4, and Column 1 reports the baseline OLS results. It shows that receiving a payout is associated with a 9% increased chance of repurchasing insurance the following year, which means that those who receive an insurance payout are more than twice as likely to purchase insurance the following year as those who did not receive a payout. The dummy for 2006 is negative and significant, which is expected due to the minimum sum insured rules imposed in 2007.

One may be concerned that the linear probability model may give biased estimates, especially since such a small percentage of the sample were repeat buyers. Column 2 reports the results of our basic specification using a probit model. The marginal effect on receiving a payout is now 10%, which is very similar to that extimated by OLS.[14]

The loss aversion theory suggests that peoples' reference points may change when they experience a perceived gain from an insurance payout. Whether or not a customer perceives an insurance payout as a gain may depend on the amount of the payout in relation to the premium paid. In Columns 3 and 4 I add two new continuous variables to the regression: the ratio of the payout received to the premium paid (which I will call the "payout ratio") and the payout ratio squared. In both OLS and probit specifications, the payout ratio has a positive and strongly significant effect, while the squared term is smaller and negative. This suggests that higher insurance payouts result in greater propensity to purchase the following year, but that the marginal effects flatten out for larger payouts. Also, the simple dummy of receiving a payout flips to negative, suggesting that small payouts have a negative affect on purchasing. In fact, payouts have a positive effect only when the payout ratio nears 1. As it makes sense that customers would need to receive a net profit on their insurance transaction to experience a reference-changing "gain", this result fits in well with the loss aversion model.

One point of concern with these results is that there are many cases where there are multiple purchasers of insurance in a certain village in one year, and then zero in the next year. While this could be the result of people simply being unsatisfied with insurance, the large amount of villages that suddenly drop to zero purchasers is suspicious. As noted before, I don't know if BASIX marketed rainfall insurance in a particular village, or even

[12]It makes sense to assume that the error $\epsilon_{t,i}$ is correlated for the same person across time, as well as across people in a given year. Ideally, we would like to include individual fixed effects to account for individual heterogeneity. However, in order to exploit this variation we would need to look at customers who purchased insurance in both 2005 and 2006, and received payouts in only one of those years. Unfortunately, due to the very low repurchase rate, this results in very little variation and is therefore an unsuitable method of analysis.

[13]Basix's insurance coverage area varied somewhat from year to year. Results do not change significantly if all areas are included in the regression.

[14]Standard errors of all regressions are clustered at the village level. Clustering standard errors at the weather station level still yields highly significant coefficients.

TABLE 2.4: INSURANCE REPURCHASING

	Dependent Variable is Customer Re-Purchasing Insurance							
	Full Sample				Marketing Restricted Sample			
	OLS	PROBIT	OLS	PROBIT	OLS	PROBIT	OLS	PROBIT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Received Insurance Payout	0.0902***	0.107***	-0.0889***	-0.0611***	0.222***	0.254***	-0.195***	-0.133**
	(0.0244)	(0.0328)	(0.0203)	(0.0146)	(0.0442)	(0.0572)	(0.0666)	(0.0537)
Ratio of Payout to Premium			0.124***	0.0935***			0.246***	0.203***
			(0.0237)	(0.0156)			(0.0405)	(0.0396)
Ratio of Payout to Premium ^2			-0.0120***	-0.00868***			-0.0243***	-0.0197**
			(0.00264)	(0.00154)			(0.00409)	(0.00368
Year 2006 Dummy	-0.0251**	-0.0271**	-0.0396***	-0.0401***	-0.0269	-0.0246	-0.0336	-0.0326
	(0.0111)	(0.0115)	(0.0107)	(0.0116)	(0.0274)	(0.0287)	(0.0274)	(0.0289)
Marketing Restricted Sample	NO	NO	NO	NO	YES	YES	YES	YES
Observations	10776	10776	10776	10776	4194	4194	4194	4194
R-squared	0.013		0.033		0.035		0.057	

Robust standard errors in parentheses — All regressions include state fixed effects. Errors clustered at village level.
*** p<0.01, ** p<0.05, * p<0.1 — Probit regressions report marginal effects

if a certain village was visited by BASIX at all. For all the villages that had purchasers in one year and then none in the next year, it is quite likely that no BASIX representative visited the village, and therefore the customer didn't really have a chance to purchase the insurance. If this was the case it would make sense to exclude these villages from the analysis, as the previous year's payout would have no way to possibly influence a customer's purchase decision.

In Columns 4-8 I exclude villages that had no purchasers the following year from the analysis, creating what I call the 'Marketing Restricted Sample'. For instance, say village A had 10 purchasers in 2005, 13 purchasers in 2006, and 0 in 2007. In this case, the buyers from 2005 would be included in the sample since they obviously had opportunity to purchase the next year. However, the 2006 buyers would be excluded because I make the assumption that they didn't have the opportunity to buy in 2007. Restricting the sample this way results in a drop of the number of observations from 11002 to 4202, and causes the coefficient on receiving a payout in the baseline specification to more than double to .22. This lends some credence to the argument that the omitted information about whether a village received marketing was downward biasing the results.

The coefficients generated in this restricted sample may be incorrect, as the decision to market to certain villages and not others is most likely not exogenous. If the marketing teams decided whether or not to market to certain villages based on the previous year's rainfall or experience with insurance then the results could be biased. For instance, assume that there were a number of villages that experienced a rainfall shock but received very low payouts, making them unhappy with insurance. If the marketing team knew this they may have decided to not market to as many of these villages, therefore censoring villages that received a payout but were likely to have few repeat buyers. Regressions that use previous years' payout characteristics to try to predict whether insurance is sold in a village the following year do not reveal any patterns that would suggest selection bias, but they may miss more subtle selection patterns. It is possible that the coefficient for the marketing restricted sample is upward biased and it therefore would be reasonable to regard the coefficients in Columns 1 and 2 as lower and upper bounds respectively.

As mentioned earlier, the dependent variable in this regression was generated by manually matching customers from one year to another, and therefore may be measured with error which could bias the coefficients downward. In order to get a feel for the potential magnitude of this error I run simulations where I assume that the BASIX data has been matched completely correctly, and then induce 'measurement error' by randomly changing the dependent variable of whether people purchased the following year or not. With the introduction of 10% matching errors (with an equal probability of a mismatch for buyer or non-buyers), the coefficient on receiving a payout in the full sample (Column 1) falls from .090 to an average of .072 over 1000 simulations. For the marketing restricted sample in Column 2, it drops from .222 to .178. In other words, if we assume 10% matching errors, then the estimated coefficients are likely to be underestimated by around 20%. It also may be possible that most of the error came from being unable to identify positive matches, possibly due to different members of a household signing the insurance contract from year to year. Repeating the above simulation but assuming that only people who were found not to have bought the next year could have been errors, the coefficients become underestimated by around 10%. While the exact form and structure of the matching errors cannot be known, it is likely that the reported coefficients are somewhat lower (in absolute value) than the true coefficients.[15]

Overall, the results indicate that receiving an insurance payout correlates with a roughly 9-22% higher chance of repurchasing the next year compared with someone who purchased insurance but did not receive a payout. They also suggest that higher payouts lead to a greater chance of repurchasing, and that very low payouts may actually have a negative effect. While all these results are consistent with the theory of shifting reference points, they are also consistent with a number of other explanations, such as receiving payouts causing increased trust in the insurance company. The next section will attempt to empirically isolate some of these other mechanisms to see if they can explain the effects found in this section.

2.5 Alternative Explanations

While observing that people who receive insurance payouts are more likely to purchase insurance the following year is consistent with this paper's loss aversion theory, it is also consistent with many other explanations. In this section I attempt to isolate some of these other mechanisms to see whether they might instead be driving the results. I first consider the hypothesis that a rainfall shock as opposed to the insurance payout may cause people to be more likely to purchase the following year. To do this I look at villages in the first year they were offered insurance, and see if a rainfall shock the previous year correlates with greater insurance take-up. On the contrary, I find that villages that had a rainfall

[15] An alternative way the think about this is to realize that the regressions in Table 2.4 are generated using estimated variables, and therefore it would be appropriate to adjust the standard error to account for this fact. Since I don't have a good estimate of the extent of the error this is difficult to do quantititively, but the previous analysis suggests that the upper bound of the 95% confidence interval of the coefficient estimates is likely higher than what is suggested by the reported standard errors.

shock the previous year were actually less likely to purchase insurance the following year, which provides strong evidence against the argument that weather as opposed to payouts are driving the main result.

I next consider the widely hypothesized suggestion that receiving insurance payouts would cause people to gain trust in the insurance company and learn about insurance, therefore making them more likely to purchase insurance in the future. To do this I assume that in order to gain trust in the insurance company or learn how insurance works, one would not have to receive a payout themselves; witnessing a neighbor receive a payout should also have the same effect. I therefore look for evidence of spillovers within a village and do not find evidence that witnessing a payout without actually receiving it yourself has a significant effect on the propensity to purchase the following year.

I also consider the possibility that payouts cause increased take-up due to direct wealth and/or liquidity effects as opposed to psychological effects. While I do not have data to empirically separate these possible mechanisms, I argue that due to the timing and circumstances of rainfall insurance payouts, wealth and liquidity are unlikely to play an important role.

Finally, I address the concern of unobserved marketing variation. While the effect of this omitted variable is admittedly difficult to measure, I argue it is unlikely to be driving the central results.

2.5.1 Direct Effects of Rainfall

Since most rainfall insurance payouts come at the same time as a rainfall shock, it is possible that the rainfall shocks themselves as opposed to the insurance payouts are what is driving increased take-up the following year. There is some evidence for this happening in developed markets, as Kunreuther et al. (1985) note that purchases of flood and earthquake insurance in the US spike after a recent event, even if people were not insurance customers before.

There are a number of theories that could explain this behavior. First, recent experiences with rainfall could change subjects' beliefs about the probability of a rainfall shock the following year. If there is actual autocorrelation of rainfall events or if the subject has limited knowledge about rainfall shocks, people may update their beliefs about shocks and therefore have more desire for insurance the following year. Alternatively, recently experiencing a rainfall shock could make shocks more salient, increasing the chance they will buy insurance the following year. Also, rainfall shocks may affect the wealth of the farmers. If farmers become poorer due to bad rainfall, CRRA utility would suggest that they would be even more risk averse the next year as a second shock would cause greater disutility.

I start by examining whether there is actual autocorrelation in the rainfall data. To test for autocorrelation, I create a panel of various rainfall indicators from 1961-2004 for each weather station. For each indicator, I run a regression of six lags of the variable on the current value, including weather station fixed effects. These results are presented in Column 1 of Table 2.5, with just the coefficient on the first lag shown. While a fixed effects

Table 2.5: First Order Autocorrelation of Weather Variables

	Fixed Effects	Arellano-Bond
	(1)	(2)
Total Rainfall	-0.106***	-.086***
	(.030)	(.021)
Phase 1 Rainfall	-.090***	-.075***
	(.030)	(.029)
Phase 2 Rainfall	-.018	-.026
	(.030)	(.028)
Phase 3 Rainfall	-.029	.007
	(.030)	(.028)
Would Have Been Payout	.023	.017
	(.030)	(.022)
Total Insurance Payout	-.0353	.004
	(.030)	(.028)
Weather Station Fixed Effects	YES	YES

Coefficients reported are from separate univariate regressions
Observation are years 1967-2004 for Fixed Effects Regression
Observation are years 1962-2004 for Arellano-Bond Regression
Fixed Effects regression contains six lags, Coefficient of First Lag Displayed
Arellano-Bond Regression contains one lag
Standard Errors are in Parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

regression with a lagged dependent variable is not generally consistent, it will converge to the true value as $T \to \infty$. As T is relatively large (38), these estimates are likely to suffer from little bias. I also run a regression of the first lag using previous lags as instruments, using the methodology proposed by Arellano and Bond (1991), with results presented in Column 2. The results from both specifications are similar, and show a negative first-order autocorrelation in rainfall that appears to be driven by rains early in the season. The bottom two rows test for autocorrelation of rainfall shocks using the parameters of the 2005 insurance policy to determine shocks. "Would Have Been Payout" is a dummy variable that takes a value of 1 if the insurance policy of 2005 would have given a payout, while "Total Insurance Payout" is the size of this payout. By these measures, shocks do not appear to exhibit significant positive first-order autocorrelation.

This evidence casts doubt on the hypothesis that positive autocorrelation of weather events is driving increased insurance purchasing. It appears that total rainfall is actually negatively autocorrelated, while shocks (which are proxied by the insurance contract giving a payout) do not appear to be correlated at all.

Even if there is no positive autocorrelation of rainfall, there may be other aspects about experiencing a shock that result in people having a higher propensity to purchase insurance. In order to look at the results of weather separately from the effects of insurance, I analyze how previous weather events affected insurance purchase decisions in the first year that insurance was offered to BASIX customers, which was 2005. To accomplish this, I first aggregate the purchasing data to the village level and then test to see whether villages that experienced a rainfall shock in 2004 had more insurance purchasers in 2005 than villages who did not experience a rainfall shock. A shock is defined using each location's insurance policies in 2005: If insurance would have paid out in 2004 based on the structure of the 2005 weather policy, this is deemed a rainfall shock. As the quality of the rainfall data is

Table 2.6: Effect of Shocks on Purchasing

	Dependent variable is number of buyers in 2005			
	(1)	(2)	(3)	(4)
Would Have Been Payout in 2004	-3.843***	-4.592***	-5.045**	-3.788*
	(0.987)	(1.039)	(2.173)	(1.898)
Ratio of 2004 Payout to 2005 Premium			4.365	-0.755
			(4.610)	(5.543)
Payout Ratio Squared			-1.991	-0.279
			(1.814)	(2.064)
Constant	8.001***	0.651	7.985***	1.015
	(0.714)	(6.341)	(0.713)	(6.494)
Weather Station Constants	NO	YES	NO	YES
Observations	733	733	733	733
R-squared	0.073	0.094	0.075	0.097

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Obervations weighted by quality of rainfall data
Errors Clustered at Weather Station Level
All Regressions Include State Fixed Effects

related to the amount of nearby weatherstations, I weight the observations based on the number of nearby rainfall stations.[16] Also, I create a hypothetical payout ratio, similar to the "Ratio of Payout to Premium" variable presented in Table 2.4. This is the ratio of the amount that the 2005 policy would have paid out in 2004 divided by the premium of the policy.

The results of this regression are presented in Table 2.6.[17] Column 1 presents the baseline regression, which shows that villages that experienced a rainfall shock in 2004 actually had an average of 3.8 *fewer* purchasers in 2005. One worry with this regression may be that since the insurance policies and rainfall patterns of each location are different, the definition of a shock may vary from one place to another. Therefore, the estimates may be improved with the inclusion of location and policy-specific covariates, which I title "Weather Station Constants". In Column 2 I add controls for the historical average rainfall, historical rainfall standard deviation, the policy premium in 2005, historical average payout of the policy, and the percentage of historical years there would have been a payout. Note that all the "historical" data is calculated from 1961-2000. With the addition of these controls, the coefficients on having a rainfall shock in 2004 remains negative, and even

[16]The APHRODITE weather data provides information about how many local weatherstations contributed to a certain rainfall reading. Since some of the rainfall observations are likely to be more accurate than others, I weight them according to accuracy. If there are no rainfall stations contributing to the APHRODITE data within a $.75^o$x$.75^o$ grid around the desired BASIX weather station, the observation is given a weight of 1. If there is a least one weather station in this $.75^o$x$.75^o$ grid, the observation is given a weight of 1.5. If there is a rainfall station within the $.25^o$x$.25^o$ grid, the observation is given a weight of 2. The weighted results to not differ significantly from the unweighted results.

[17]Note that while it is reasonable to think that village-specific characteristics (such as village size) may have an effect on village-level insurance take-up, village-level co-variates are not included in the regression. When the regressions are run with the village characteristics from the 2005 Indian census, the coefficients of interest do not change significantly. Also, most village-level characteristics had insignificant coefficients, with the exception that a more literate population was correlated with higher takeup. Since village-level coefficients were only available for around 50% of the villages, these variables are not included in the main specifications.

decreases slightly.

Following previous results that suggest that the size of the insurance payout is important, in Columns 3 and 4 I include variables for the severity of the shock in 2004 using the ratio of the hypothetical payout to the premium (the payout ratio) and the payout ratio squared. In both specifications these variables are insignificant, suggesting that most of the variation in purchasing in 2005 is explained by our binary shock variable.

The main conclusion to be drawn from these regressions is that the data does not support the hypothesis that bad weather induces people to purchase insurance in the following season. If anything, it seems to decrease insurance purchases. I can only speculate on the reasons for this; it may be due to the fact that people recognize the actual negative autocorrelation of rainfall, or it may be that the rainfall shocks decrease the available liquidity to purchase insurance the following year. Regardless, this data provides relatively convincing evidence that the direct effect of weather is not causing people who receive insurance payments to purchase again the following year.

2.5.2 Trust, Learning, and Spillover Effects

It is also possible that the propensity to purchase insurance after receiving a payout results from learning about insurance and trusting the insurance company, as opposed to being a direct result of the payout. In order to separate the effects of trust and learning from that of receiving the payout, I make the assumption that if trust and learning are playing an important role in causing people to purchase insurance after they have received a payout, then we should be able to see a positive spillover effect of payouts within the village.[18] This is because one shouldn't need to actually receive a payout to gain the effects of trust and learning, as someone who witnesses a payout gains all the same information as someone who receives a payout. But witnessing a payout would not give the psychological effect of gain from the insurance company utilized in the loss aversion model.

To perform this analysis I aggregate all buyers to the village level, but divide them into two types: repeat buyers and new buyers, where repeat buyers are people who purchased insurance the year before. I then regress the number of each type of buyer on payout statistics and the total number of buyers in the previous year. When there is an insurance payout in the previous year, most of the repeat buyers the following year received money from the insurance company, while new buyers didn't receive anything.[19] These results are presented in Table 2.7.

In order to compare results with the main specification in Table 2.4, I again provide a dummy for whether there was a payout in the village along with a quadratic effect of the

[18]This makes the assumption that people in the village learn of others receiving payouts, which may not be true of people hide the money to guard against claims from friends and relatives. If people can only gain trust and learning by actually receiving a payout themselves, then then data gives us no way to separate trust and learning from other possible mechanisms of receiving a payout.

[19]Some buyers may not have received money if they bought one phase of the insurance policy but one of the other phases paid out. This happened in 427 cases, and removing these individuals does not change there results.

Table 2.7: New Buyers In A Village

	Dependent Variable is the Number of Buyers in a Village the Following Year					
	Panel A: All Villages			Panel B: Villages With At Least 1 Repeat Buyer		
	Total Buyers	New Buyers	Repeat Buyers	Total Buyers	New Buyers	Repeat Buyers
	(1)	(2)	(3)	(4)	(5)	(6)
Was Payout in Village	-0.334	0.376	-0.709**	3.353	5.799*	-2.446***
	(2.252)	(2.121)	(0.338)	(3.275)	(3.413)	(0.433)
Mean Ratio of Payout to Premium	1.319	0.481	0.838*	0.522	-2.131*	2.653***
	(1.330)	(1.050)	(0.463)	(1.094)	(1.095)	(0.168)
Mean Payout Ratio Squared	-0.135	-0.0617	-0.0729	-0.0832	0.166	-0.249***
	(0.141)	(0.106)	(0.0539)	(0.120)	(0.114)	(0.0167)
Number of Buyers in Village	0.131***	0.0795*	0.0512***	0.197***	0.0959	0.101**
	(0.0480)	(0.0434)	(0.0178)	(0.0530)	(0.0596)	(0.0427)
Year 2006 Dummy	-2.994*	-2.738*	-0.256	-5.231	-5.241*	0.0105
	(1.661)	(1.448)	(0.236)	(3.284)	(2.929)	(0.514)
Constant	3.445***	3.301***	0.144	10.48***	10.00***	0.483*
	(1.140)	(1.065)	(0.136)	(1.599)	(1.614)	(0.249)
Observations	1534	1534	1534	459	459	459
R-squared	0.061	0.047	0.118	0.084	0.069	0.285

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Errors clustered at the weather station level. All regressions include state fixed effects
Data is aggregated to the Village Level
Includes all villages in 2005 and 2006 where there was insurance coverage the following year

ratio of payouts to the premium. When aggregating the village data I used the mean of the payout ratios in the village to create a payout ratio for the village.[20] The overall results of the table tell a clear story: payouts drive repeat buyers but not new purchasers, showing few spillover effects. Column 3 shows how payouts affect the number of repeat buyers the next year, and the results are very consistent with the baseline results from Table 2.4. A dummy for whether there was any payout is negative and significant, but the payout size has a positive effect. This suggests that low payouts have a marginally negative effect on the number of repeat purchasers, but this effect flips to positive as the size of the payout ratio increases above 1. Column 2 shows the effect of payouts on new buyers in a village. Here all the payout coefficients are insignificant, but due to large standard errors I cannot reject that they are the same as the effects on repeat buyers.

In Panel B I restrict the analysis to villages that had at least one buyer the year after insurance outcomes, creating a sample analogous to the 'Marketing Restricted Sample' in Table 2.4. The logic behind this is that if a village had zero buyers it is likely that insurance was not marketed in the village in that year, and therefore customers did not have an opportunity to purchase insurance. Restricting the data set in this way gives a much clearer pattern. Column 6 now shows much stronger effects of payouts on repeat buying, though the pattern is the same as in Column 3. Small payouts have a negative effect, while increasing the payout ratio increases repeat buying. The squared term on the payout ratio is now negative and significant, indicating that high payout ratios have diminishing effects.

These coefficients are now all significantly different from the coefficients for new buyers found in Column 5. In fact, the coefficients in Column 5 flip signs, suggesting that payouts

[20]The results are not sensitive to using the mean, and are very similar using the median, maximum, and mode.

have the opposite effect on people who did not receive payouts. These results suggest that low payouts actually induce more new buyers, but that these effects decrease and then turn negative as the payout in the village increases. This effect isn't consistent with any of the theories I have advanced, but is especially inconsistent with the hypothesis that people are purchasing insurance after receiving a payout due to the effects of trust and learning, as we have seen that higher payouts increase peoples' propensity to purchase again the next year.

One important clarification of these results is that most of the potential "new buyers" living in a village that had experienced payouts would have also experienced uninsured rainfall shocks during the same season. Therefore it may be possible that there are effects of trust and learning, but they are outweighed by opposite effects of the weather. As we saw in the previous section, rainfall shocks tend to have a negative effect on insurance demand, so the (lack of) evidence of spillovers may be a result of a more complex interaction between trust/learning and direct effects of weather.

This fact may possibly explain the unexpected pattern of coefficients seen in Column 5 of Table 2.7. If heavy rainfall shocks (and therefore high insurance payouts) cause liquidity high constraints but low rainfall shocks allow an increase in trust without the liquidity constraints, this would be consistent with the results in Column 5 and provide some evidence for the existence of spillovers. But it is not clear why this mechanism would only exist in the restricted sample, and only for new buyers.

Overall, these results do not support the hypothesis that trust, learning, or any other effects of simply witnessing insurance payouts are driving increased purchasing. While it is possible that our measurements of spillovers are too crude and miss more subtle effects, it is telling that there is no sign of spillovers in villages that received the largest payouts. Since we do not see these spillover effects, this provides further evidence that increased purchasing of insurance is instead driven by the actual reception of money from the insurance company, and is consistent with the proposed loss aversion model.

2.5.3 Direct Effects of Payouts on Wealth and Liquidity

The previous two sections discount the possibility that trust, learning, or weather effects are driving the result that receiving an insurance payout is correlated with purchasing insurance the following year. This points to the actual reception of money from the insurance company as being the driving force behind greater takeup. However, this paper's model of loss aversion is not the only explanation that could explain why an influx of money could drive greater insurance uptake. Instead, one might think that receiving an insurance payout could directly affect choices the next year due to its effects on wealth and liquidity. For instance, if insurance were a normal good then increased wealth would result in greater insurance demand.[21]

While the BASIX data set does not offer the opportunity to test the direct effects of a cash payment separately from an insurance payout, there are a number of reasons why it is

[21] This is consistent with the empirical findings of Cole et al. (2010).

unlikely that wealth or liquidity effects are driving the results. Most importantly, insurance payouts are given in the context of a rainfall shock, which would most likely result in a loss of income. It may help to recall that the empirical results are being driven by variation in rainfall across locations, not by levels of insurance within a village. Therefore, for wealth effects to be driving the results, one would need to think that experiencing an insurance payout in the context of a rainfall shock resulted in people becoming wealthier than those people who didn't experience a shock at all. Given the fact that most buyers bought a relatively low amount of insurance coverage relative to their incomes, experiencing a rainfall shock, even when insured, would likely decrease future wealth. Therefore, wealth effects seem like a poor explanation as to why receiving payouts spur future insurance sales.

If people who received insurance payouts had a decrease in wealth it is also unlikely that receiving the insurance payout would increase their liquidity the next season. Insurance payments were generally made in January, while people had the opportunity to purchase insurance for the next season only in May. It is doubtful that these payments would have a lasting enough liquidity effect to influence insurance buying decisions five months later.

While I can't provide direct empirical evidence against the hypothesis that insurance payments drive increased take-up due to wealth or liquidity effects, given the structure and timing of insurance payments this explanation seems extremely unlikely.

2.5.4 Omitted Marketing Intensity

As mentioned earlier, the data set does not contain the exact marketing practices that BASIX undertook in each village in each year. If the intensity of marketing was correlated with both previous years' insurance payouts and current years' sales, this omitted variable could be biasing the results. For instance, assume that the marketing staff at BASIX think that people who have just received a payout are more likely to repurchase insurance. In this case, as the marketing team has limited resources, it may make sense for them to direct these resources towards the area of highest return, which would be people who have already received payouts. If this was the case, the increased take-up rates from people who received payouts could simply result from increased marketing attention from the BASIX team.

While the results could be picking up some of this effect, there are a couple of reasons I believe it is unlikely to be a significant factor. First, regressions of observable marketing factors (such as a dummy of whether there were any purchasers in the village) do not show any significant correlations with payouts. Next, the BASIX marketing staff claim to not give any special marketing treatment to previous payout recipients.[22] As they are trying to build long-term business, BASIX claims that they do not change their marketing practices for villages that have recently received a payout. Finally, if BASIX targeted payout recipients and they didn't really have a higher tendency to purchase, one would think that the marketing team would quickly learn that this strategy wasn't effective and would stop it. While I only observe two marketing cycles and erroneous beliefs could

[22]Conversation with Sridhar Reddy, Assistant Manager for Insurance at Basix, Jan 09.

survive throughout this short time span, it is telling that the effect of payouts on take-up is greater in 2006 than 2005, suggesting that the effect is increasing over time.[23] If it was caused by erroneous expectations of the marketing team, we would expect the effect to decrease over time. Overall, while I must accept the possibility that increased marketing is driving the results, I regard it as unlikely.

2.6 Conclusion

After receiving an insurance payout, customers of rainfall insurance in India are 9-22% more likely to purchase insurance again the next year. This behavior seems to be driven by actually receiving the money from the insurance company, and is consistent with a loss aversion model where previous insurance gains shift the subject's reference point. In this model, a subject views future insurance premiums as deductions from his previous gains, as opposed to a true loss. Therefore, after receiving an insurance payout, future insurance purchases are more attractive. While there are other possible explanations for this phenomenon, the BASIX customer data does not provide support that any of these other possible mechanisms are driving the results.

First, direct effects of a rainfall shock do not appear to drive insurance purchases. Looking at villages in their first year of insurance availability, I find that locations that experienced a rainfall shock the year before are actually less likely to purchase insurance. Next, I do not find evidence that repeat purchasing is driven by trust, learning, or any other effect that one would expect to spill over to other members of the community. Taken together, these results point toward the actual reception of money from the insurance company as the dominant driver of repeat purchases. While this does not constitute direct evidence of the loss aversion model presented in the paper, it does provide empirical evidence that is consistent with the predictions and mechanisms of the model.

This study brings to light a number of questions that would be ripe for future research. First of all, it would be interesting to understand whether insurance payouts have long-term effects on future purchases, and also whether payouts continue to have similar effects for people who have years of experience with insurance. To answer these questions one would need a data set with a longer time frame. Also, a longer data set could shed further light onto the question of whether customers learn about insurance over time. It is possible that people need a few years of experience with insurance to really learn about the product and gain trust in it, which would explain why this paper fails to see any spillover effects.

These results point to a number of policy recommendation for the Indian rainfall insurance market, and possibly for insurance markets in general. One of the main arguments made for the slow adoption of insurance in India is that people do not understand insurance and do not trust the insurance companies. If trust and learning were the crucial determinant of insurance adoption, then incentives could be given to encourage early adoption and over time as people witnessed and experienced payouts we would expect insurance

[23] Results not shown.

adoption to grow. This paper fails to find any evidence of increased trust and learning driving insurance decisions, which suggests that incentivizing early adopting is unlikely to quickly spur insurance take-up.

Historical evidence (as in Kunreuther et al. 1985) has suggested that an effective policy to spur insurance markets would be to target areas that have recently experienced a large shock. This paper does not support this notion in the case of rainfall index insurance in India, as places that recently experienced a shock were less likely to purchase insurance.

Instead I suggest that the mechanism that drives increased purchases after a payout is the feeling of winning money from the insurance company. This result is not entirely helpful in terms of policy prescriptions. If the propensity to purchase insurance after receiving a payout is truly due to the reception of money, in order to maintain customers the insurance company would have to give out significant payouts each year. Since this would result in losses for the insurance company, such a scheme would not be sustainable in the long run.

With relation to the future of rainfall index insurance in India, one stark result is that the raw numbers of continuing customers of insurance are very low, calling into question the sustainability of the product. Even among people who received payouts in excess of twice their premium in 2006, only 18% bought again in 2007. With the proportion of repeat buyers so low, one would have to assume that many people are not satisfied with their experience of insurance, which suggests that the product or marketplace will need to evolve in order to survive.

One factor to note is that this study looks at the first major scale-up of rainfall insurance in the world. Rainfall insurance is still a young product, and is still evolving to meet the needs of customers. One particular point of attention is the massive loading on most policies offered. As we saw in Table 2.2, the many BASIX insurance policies had premiums of up to six times the actuarially fair rate. With premiums this high, it is unsurprising that people are not signing up. Also, one may argue that the correlation between insurance payouts and crop outcomes were less than ideal in these early products. Around the world, index insurance policies are constantly evolving to better correlate with crop outcomes and avoid basis risk. While this study predicts that rainfall insurance in the form of BASIX's policies from 2005-2007 are likely to fail, it is quite possible that innovations in products and pricing can create an insurance product that better meets the needs of small scale farmers.

The insurance pilot in Gujarat described in Chapter 1 has conducted ongoing rainfall insurance marketing to the same sample from 2006-2011. As there have been some payouts in recent years, we hope to study the dynamics of these purchasers in a future paper. That study should lend more evidence to the dyanmics of insurance behavior, and will allow us to further test the loss aversion model presented in this paper.

2.7 Appendix

2.7.1 Solving the Model for the General Case

Section 2.2 of the paper solved our reference-dependent model for a specific case of $r_1 = 0$, $f(r_1, 1, 0) = r_1$. This is the simple situation where the reference point starts at zero and we assume that if someone has bought insurance but does not receive a shock then their reference point will not change. In this Appendix I solve for the case of a general r_1 and $f(r_1, 1, 0)$, though still keeping the restriction that $-X < f(r, I, S) < 0$, as this condition ensures nonlinear (and therefore interesting) utility functions over the relevant range. Remember, $B_2(f(r_1, 1, 1)) > B_2(f(r_1, 1, 0))$ means that the benefit of buying insurance in the second period is greater for an insurance customer in period one who has received an insurance payout than for a customer who did not receive a payout.

Proposition 2.2. $B_2(f(r_1, 1, 1)) > B_2(f(r_1, 1, 0))$ iff

$$\begin{cases} f(r_1,1,0) > f(r_1,1,1) > \frac{p-1}{p} f(r_1,1,0) - (1+\lambda)X & if \quad -(1+\lambda)pX < f(r_1,1,0) < 0 \\ f(r_1,1,0) < f(r_1,1,1) < -\frac{p}{(1-p)}(f(r_1,1,0) + (1+\lambda)X) & if \quad -X < f(r_1,1,0) < -(1+\lambda)pX \end{cases}$$

Proof. To prove Proposition 2.2 , we must proceed case by by case. Start with the case where $-(1+\lambda)pX < f(r_1, 1, 0) < 0$, which means that customers who don't get an insurance payout in period 1 have a reference point in period 2 greater than the insurance premium. In this case the second period benefit for a customer who does not receive an insurance payout is:

$$B_2(f(r_1,1,0)) = -\lambda\beta pX + (1-p)(\alpha-\beta)f(r_1,1,0) \tag{2.9}$$

For those who did receive a payout, benefits are:

$$\begin{cases} B_2(f(r_1,1,1)) = -\lambda\beta pX + (1-p)(\alpha-\beta)f(r_1,1,1) & if \quad -(1+\lambda)pX < B_2(f(r_1,1,1)) < 0 \\ B_2(f(r_1,1,1)) = -(1+\lambda)\alpha pX + p\beta X - p(\alpha-\beta)f(r_1,1,1) & if \quad -X < f(r_1,1,1) < -(1+\lambda)pX \end{cases} \tag{2.10}$$

Combining Equations 2.9 and 2.10 proves the first case of Proposition 2.2. Turning to the case where $-X < f(r_1, 1, 0) < -p\lambda X$, the benefits for people who did not receive a payout are

$$B_2(f(r_1,1,0)) = -(1+\lambda)\alpha pX + p\beta X - p(\alpha-\beta)f(r_1,1,0) \tag{2.11}$$

The benefits for those who did receive a payout are:

$$\begin{cases} B_2(f(r_1,1,1)) = -\lambda\beta pX + (1-p)(\alpha-\beta)f(r_1,1,1) & if \quad -(1+\lambda)pX > f(r_1,1,1) > 0 \\ B_2(f(r_1,1,1)) = -(1+\lambda)\alpha pX + p\beta X - p(\alpha-\beta)f(r_1,1,1) & if \quad -X < f(r_1,1,1) < -(1+\lambda)pX \end{cases} \tag{2.12}$$

Combining Equations 2.11 and 2.12 proves the second case of Proposition 2.2. ■

Proposition 2.3. $\overset{Max}{r} B_2(r) \to r = -(1+\lambda)pX$

Proof. In words, this means that a customer achieves the greatest benefit from insurance when the reference point is equal to the negative of the premium of the insurance policy.

For an arbitrary reference point r, the benefits of insurance can be calculated by combining Equations 2.2 and 2.3.

$$\begin{cases} B(r) = -\lambda p\beta X - r(\beta+\beta\alpha+p(\beta-\alpha)) & if \quad -(1+\lambda)pX < r \\ B(r) = -(1+\lambda)\alpha pX + p\beta X + rp(\beta-\alpha) & if \quad -(1+\lambda)pX > r \end{cases}$$

In the first case, $B(r)$ is maximized at the boundary of $r = -(1+\lambda)pX$ since $(\beta+\beta\alpha+p(\beta-\alpha)) > 0$. In the second case $B(r)$ is maximized at the boundary of $r = -(1+\lambda)pX$ since $(\beta-\alpha) > 0$ ■

Chapter 3

Weather Insured Savings Accounts[1]

3.1 Introduction

As rainfall index insurance seems theoretically valuable but receives low demand, we propose that the problem may lie not with the product itself but instead with its "packaging": perhaps insurance would be more attractive when bundled with a more familiar financial product. In fact, some of the most successful trials of rainfall insurance (in terms of take-up) have come when it has been tied to credit.[2] We alternatively propose the idea of a WISA (Weather Insured Savings Account), which is a financial product that combines features of a savings account with rainfall index insurance. Money invested in a WISA would be partially allocated to insurance, with the rest allocated to savings. A WISA would provide insurance payouts when a rainfall shock occurred, but would also allow money to accumulate regardless of the state of the world. This paper develops a model to understand the theoretical demand for different types of WISA, and then conducts a lab experiment to test participants' relative demand for insurance, savings, and WISAs.

We develop a simple two-period model with risk averse agents that shows how people would value different types of WISAs allocated to them, assuming that people have access to savings but not insurance apart from the WISA. We define the proportion of insurance to savings as a WISA's type, and define a consumer's valuation of a WISA to be the minimum amount of money they would be willing to accept (WTA) to give it up. The central prediction of the model is that there is an ideal WISA type for which a consumer has a maximum WTA, and WTA alway decreasse as one moves away from this ideal. This ideal type increases with the discount factor, and under certain conditions it increases with risk aversion.

We then present the results from a laboratory experiment in Gujarat, India that tests

[1]This chapter based on joint work with Jeremy Tobacman.

[2]The Weather Based Crop Insurance Scheme (WBCIS) of The Agriculture Insurance Company of India (AICI) saw large take-up of rainfall index insurance when it was required to receive agricultural loans. Similarly, the NGO Microensure provides weather insurance exclusively tied to loans. But in a cautionary note, Giné and Yang (2009) find that requiring insurance as part of a loan decreases demand for the loan.

these predictions. We invited 322 farmers into a computer lab in Ahmedabad where they were asked to assess their valuations of rainfall insurance policies, savings vehicles, and WISAs using the Becker-DeGroot-Marschak (BDM) mechanism. We measured the WTA for four financial products: pure insurance, 1/3 savings + 2/3 insurance, 2/3 savings + 1/3 insurance, and pure savings. Both the savings and insurance products used in the experiment were real, in that they offered significant monetary payouts to the participants that could be collected after the monsoon season. Contrary to the predictions of our model, we find that a strong plurality (39%) of participants value both pure savings and pure insurance more highly than any mixture of the two. Additionally, more risk averse farmers have a stronger preference for pure products, which again does not conform to theoretical predictions.

We test a number of alternative explanations for this phenomenon in an attempt to explain participants' preferences for pure products over mixtures. First, it may be possible that people value the WISAs less because they do not understand them as well as the pure products. We test whether the preference for pure products varies based on different framing strategies for the WISAs which vary in their complexity, and find that this does not have an effect, casting doubt on lack of understanding as a driver of our results. We also test the hypothesis that the results are driven by the expectation that small payouts are less likely to be collected, which could make the small guaranteed payouts of the savings/insurance mixtures less attractive. After the monsoon, farmers with higher payouts were not more likely to collect their their money than farmers with smaller payouts, making this explanation unlikely.

One way to explain the preference for pure products is to drop the assumption of concave utility, allowing for convex utility in the loss domain as proposed by prospect theory. The intuition behind this is that people may not value small insurance payouts, as they view it as an insignificant contribution to a large loss. Therefore they value the WISAs less, as they provide insignificant amounts of insurance coverage.

The idea to combine insurance and savings is inspired by a few strands of literature, as well by observing various insurance markets. Slovic et al. (1977) have suggested that many people view insurance as a form of investment, rather than a pure risk mitigation tool. As market priced insurance generally gives a negative return on the invested premium, insurance is clearly a poor investment, and people who view it as such will tend to be dissatisfied with standard insurance options. If consumers do view insurance as an investment, then it may make sense to design insurance products that provide a positive payment in most states of the world so that consumers feel they are getting some return on their investment. Even in a lab setting, identifying investment as a motivation for purchasing insurance is very difficult, and experiments have given mixed results. For instance, Connor (1996) finds strong evidence that people view insurance as an investment, but experiments by Schoemaker and Kunreuther (1979) do not support the claim.

Despite inconclusive results in the literature, the private insurance marketplace does provide many insurance products that offer a guaranteed return on the premium through policies that offer "no claims refunds". With this type of insurance, policy holders receive part (or all) of their premium refunded to them if they do not make an insurance claim.

One example is a "whole life" insurance policy, in which customers pay monthly premiums for life insurance, but receive a lump sum of all the premiums paid if they are still alive at a certain age. Customers pay extra for this service, and insurance companies make money off the ability to invest the held premiums.

If people choose "no claims refunds" policies, they show a preference for using insurance as a vehicle for savings. Similarly, people may also view savings as a type of insurance. Many studies have pointed towards preparing for potential income shocks as a primary motive for savings, especially in developing countries (Karlan et al., 2010; Rosenzweig, 2001; Fafchamps and Pender, 1997; Carroll and Samwick, 1997; Lusardi, 1998; Guiso et al., 1992). Despite the frequent usage of savings to protect against shocks, the meager savings of the rural poor are generally insufficient to guard against large aggregate shocks such as a drought. Townsend (1994) shows that rural villagers in India do a good job of informally protecting themselves against idiosyncratic shocks, but that they are still affected by aggregate shocks. In a survey of farmers participating in a rainfall insurance pilot in Andhra Pradesh, India, 88% listed drought as the greatest risk they faced (Giné et al., 2008). If people are saving primarily to protect against shocks yet these savings are not enough to buffer against the most important risk they face, they might find a savings account with an insurance component especially attractive.

While there are no products (to our knowledge) combining weather insurance with savings accounts, savings accounts offering other types of insurance do exist. In the 1990s the China Peoples' Insurance Company (CPIC) offered a savings account where customers received various types of insurance coverage instead of interest on savings (Morduch, 2006). This is potentially attractive to customers, as those with money illusion may perceive this as resulting in "free" insurance coverage, but has the drawback that small savings balances will result in minimal coverage. Similarly, many banks and credit unions in the West offer savings accounts that give the depositor auto, renters, or other types of insurance as benefits.

Savings accounts that offer some insurance in lieu of interest (such as the one offered by CPIC described above), and insurance policies offering "no claims refunds" can be seen as lying along a spectrum between insurance and savings. The CPIC savings accounts are mostly savings, while the "no claims refunds" policies are mostly insurance. Seemingly there is scope for these mixtures in many insurance markets, so finding the correct balance between savings and rainfall insurance can potentially result in a financial product that best meets farmers' needs for dealing with rainfall risk. The lab experiment in this paper attempts to determine which mix of savings and insurance farmers would prefer.

This paper will proceed as follows: Section 3.2 introduces a simple insurance demand model to explain how people choose between savings and insurance. Section 3.3 outlines the experimental procedure and provides summary statistics of our sample. Section 3.4 presents the results, and section 3.5 provides discussion of these results. Section 3.6 concludes.

3.2 Theory

This model provides a formal framework which shows how people would value different formulations of a WISA. In order to concentrate on the consumer's relative valuation of different WISA types, we consider a scenario where a consumer receives a gift of a fixed amount of money invested in a WISA. We then analyze how the certainty equivalent of this gift changes with the WISA type, and how the optimal WISA type varies with risk and time preference.

The model assumes that the consumer has access to savings outside of the experiment, and can therefore adjust his savings in response to any WISA or cash payment that he receives. However, it assumes that he does not have outside access to insurance.[3]

3.2.1 Basic Model Setup

A consumer lives in a two period world where he is subject to a negative income shock $\tilde{x}$ in the second period. In the world there are two types of investment technologies: standard savings, in which an investment of s in the first period pays net return Rs in the second period, and a WISA which consists of a mix of savings and insurance. The structure of the WISA is determined by the parameter $\gamma \in [0, 1]$, which determines the relative amount of savings and insurance that the WISA provides. An investment of w in a WISA results in $(1 - \gamma)w$ being invested in savings (which has the same interest rate R as standard savings[4]), and γw being invested in insurance. The insurance is standard proportional coinsurance (as in Schlesinger [2000]), where the premium is equal to the expected payout times $1+\lambda$, with λ being the loading factor. This means that if γw is invested in insurance, the customer receives a payout of $\frac{\gamma w \tilde{x}}{(1+\lambda)E(\tilde{x})}$ in the event of income shock $\tilde{x}$. We can define the payout from a WISA as follows:

$$g(w, \tilde{x}, \gamma) = (1 - \gamma)wR + \gamma w \frac{\tilde{x}}{(1 + \lambda)E(\tilde{x})} \tag{3.1}$$

It is worth noting that full insurance is achieved when $\gamma w = (1 + \lambda)E(\tilde{x})$. Since γ is bounded above by 1, if $w < (1 + \lambda)E(\tilde{x})$ there is no WISA which provides full insurance. To mimic our lab setup, we assume that the amount w invested in a WISA is both fixed and comes free of cost to the consumer.[5] The consumer chooses s, which is savings made outside of the WISA. The timing thus proceeds as follows:

1. The consumer is endowed with first period income Y_1, and chooses the amount of savings s. He consumes the rest of his income and realizes first period utility.

[3] We think this is a realistic assumption, as local farmers can save informally but have very few (if any) formal insurance options.

[4] It is easy to show that none of our results are sensitive to this assumption, but it maintained for clarity.

[5] While the imposition of w may seem strange, it was fixed in our laboratory experiment so that we could focus on varying γ. One could easily generalize the model to make w a choice variable.

2. Shock is realized and the consumer receives second period income $Y_2 - \tilde{x}$. He also receives returns of Rs from savings and $g(w, \tilde{x}, \gamma)$ from WISA. He consumes all income and realizes second period utility.

The participant has a concave utility function[6] $u'(c) > 0$,$u''(c) < 0$ and discount factor β. Expected utility U over the two periods is:

$$U = u(Y_1 - s) + \beta E(u(Y_2 - \tilde{x} + Rs + g(w, \tilde{x}, \gamma))) \tag{3.2}$$

The customer chooses savings s to maximize expected utility. We can define his indirect expected utility V as a function of his first period endowment Y_1 and the WISA payment function $g(w, \tilde{x}, \gamma)$ as follows:

$$V(Y_1, g(w, \tilde{x}, \gamma)) = \underset{s}{max}\, u(Y_1 - s) + \beta E(u(Y_2 - \tilde{x} + Rs + g(w, \tilde{x}, \gamma)))\ s.t.\ 0 \leq s \leq Y_1 \tag{3.3}$$

Define the optimal value of s as $s^*(\gamma)$. For simplicity define:

$$c_1 = Y_1 - s^*(\gamma)$$

$$c_2 = Y_2 - \tilde{x} + Rs^*(\gamma) + g(w, \tilde{x}, \gamma)$$

Assuming an interior solution (and valid second order condition), the following first order condition holds for $s^*(\gamma)$

$$\left.\frac{dU}{ds}\right|_{s=s^*(\gamma)} = -u'(c_1) + \beta RE(u'(c_2)) = 0 \tag{3.4}$$

We are interested in understanding how valuations of a WISA change as γ is varied. To do this, we define the willingness to accept (WTA) $A(\gamma)$, which makes a customer indifferent between receiving a monetary payment of $A(\gamma)$ or receiving an endowment of a WISA with parameter γ . By definition, $A(\gamma)$ satisfies the following equation:

$$V(Y_1 + A(\gamma), 0) = V(Y_1, g(w, \tilde{x}, \gamma)) \tag{3.5}$$

3.2.2 Characteristics of WTA

We are interested in how WTA changes with the WISA type. Applying the implicit function theory to Equation 3.5:

$$\frac{dA(\gamma)}{d\gamma} = \frac{\beta E((\frac{dg(w,\tilde{x},\gamma)}{d\gamma})u'(c_2))}{u'(c_1)} = \frac{\beta}{u'(c_1)}\left[E(u'(c_2))\left[\frac{1}{1+\lambda} - R\right] + \frac{1}{(1+\lambda)E(\tilde{x})}cov\{u'(c_2), \tilde{x}\}\right] \tag{3.6}$$

[6]We also assume the utility function is globally continuous and differentiable.

This expression reveals two effects.[7] Assuming $\frac{1}{1+\lambda} < R$, the first term represents loss from substituting away from savings, while the second term represents the gain from acquiring more insurance. $\frac{dA(\gamma)}{d\gamma}$ is of ambiguous sign, and its sign can change over the range of γ. However, we can still isolate other properties of $A(\gamma)$.

We next show that the function $A(\gamma)$ cannot contain any local minima. From the extreme value theorem, we know there must be a $0 \leq \gamma \leq 1$ which maximizes $A(\gamma)$ over this range. If $A(\gamma)$ has no local minima, $A(\gamma)$ weakly decreases as one moves away from this optimum γ.

Proposition 3.1. *$A(\gamma)$ has no local minima.*

Proof. We will show this in two steps:

1. Prove that if $\frac{d^2}{d\gamma^2}V(Y_1, g(w, \tilde{x}, \gamma)) < 0$, $A(\gamma)$ has no local minima.
2. Prove that $\frac{d^2}{d\gamma^2}V(Y_1, g(w, \tilde{x}, \gamma)) < 0$

Step 1: Show that if $\frac{d^2}{d\gamma^2}V(Y_1, g(w, \tilde{x}, \gamma)) < 0$, $A(\gamma)$ has no local minima.

Using the definition of $V(Y_1, g(w, \tilde{x}, \gamma))$ from Equation 3.5 and applying the envelope theorem:

$$\frac{dV(Y_1, g(w, \tilde{x}, \gamma))}{d\gamma} = \frac{d(Y_1 + A(\gamma), 0)}{d\gamma} = \frac{dA(\gamma)}{d\gamma}u'(Y_1 + A(\gamma) + s^*(\gamma)) \tag{3.7}$$

$$\frac{d^2V(Y_1, g(w, \tilde{x}, \gamma))}{d\gamma^2} = \frac{d^2A(\gamma)}{d\gamma^2}u'(Y_1+A(\gamma)+s^*(\gamma))+\frac{dA(\gamma)}{d\gamma}(\frac{dA(\gamma)}{d\gamma}+\frac{ds^*(\gamma)}{d\gamma})u''(Y_1+A(\gamma)+s^*(\gamma)) \tag{3.8}$$

In general, the sign of the second term in Equation 3.8 is unclear. Since $A(\gamma)$ is continuous and differentiable, at any local extrema $\frac{dA(\gamma)}{d\gamma} = 0$ and the second term goes to zero. At any extrema the following equation holds:

$$\frac{d^2A(\gamma)}{d\gamma^2} = \frac{d^2V(Y_1, g(w, \tilde{x}, \gamma))}{d\gamma^2}\frac{1}{u'(Y_1 + A(\gamma) + s^*(\gamma))}$$

When $\frac{d^2V(Y_1, g(w, \tilde{x}, \gamma))}{d\gamma^2} < 0$, $\frac{d^2A(\gamma)}{d\gamma^2}$ will be less than zero because $u' > 0$. This means that any local extremum must be a maximum, and therefore no local minimum can exist. Note that Equation 3.7 also shows that the γ which locally maximizes $V(Y_1, g(w, \tilde{x}, \gamma))$ will also locally maximize $A(\gamma)$.

Step 2: Prove that $\frac{d^2}{d\gamma^2}V(Y_1, g(w, \tilde{x}, \gamma)) < 0$

Using the envelope theorem:

$$\frac{d}{d\gamma}V(Y_1, g(w, \tilde{x}, \gamma)) = \beta E(\frac{dg(w, \tilde{x}, \gamma)}{d\gamma})u'(c_2)) \tag{3.9}$$

[7] Also, note that continuity and differentiability of the utility function guarantee that $\frac{dA(\gamma)}{d\gamma}$ is defined everywhere, and therefore $A(\gamma)$ is globally continuous and differentiable.

$$\frac{d^2}{d\gamma^2}V(Y_1, g(w,\tilde{x},\gamma)) = \beta E\big((\frac{dg(w,\tilde{x},\gamma)}{d\gamma})^2 u''(c_2)\big) + \beta R E\big(\frac{dg(w,\tilde{x},\gamma)}{d\gamma}\frac{ds(\gamma)^*}{d\gamma}u''(c_2)\big) \quad (3.10)$$

The first term is negative, but the second is of ambiguous sign. In order to sign the expression, we can leverage the first order condition for s. Applying the implicit function theorem to Equation 3.4, we get the following expression for $\frac{ds(\gamma)^*}{d\gamma}$:

$$\frac{ds^*(\gamma)}{d\gamma} = -\frac{\frac{d}{d\gamma}\frac{dU}{ds}}{\frac{d}{ds^*(\gamma)}\frac{dU}{ds}} = -\frac{\beta R E(\frac{dg(w,\tilde{x},\gamma)}{d\gamma}u''(c_2))}{u''(c_1) + \beta R^2 E(u''(c_2))} \quad (3.11)$$

Rearranging terms and multiplying both sides by $\frac{ds^*(\gamma)}{d\gamma}$ yields the following equation.

$$(\frac{ds^*(\gamma)}{d\gamma})^2 u''(c_1) + \beta E\big((R\frac{ds^*(\gamma)}{d\gamma})^2 u''(c_2)\big) + \beta E\big(R\frac{ds^*(\gamma)}{d\gamma}\frac{dg(w,\tilde{x},\gamma)}{d\gamma}u''(c_2)\big) = 0$$

As the above expression is equal to zero, we can add it to the right hand side of Equation 3.10.

$$\frac{d^2}{d\gamma^2}V(Y_1, g(w,\tilde{x},\gamma)) = \beta E\big((\frac{dg(w,\tilde{x},\gamma)}{d\gamma})^2 u''(c_2)\big) + \beta R E\big(\frac{dg(w,\tilde{x},\gamma)}{d\gamma}\frac{ds(\gamma)^*}{d\gamma}u''(c_2)\big) +$$

$$(\frac{ds^*(\gamma)}{d\gamma})^2 u''(c_1) + \beta E\big((R\frac{ds^*(\gamma)}{d\gamma})^2 u''(c_2)\big) + \beta E\big(R\frac{ds^*(\gamma)}{d\gamma}\frac{dg(w,\tilde{x},\gamma)}{d\gamma}u''(c_2)\big)$$

Collecting and factoring the terms under the expectation operator:

$$\frac{d^2}{d\gamma^2}V(Y_1, g(w,\tilde{x},\gamma)) = \big(\frac{ds^*(\gamma)}{d\gamma}\big)^2 u''(c_1) + \beta E\big((\frac{dg(w,\tilde{x},\gamma)}{d\gamma} + R\frac{ds^*(\gamma)}{d\gamma})^2 u''(c_2)\big) \quad (3.12)$$

Both terms are clearly negative due to the concavity of the utility function. Therefore, $\frac{d^2}{d\gamma^2}V(Y_1, g(w,\tilde{x},\gamma)) < 0$. This combined with the proof in Step 1 shows that $A(\gamma)$ cannot have any local minima. ■

3.2.3 Risk Aversion

Models of classical insurance demand (such as Schlessinger [2000]) predict insurance demand to be increasing in risk aversion. This is also true in our model. The following exposition closely follows the proof of Proposition 3 in Schlessinger (2000), which proves (in a world without external savings) that an increase in risk aversion increases insurance demand.

Proposition 3.2. $\underset{\gamma}{argmax}\, A(\gamma)$ *is weakly increasing in risk aversion*

Proof. Define a function $v(c)$ which is globally more risk averse (as defined in Pratt[1964]) than the original utility function $u(c)$. We would like to to understand how $A(\gamma)$ will differ for a person with utility function $v(c)$ compared with someone with utility function $u(c)$.

Pratt (1964) guarantees the existence of a function h such that $v(c) = h(u(c))$, $h' > 0$, $h'' < 0$. Analagous to the definition in Equation 3.3, we can define the indirect utility V_2 as the indirect utility for someone with utility function $v(c)$.

$$V_2(Y_1, g(w, \tilde{x}, \gamma)) = \underset{s}{max}\, h(u(Y_1 - s)) + \beta E(h(u(Y_2 - \tilde{x} + Rs + g(w, \tilde{x}, \gamma)))) \; s.t.\; 0 \leq s \leq Y_1 \tag{3.13}$$

Define γ^* as the value of γ that maximizes V (the indirect utility function with utility function $u(c)$). Assuming an interior solution γ^* will adhere to the following first order condition:

$$\frac{dV}{d\gamma} = 0 = \beta E(\frac{dg(w, \tilde{x}, \gamma^*)}{d\gamma} u'(c_2)) \tag{3.14}$$

For someone with utility function v, how does the choice of γ^* compare to their optimal γ? Taking the derivative of Equation 3.13, consider how utility changes when we increase γ above γ^*

$$\left.\frac{dV_2}{d\gamma}\right|_{\gamma=\gamma^*} = \beta E\left[h'(u(c_2))u'(c_2)\frac{dg(w, \tilde{x}, \gamma^*)}{d\gamma}\right] > 0 \tag{3.15}$$

To see why this expression is greater than zero, define F as the probabilty distribution of $\tilde{x}$ and assume that $0 < \tilde{x} < \bar{x}$. Consider the level of shock $x_0 = R(1+\gamma)E(\tilde{x})$, which is where $\frac{dg(w, x_0, \gamma^*)}{d\gamma} = 0$. Define $c_0 = Y_2 - x_0 + Rs^* + g(w, x_0, \gamma^*)$, which is the level of consumption in the second period when $\tilde{x} = x_0$. Substituting the above definitions into Equation 3.15 , we get

$$\left.\frac{dV_2}{d\gamma}\right|_{\gamma=\gamma^*} = \beta \int_0^{\bar{x}} \frac{dg(w, \tilde{x}, \gamma^*)}{d\gamma} h'(u(c_2))u'(c_2)dF >$$

$$\beta h'(u(c_0))\left(\int_0^{x_0} \frac{dg(w, \tilde{x}, \gamma^*)}{d\gamma} u'(c_2)dF + \int_{x_0}^{\bar{x}} \frac{dg(w, \tilde{x}, \gamma^*)}{d\gamma} u'(c_2)dF\right) = 0 \tag{3.16}$$

The right hand side of Equation 3.16 is equal to zero due to the first order condition that defines γ^* (assuming that the constraints on γ are not binding.) The inequality holds due to the fact that $h'(u(c_0))$ is always increasing in $\tilde{x}$. The integral on the right hand side has been split into two to show its negative and positive regions, which illustrates why the inequalilty holds.

Since V_2 is convex in γ (shown in Proposition 3.1), if $\frac{dV_2}{d\gamma}\Big|_{\gamma=\gamma^*} > 0$ then γ^* lies below the new γ which maximizes V_2. This means that an increase in risk aversion increases $\underset{\gamma}{argmax}\, A(\gamma)$.

We assumed before that the restrictions on γ^* were not binding. If one of the restrictions on γ^* is binding, then a marginal change in risk aversion will not change γ^*, making $\underset{\gamma}{argmax}\, A(\gamma)$ constant in risk aversion. ■

One important point is that our model does not include basis risk, meaning there is no chance that a customer suffers a shock yet receives no payout. With the index insurance used in our experiment, this is certainly possible. As discussed in Chapter 1, a model of insurance demand incorporating basis risk can cause insurance demand to decrease with risk aversion, potentially skewing the above results.

3.2.4 Discount Rates

Proposition 3.3. $\underset{\gamma}{argmax}\, A(\gamma)$ *is constant in* β

Proof. Define $\beta_1 > \beta$. As before, γ^* is the value of γ that maximizes expected utility when the discount rate is β. As before, we define a new indirect utility function V_3, for someone with a discount factor of β_1.

$$\frac{dV_3}{d\gamma}\Big|_{\gamma=\gamma^*} = \beta_1 E((\frac{dg(w,\tilde{x},\gamma^*)}{d\gamma})u'(c_2) = \beta E((\frac{dg(w,\tilde{x},\gamma^*)}{d\gamma})u'(c_2)) = 0$$

The right hand side of the equation zero due to the first order condition γ in Equation 3.14. Since $\frac{dV_3}{d\gamma}\Big|_{\gamma=\gamma^*} = 0$ regardless of the value of β_1 this means the γ^* does not depend on β. ■

Next, we present the results from a laboratory experiment in which we test the predictions of our model.

3.3 The Experiment

3.3.1 Laboratory Experiments

In order to test what is the optimal WISA type, we invited 322 farmers from rural areas surrounding Ahmedabad, India to participate in a laboratory experiment. The session was conducted entirely on a computer, where the subjects participated in games designed to elicit their preferences about risk, time, savings, and insurance. The participants were recruited using personal connections and are not meant to represent a random sample of Gujarati farmers. Since many of the participants were uncomfortable using computers, each participant was paired with an enumerator who read all the questions out loud and entered the answers into the computer. Summary statistics on the experimental population are presented in Table 3.1.

Table 3.1: Summary Statistics

Summary Statistics	
Personal Characteristics	
Sex is Male	100.00%
	(0)
Age	43.14
	(13.48)
Land Owner	86.65%
	(0.34)
Village Distance from Ahmedabad (Km)	23.20
	(3.12)
Have a Telephone	43.17%
	(0.50)
Decision Parameters	
Discount Factor between experiment and Post-Monsoon	0.78
	(0.65)
Estimate of Coefficient of Partial Risk Aversion	2.23
	(3.07)
Rainfall Risk Exposure	
If there was a severe drought during the upcoming monsoon season, would the income of you or your family be affected?	
Yes, A Lot	82.92%
Yes, A Little	16.46%
No	0.31%
Have Government Crop Insurance	12.73%
Roughly how much money could you gain from drawing on savings and selling assets if there was an emergency? (Rs.)	7574.34
	(3366.06)
Roughly how much money could you borrow if there was an emergency? (Rs.)	5559.78
	(3654.97)
If there was a serious drought in the upcoming monsoon, how would you and your family cope?	
Draw upon cash savings	36.65%
Sell Assets Such as Gold, Jewlery, Animals	40.68%
Rely on help from friends and family	49.38%
The government would step in to help	52.80%
Take a Loan	44.10%
Number of Respondents	322

Standard deviations are in parenthesis

The experiment consisted of two primary parts: eliciting risk and discount parameters, and then eliciting valuations for various WISA types. In this section we give a brief overview of the methodology used, while further details can be found in the Appendix. Risk preferences were calculated using a Binswanger lottery, as in Binswanger (1981). Subjects were asked to pick from a menu of lotteries where the payout would be determined by a (virtual) coin flip. At the end of the session the coin flip was performed, and subjects were paid their outcome on the spot. The maximum payout was Rs 200 (around US $4), which roughly corresponds to the wages of 2-3 days of agricultural labor.

Discount rates were calculated using a set of hypothetical questions about whether farmers would rather accept Rs 80 (around US $2) now or a certain sum later. The sum to be paid later is increased question by question up to a maximum of Rs 280. Assuming that the subject starts by preferring Rs 80 now and then at some point switches to preferring money in the future, we can establish bounds on the discount rate.

The central part of the experiment revolves around figuring out participants' valuations for four savings and insurance products by eliciting their willingness-to-accept (WTA) using a Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964). Establishing the WTA using BDM is relatively straightforward, and we will give an example of a participant playing the BDM game for an insurance policy. The participant is given an insurance policy as a gift, and is then asked how much money he would be willing to accept to give up the policy (this is his "bid"). He then draws a random "offer price". If the offer price is above the bid, the participant must give up the insurance and instead receives the amount of the offer in cash. If the offer is less than the bid, he simply keeps the insurance.

Under standard expected utility theory, the dominant strategy under BDM is to state the true minimum WTA, but there is plenty of criticism as to its ability to assess accurate valuations in practice. We covered various criticisms of the technique in Chapter 1. Regardless of BDM's ability to elicit "true" valuations for goods, we believe our experimental design warrants the use of BDM since it relies not on the absolute valuations of the products but instead their relative valuations. Even if BDM is giving biased valuations, our experiment will still be valid as long as those biases do not change based on the type of product being offered.[8]

In our experiment, farmers are asked to play the BDM game to elicit their WTA for gifts of savings (in the form of delayed payments), insurance, and mixes between the two. For instance, in one question the farmers are asked to express their WTA for a "large" rainfall insurance policy, which is equivalent to owning three units of the insurance policy described in the Appendix. The exact question text (translated into English) is as follows:

> Consider you have been given a gift of 1 large rainfall insurance policy. This policy can be purchased for Rs 180 and pays out a maximum of Rs 1500 in the event of bad rainfall. What is the minimum amount of immediate payment you would accept to give up the insurance policy? Our offer to purchase this policy

[8] This assertion may be debatable since Karni and Safra's critique of BDM is based around how subjects would express WTA for lotteries, and our WISAs are all different types of lotteries. However, it is difficult to think of any way that this critique could explain our central results.

> from you will be between Rs 10 and Rs 250. You would receive the payout at the end of today's session.

As Indian law has very strict regulations about the holding of deposits, we were not able to officially create savings accounts for the participants. However, we proxied for savings by giving the participants coupons which could be redeemed for cash after the monsoon, at the same time as insurance payouts are given. These guaranteed payouts are theoretically equivalent to giving the farmers a gift of a fixed-term savings account that matures after the monsoon. However, this strategy has the drawback that farmers may not perceive guaranteed payouts in the same way they would perceive savings.

Since the point of the exercise is to understand relative valuations of savings and insurance, the farmers are asked to give their valuations for four different products:

- A "large" insurance policy with maximum sum insured of Rs 1500
- A "medium" insurance policy with maximum sum insured of Rs 1000 plus a guaranteed payment of Rs 60
- A "small" insurance policy with maximum sum insured of Rs 500 plus a guaranteed payment of Rs 120
- A guaranteed payment of Rs 180 after the monsoon

The market price of Rs 500 of insurance coverage is Rs 60, making all the bundles of roughly equal monetary value.[9]

The participants gave their minimum WTA for each of these 4 bundles, and then the computer randomly selected one of the games to be played for real.[10] After the selection of the "real" game was made, the offer choice was shown, and the farmer either kept the bundle or was given money equal to the offer price.

We also undertook a framing experiment, where subjects were randomly shown one of the three descriptions of the financial products. The text shown above, given to 25% of the participants, is the "Bundle Frame", where the WISAs are described as an insurance projects plus a voucher for guaranteed money. 25% of the participants were shown the "Insurance Frame", in which the WISAs were presented as an insurance policy with a minimum payout equal to the voucher size. This frame was designed to mimic "no claim refund" insurance policies. The rest of the participants were shown the ICICI Bundle Frame, which is the same as the bundle frame. but explicitly mentions that the farmer could purchase the policy from the insurance company ICICI-Lombard. Full text of all these frames is given in the Appendix.

[9]The actual quoted price for the policy was Rs 66, but the prices were rounded to make comparisons between bundles easier.

[10]Additionally, for the first three bundles (the ones which contain some insurance), the subjects were asked to give their WTA under the circumstance that the money paid to give up the bundle would be paid not on the day of the experiment, but post-monsoon. These results are not reported in this paper.

We took a number of steps to ensure that the subjects understood the BDM game. Prior to playing the game with insurance, subjects played a BDM game to elicit their WTA for a bar of chocolate. This game was immediately resolved, with subjects receiving either real money (up to Rs 15) or keeping the chocolate bar.[11] During the BDM game for insurance products, the enumerators began by reading the question aloud. They then used props of sample insurance enrolment forms and/or vouchers to simulate the products being given as gifts. For instance, when a subject was playing the game for an insurance policy with a maximum sum insured of Rs 1000 and a guaranteed payment of Rs 60, they were physically handed two insurance policies with a sum insured of Rs 500 plus a voucher worth Rs 60. The enumerator then explained that they would keep this gift if the computer's offer was less than their bid, but would get the value of the offer at the end of the session if the computer's offer was greater than their bid. In informal conversations with the farmers after the sessions, all farmers who we spoke to claimed that they understood how the BDM game worked.

3.3.2 Delayed Payments

One challenge of conducting a laboratory experiment with products such as savings and insurance is that an experiment must take place in a short amount of time, while real insurance and savings products have delayed benefits. To increase the realism of the lab experiments we offered real financial products that paid out money after the monsoon. Delayed payouts (which proxied for savings) were delivered in the form of a voucher, which could be redeemed for cash by bringing the voucher to our Ahmedabad laboratory. Participants had two months after the end of the monsoon to come to Ahmedabad to redeem their vouchers.[12] They also had the options of sending the voucher with someone else to collect the money.

There are two main problems with making delayed payments a part of our experiment. First the participants may not believe that the lab would actually provide the promised delayed payouts. This belief would most likely hold for both voucher payouts and insurance payouts, which means that the relative valuations of insurance and savings would be unaffected. The second problem is that since the farmers live in surrounding villages, the cost of coming to redeem the voucher may be greater than the amount of money they would receive. We dealt with this problem by providing a large span of time to collect the vouchers, and also allowing participants to send friends or relatives to redeem the vouchers. Many farmers living in local villages travel occasionally to Ahmedabad to visit family or do business, and once in Ahmedabad the marginal cost of coming to our office to pick up the voucher would be quite low.

Once the vouchers were ready to be redeemed, we called all participants who had given us a phone number multiple times to remind them that they had money to pick up and

[11] Similar to the results shown in Table A1.3 in Chapter 1, the results of the chocolate game did seem to affect BDM bids for insurance.

[12] Specifically, they were told that they could redeem their vouchers after the Hindu holiday of Dashera, which corresponded roughly with the end of both the monsoon season and the insurance policy.

Figure 3.1: Average Willingness to Accept for all Products

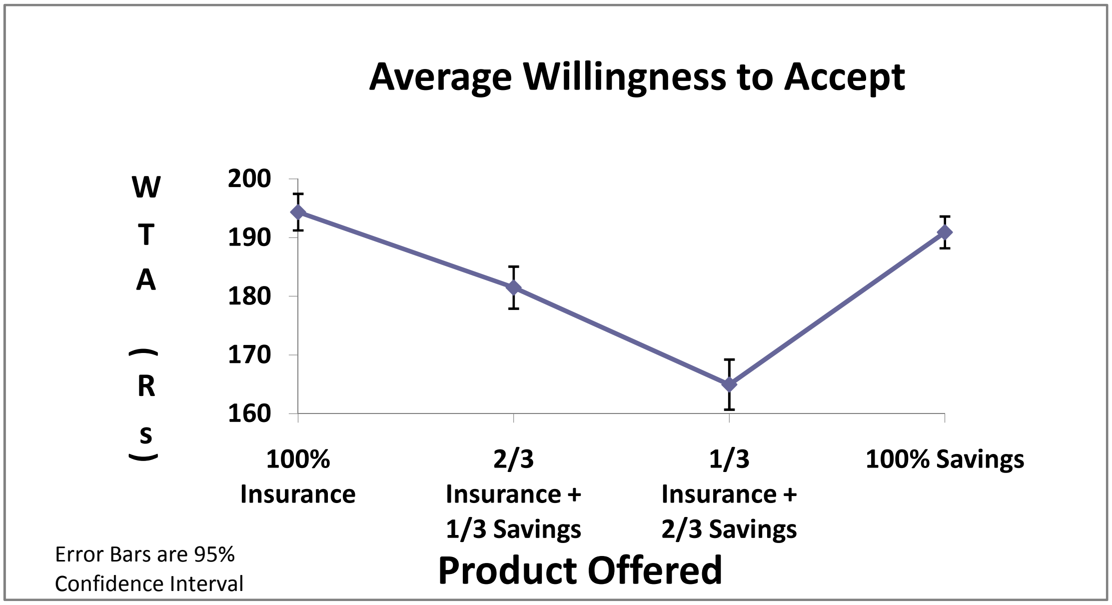

explained to them the procedure for redeeming the voucher. On the second phone call, we stressed that the farmer did not have to come in themselves to pick up the money but could instead send it with friends or family. Despite these attempts, only 42% of people with vouchers came to pick them up. We will discuss the possibility that uncertainty over whether the vouchers would actually be redeemed could be driving our results in Section 3.5.

3.4 Results

In this section we will present a number of results from the experiment. The first section will summarize the main results, which show average WTA of insurance, savings, and their mixtures, and will also explore the heterogeneity of the WTA patterns. We will then look at how the relative valuations of insurance and savings change with risk aversion and discount factors.

3.4.1 Main Results on Insurance and Savings Preferences

Our main empirical result is that most farmers prefer both pure savings and pure insurance to any mixture of the two. Figure 3.1 shows a plot of the average valuation versus the ratio of savings to insurance.

As Figure 3.1 shows, participants have the highest valuation of pure savings or pure insurance, with these bids being statistically indistinguishable. Valuations for both mixtures of savings and insurance are significantly lower than those for pure products. This graph

Table 3.2: Heterogeneity of Preferences

Preference	Percentage of Respondents
Indifferent	18%
Prefer Savings	7%
Prefer Insurance	13%
Prefer Mix	11%
Prefer Pure Product	39%
Other	12%

shows that the WTA has a local minimum in the percentage of insurance, which is a clear violation of Proposition 3.1 of our model.

Figure 3.1 indicates that, on average, subjects show a preference for either pure insurance or pure savings over a combination, but this average could obscure heterogeneity. To explore this further we group the subjects according to various patterns of the bids, which indicate distinct preferences over insurance or savings. These groups are shown in Table 3.2.

18% of respondents were indifferent, which means they had the same valuation for each product. 7% preferred savings, which means their bids were weakly decreasing in the percentage of insurance contained in the product. 13% showed a preference for insurance, meaning their bids were weakly increasing in the percentage of insurance contained in the product. 11% preferred a mix, which means they had the highest bid for one of the mixture products, with the bids weakly decreasing as one moves away from the highest bid. A strong plurality (39%) of the subjects had preferences that corresponded to the average, meaning they showed a preference for both pure insurance and pure savings over any of the mixtures. This means that the results are being driven by the within-person variation of bids, and therefore cannot be explained by heterogeneity of preferences across subjects. 12% of subjects did not express clear preferences, meaning that their bids changed directions twice as the percentage of insurance increased.

To get a more quantitative estimate of how product valuations vary based on the proportion of insurance we can use one observation for each bid, and regress the WTA on the percentage of insurance in the product. Based on the curvature of bids seen in Figure 3.1, we adopt a quadratic functional form for the percentage of insurance. Results are shown in Table 3.3. Column 1 contains only the linear term of the percentage of insurance, and we find it enters positively and significantly. In Column 2 we add the squared term, and now the linear term is negative while the squared term is positive, which is consistent with the U-shape seen in Figure 3.1.

3.4.2 Risk and Time Preferences

Our theoretical model predicted that people with higher levels of risk aversion or higher discount factors would have an optimal WISA with a higher percentage of insurance than

Table 3.3: Proportion of Insurance and Willingness-to-Accept

	Dependent Var is WTA Bid	
	(1)	(2)
Percentage of Insurance	0.0820**	-0.783***
	(0.0302)	(0.101)
Percentage of Insurance Squared		0.00864***
		(0.000970)
Constant	178.8***	188.5***
	(1.874)	(2.104)
Observations	1,288	1,288
R-squared	0.695	0.718
Robust Standard errors in parentheses	Percentage of Insurance is 0-100	
*** p<0.01, ** p<0.05, * p<0.1	Individual Fixed Effects Included	

those with lower risk aversion. As the previous section showed that most people do not have a unique optimal WISA type, we know that it is going to be impossible to measure how the movement of this optimum changes with risk or time preferences. However, in Table 3.4 we look at how risk and time preferences affect relative WTA.

We first look at the direct correlations between WTA and risk and discount factors for all products. While we did not address the direct effects of these parameters on WTA in the theory section, it is trivial to show that the theory predicts higher WTA for people with higher discount factors, while there are no clear predictions for risk aversion. Column 1 shows that people with higher discount factors have higher WTA and those who are more risk averse have lower WTA. In Column 2 we interact the risk parameter with the proportion of insurance offered, and find that people with higher risk aversion have a stronger preference for pure products over mixtures. Column 3 interacts the proportion of insurance with the discount rate, and the interaction terms have no significance. Column 4 includes both risk and discount factors.

The main conclusion from this analysis is that people who were more risk averse tended to have a relative preference for pure products as opposed to the mixtures. Time preference had no significant effect on relative preferences.

3.5 Discussion

Our theoretical model predicted that participants' WTA would not have a local minimum in the WISA type. However, the results showed that most people preferred pure savings and pure insurance to any mixture of the two, which is not a result not anticipated by our expected utility model. Looking at the heterogeneity of preferences in Table 3.2, we see that only 49% of the respondents gave results consistent with our theoretical model.

In the next section we consider a few possible explanation for these unexpected results.

TABLE 3.4: RISK FACTORS

	Dependent Variable is Willingness to Accept			
	(1)	(2)	(3)	(4)
Percentage Insurance	-0.843***	-0.625***	-0.789***	-0.512**
	(0.131)	(0.186)	(0.226)	(0.251)
Perc. Insurance Squared	0.00926***	0.00726***	0.00897***	0.00645***
	(0.00129)	(0.00184)	(0.00222)	(0.00248)
Perc. Ins. X Risk Aversion		-0.0977*		-0.102**
		(0.0508)		(0.0507)
Perc. Ins. Sq X Risk Aversion		0.000896*		0.000924*
		(0.000505)		(0.000505)
Perc. Ins X Discount Factor			-0.0698	-0.135
			(0.233)	(0.234)
Perc. Ins Sq X Discount Factor			0.000371	0.000967
			(0.00233)	(0.00233)
Risk Aversion	-2.283**			
	(0.989)			
Discount Factor	10.81**			
	(4.907)			
Constant	187.0***	190.2***	190.2***	190.2***
	(5.445)	(2.245)	(2.263)	(2.244)
Individual Fixed Effects	NO	YES	YES	YES
Observations	1,104	1,104	1,104	1,104
R-squared	0.057	0.715	0.713	0.715
Robust standard errors in parentheses	All Errors custered at Individual Level			
*** p<0.01, ** p<0.05, * p<0.1	Percentage of Insurance Ranges from 0-100			

3.5.1 Uncertainty Aversion

One explanation for our results could be that people were simply confused about the WISAs, as they are more difficult to understand than the pure products. If this was true, uncertainty about the WISAs could cause people to value them less due to uncertainty aversion, as in Gneezy et al. (2003).

In order to provide an empirical test on whether uncertainty was driving our results, we can draw some information from our results on framing. As explained earlier, participants were shown one of three frames describing the WISAs: the Bundle Frame, Insurance Frame, and ICICI Bundle Frame. While the Bundle Frames explained the financial product as an insurance product plus a voucher, the insurance frame explained them as simply an insurance policy with a minimum payout. Arguably, the insurance frame is much simpler to understand, as it presents the farmers with just one product instead of two. If this is true, we would expect the preference for pure products to be greater for participants shown the Bundle Frames as opposed to the Insurance Frame. However, as shown in Table 3.5, framing had little effect on participants' bids.

In Column 1 we introduce dummies for the frames directly to see how they affected subjects' average bids (the Bundle Frame is the omitted category). We see that compared to the Bundle Frame, the other two frames received modestly higher bids, but only the ICICI Frame is significant (and even then, only marginally so). In Column 2 we interact dummies for each of the frames with the percentage insurance (linear and squared) to see whether the framing affects the relative valuation of savings versus insurance. None of the coefficients are significant, indicating that framing of the question had little effect on the

Table 3.5: Framing

	(1)	(2)
Percentage Insurance	-0.796***	-1.016***
	(0.120)	(0.294)
Percentage Insurance Squared	0.00877***	0.0113***
	(0.00119)	(0.00297)
Insurance Frame X Percentage Insurance		0.281
		(0.380)
Insurance Frame X Perc Insurance Sq		-0.00269
		(0.00381)
ICICI Frame X Percentage Insurance		0.324
		(0.356)
ICICI Frame X Perc Insurance Sq		-0.00409
		(0.00355)
ICICI Frame	12.48*	
	(7.030)	
Insurance Frame	8.904	
	(7.646)	
Constant	180.6***	188.6***
	(5.282)	(2.058)
Individual Fixed Effects	NO	YES
Observations	1,288	1,288
R-squared	0.033	0.720
Robust standard errors in parentheses	Errors Clustered at Ind. Level	
*** p<0.01, ** p<0.05, * p<0.1		

results.

Finally, one would certainly think that rainfall index insurance itself is a confusing product compared to a simple savings voucher. If uncertainty aversion was driving the results, the WTA for pure insurance should be lower than that of pure savings. However, the mean WTA for these two products are not statistically distinguishable.

While we do not have explicit tests for whether uncertainty aversion caused subjects to value the WISAs less than pure products, the available evidence suggests that this is not the case. Regardless, we must consider uncertainty aversion as a possible explanation for our results.

3.5.2 Probability of Redeeming Vouchers

Another possible explanation for why participants valued the WISAs less than pure products is that they held certain expectations about their chances of picking up vouchers. After the experiment, 197 participants (61% of total) were given a voucher that could be redeemed for cash after the monsoon, and only 42% of these people eventually redeemed their vouchers. If during the experiments the farmers took into account the possibility that they might not redeem their vouchers, this could have affected the valuations of the financial products. Specifically, if they thought that the voucher size would affect their probability of actually redeeming their vouchers, this could affect their relative valuations of savings and insurance.

One way to think about this is to assume that farmers have fixed costs for redeeming the voucher. For instance, assume that a farmer anticipates that he will not redeem any voucher worth less than Rs 130. As both savings/insurance mixtures contained guaranteed payments of less than Rs 130, the participant would not redeem the voucher in the event the insurance did not pay out, making the bundles relatively unattractive. This could explain why WTA for the financial products decreased with the percentage of insurance until the voucher was worth Rs 180.

We can test this theory by taking a look at whether the chance of picking up the vouchers was influenced by the size of the voucher to be picked up. These results are presented in Table 3.6.

In Columns 1 and 2 of Table 3.6 we see that the total voucher amount held by the individual does not have a positive effect on the chance that the farmer redeems the voucher. In fact, the point estimates are all negative. In Column 1 we include village-level fixed effects, while in Column 2 we include village-level controls for the distance a farmer lives from Ahmedabad and the total amount of vouchers to be redeemed in the village, as we expect that these factors would influence the probability that they came to redeem the voucher. However, coefficients on these variables are insignificant.

Farmers who had vouchers waiting to be redeemed were called on the telephone two times to remind them to get their vouchers (if they had provided a telephone number at the time of the experiment).[13] In the second call, farmers were explicitly reminded that they

[13] 43% of participants gave us phone numbers.

Table 3.6: Voucher Size and Picking Up Payouts

Dependent Variable:	Voucher is Picked Up		Voucher is Picked Up Before Second Phone Call		Voucher is Picked Up After the Second Phone Call	
	(1)	(2)	(3)	(4)	(5)	(6)
Log Amount of Voucher	-0.0989	-0.101	0.0436	0.0505	-0.143**	-0.147*
	(0.0789)	(0.0843)	(0.0329)	(0.0309)	(0.0660)	(0.0839)
Log of Total Vouchers in Village		0.0618		0.0292		0.073
		(0.122)		(0.0903)		(0.146)
Distance from Ahmedabad		-0.0281		-0.0119		-0.0262
		(0.0228)		(0.0150)		(0.0249)
Have Phone Number	0.0828	0.136	0.0324	-0.00294	0.0504	0.155
	(0.0704)	(0.0879)	(0.0603)	(0.0433)	(0.0676)	(0.0993)
Log of Total Village Vouchers Remaining After Second Phone Call						0.0733
						(0.146)
Constant	0.853**	1.066	-0.0881	-0.0312	0.942***	1.077
	(0.377)	(1.223)	(0.151)	(0.718)	(0.314)	(1.216)
Village Fixed Effects	YES	NO	YES	NO	YES	NO
Observations	197	197	197	197	197	171
R-squared	0.528	0.072	0.372	0.024	0.561	0.083

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

did not need to show up in person to redeem their voucher but could instead send it with a friend or relative. Therefore it may be possible that before this phone call farmers with low voucher amounts were less likely to redeem them, as they thought they had to come to Ahmedabad themselves. In fact, after the second phone call there were a few instances where one farmer from a village collected many vouchers and came to redeem them all.

In Columns 3 and 4 the dependent variable is a dummy which takes the value of 1 if the farmer redeemed his voucher before the second phone call and zero otherwise. In these regressions the coefficient on voucher size is positive, though it is not significantly different than zero. In Columns 5 and 6 the dependent variable is a dummy which takes a value of 1 if the farmer redeemed his voucher after the second phone call and zero otherwise. Here we see that farmers with lower voucher sizes are more likely to redeem their vouchers after the second phone call, reflecting the fact that in this period some villages pooled a number of small vouchers and sent a single representative to redeem them. However, the total amount of vouchers still to be redeemed in the villages does not enter significantly.

The above regressions do not exactly paint a clear picture, but they would be consistent with an argument that farmers did not expect to redeem small vouchers, but only did so when reminded on the phone that they could send them to Ahmedabad with a representative. This could mean that when they were formulating their WTA for the financial products, they considered the chance of actually redeeming vouchers to be increasing in voucher size.

If we assume that the fixed costs to redeem vouchers are heterogeneous, we would expect that farmers who did not end up redeeming their vouchers anticipated higher costs and therefore would have expressed a higher preference for pure products (and lower valuations of all products). We can test these predictions by looking at the WTAs of people who actually picked up their vouchers versus those who did not pickup. We test this hypothesis in Table 3.7.

Column 1 looks at the direct effect of picking up the voucher on WTA. Here we see that people who picked up their vouchers had an average WTA that was Rs 20 lower than those

Table 3.7: WTA and Picking Up Vouchers

	Dependent Variable is WTA		
	(1)	(2)	(3)
Proportion of Insurance	-0.645***	0.0140	-0.324
	(0.143)	(0.0493)	(0.207)
Proportion of Insurance Squared	0.00683***		0.00338*
	(0.00136)		(0.00195)
Picked Up Voucher X Prop of Insurance		0.0571	-0.752**
		(0.0747)	(0.331)
Picked Up Voucher X Prop Insurance Sq			0.00809**
			(0.00316)
Picked Up Voucher	-20.10***		
	(7.401)		
Constant	202.2***	186.0***	193.6***
	(5.130)	(1.853)	(2.641)
Individual Fixed Effects	NO	YES	YES
Observations	788	788	788
R-squared	0.040	0.741	0.760
Robust standard errors in parentheses	Errors Clustered at Indiv. Level		
	Sample Includes only those who		
*** p<0.01, ** p<0.05, * p<0.1	had vouchers to pick up		

who did not pick up their vouchers. This is counterintuitive, as we would think that people who were less likely to pick up delayed payments would have lower WTA for all products. In Column 2 we look at how WTA varies with the proportion of insurance, and interact the proportion of insurance with a dummy for those who picked up their vouchers. We find that people who picked up the voucher had no significant difference on the relative valuation of savings versus insurance. The point estimate is positive, suggesting that people who were more likely to pick up their vouchers were more likely to prefer insurance, which is not in line with our hypothesis. Finally, in Column 3 we introduce the square of the proportion of insurance and also interact this term with the dummy for picking up the voucher. Both interaction terms are significant and very large, suggesting that people who picked up their vouchers showed a higher preference for pure products. Overall the results of Table 3.7 do not support the hypothesis that peoples' bids were affected by taking the transaction costs of redeeming the vouchers into account.

One final point is that if customers expected their chance of receiving delayed payouts to be increasing in the size of those payouts, we would expect that the WISA with the smallest voucher to have the lowest valuation. This is the 1/3 Savings + 2/3 Insurance product, which has a voucher of only Rs 60. However, the bids for the 2/3 Savings + 1/3 Insurance bundle (which has a voucher of Rs 120) were significantly lower even though the voucher size was larger.

Overall, the empirical evidence does not support the hypothesis that our results are driven by participants' consideration of whether or not they will actually receive delayed payouts.

3.5.3 Diminishing sensitivity

The main assumption that drives our theoretical result that $A(\gamma)$ cannot have a local minimum is the concavity of the utility function. If we relax this assumption, then the theory would allow local minima for $A(\gamma)$. While the assumption of risk averse agents is generally standard, prospect theory (Kahneman and Tversky, 1979) predicts that people have diminishing sensitivity around a reference point, which results in risk-seeking for losses. For someone with diminishing sensitivity, the utility function for someone with reference point r satisfies the following:

$$\begin{cases} u''(c) < 0 & if\ c > r \\ u''(c) > 0 & if\ c < r \end{cases}$$

If people exhibited diminished sensitivity around a reference point, this means that their utility function is convex for losses below a reference point, and Proposition 3.1 fails to hold. In order to see this, let's take a look again at the central result of Proposition 3.1. Define the reference level of consumption in each period to be $r_1$and r_2 respectively. Equation 3.12 now becomes:

$$\frac{d^2}{d\gamma^2}V(Y_1, g(w, \tilde{x}, \gamma)) = \big(\frac{ds^*(\gamma)}{d\gamma}\big)^2 u''(c_1 - r_1) + \beta E\Big(\big(\frac{dg(w, \tilde{x}, \gamma)}{d\gamma} + R\frac{ds^*(\gamma)}{d\gamma}\big)^2 u''(c_2 - r_2)\Big) \tag{3.17}$$

In a world where people exhibit diminishing sensitivity, $u''(c)$ is no longer universally less than zero, so the above expression is not necessarily negative. Instead, the sign will be determined by the specific shape of the utility function and the choice of the reference point.

In the first period, we can consider the reference point r_1 to be the amount of first period consumption in a world where the consumer has not recieved a gift of a WISA. If the gift of a WISA causes the consumer to increase (decrease) savings, then $c_1 - r_1$ will be less than (greater than) zero. Unfortunately, the model does not contain clear predictions about how the gift of the WISA will change savings, and therefore the first term has ambiguous sign.

In the second period, the choice of the reference point is less clear. One reasonable choice would be the level of consumption reached there if was no gift of a WISA and when $\tilde{x} = E(\tilde{x})$.[14] In this case, second period consumption can be above or below the reference point, and therefore the utility function is neither globally convex or concave, making the second term also ambiguous in sign.

We can resolve this ambiguity with a few simplifying assumptions. Assume that savings is fixed and that the second period reference point is the level of consumption when $\tilde{x} = 0$. In farming situations, this reference point is not unrealistic, as losses may come during rare catastrophic events while most seasons bring good harvests. In this scenario, we can drop the first term of Equation 3.17 as first period utility is always equal to reference utility.

[14]Using expectations as references is suggested by Köszegi and Rabin (2006).

The second term is positive, as $c_2 - r_2$ is always either zero or negative.[15] In this scenario, $A(\gamma)$ can have a local minimum.

In general, the necessary conditions for $A(\gamma)$ to have a local internal minimum is that there is a γ over the range of $0 < \gamma < 1$ that solves the first order condition for γ (found in Equation 3.9), and also satisfies the following second order condition :

$$\big(\frac{ds^*(\gamma)}{d\gamma}\big)^2 u''(c_1 - r_1) + \beta E\Big(\big(\frac{dg(w, \tilde{x}, \gamma)}{d\gamma} + R\frac{ds^*(\gamma)}{d\gamma}\big)^2 u''(c_2 - r_2)\Big) > 0$$

The intuition behind this effect is as follows. When people have diminishing sensitivity to losses, partial insurance is especially unattractive because the marginal utility of wealth is very low after a large loss. For instance, a person would be willing to pay less than half the premium for an insurance policy which offered half coverage (compared to full insurance). Therefore, the low amount of insurance offered as part of a WISA is unattractive, making the WISA unattractive overall compared to the pure products.

Our experiment does not shed light on whether the above necessary conditions are satisified for people who showed a local minimum in $A(\gamma)$. However, results of our experiment are consistent with predictions of a model with agents who exhibit diminishing sensitivity around a reference point. This would be an interesting topic for further research.

3.6 Conclusion

This study has explored Indian farmers' relative preferences for savings and insurance when planning for the monsoon season. We found that, contrary to theoretical predictions, most farmers preferred both pure savings and pure insurance to any mixture of the two. This finding suggests that a combined savings/insurance product such as a WISA would not be an attractive product for most Indian farmers.

Although the reasons for these choices are not entirely clear, we propose a couple of primary explanations for the preference for pure products. First of all, it is possible that farmers were uncertainty averse and confusion about the WISAs caused them to value them less. This suggests that introducing a complex financial product such as a WISA is likely to be unsuccessful.

Alternatively, lower valuation of mixed products would be consistent with a model where participants experience diminishing sensitivity to wealth changes around a reference point. People who have diminishing sensitivity to losses would show a strong preference for full insurance over partial insurance. If this was true, then a WISA would be an inherently unattractive product as it would provide less insurance than a pure insurance product.

There are a number of drawbacks of this experiment that may cause it to underestimate the potential demand for a WISA. One possibly attractive feature of a WISA is that it

[15] Note that $u''(c-r)$ is technically undefined when $c = r$. However, we can finesse this issue by assuming that in a world where savings does not adjust, first period utility will always be zero and should therefore be removed from the indirect utility function altogether. For the second term, we simply consider the expectation for all situations where $c_2 \neq r_2$.

could be framed as a savings account and the insurance payments could just be paid using foregone interest payments. Unfortunately we were unable to adequately create this scenario in the laboratory, as this scheme requires relatively large deposits to provide meaningful coverage, and Indian banking regulations prevented us from acting as a bank and actually opening savings accounts. This scenario would present an interesting route for future study.

Another potential formulation for mixing savings and weather insurance would be to follow the example of "whole life" life insurance policies to develop a type of financial product called "whole weather". With this product, policy holders would purchase a fixed multi-year policy where they would pay premiums each year and receive insurance coverage for each monsoon. If at the end of the term they had not recovered at least the amount of premiums paid in payouts, then they would be refunded the extra premiums. The insurance would be funded by the difference between the nominal premium paid and the expected present value of the future refunds. Based on the success of whole life insurance products, the product would potentially be attractive to customers, as the premium refund makes it seem as if there is no risk of losing money with the product. We hope to explore this type of product in future studies.

Table A3.1: Policy Termsheet

Cover Phase	I	II	III
Duration	35 days	35 days	40 days
DEFICIT			
Strike (mm)	150	80	-
Exit (mm)	50	10	-
Amount Paid per mm (Rs / mm)	2.00	2.00	-
Policy Limit (Rs)	200	200	-
EXCESS			
Strike (mm)	-	-	550
Exit (mm)	-	-	650
Amount Paid per mm (Rs / mm)	-	-	1.00
Policy Limit (Rs / Acre)	-	-	100

3.7 Appendix

3.7.1 Rainfall Index Insurance

This experiment uses rainfall index insurance policies underwritten by the Indian insurance company ICICI-Lombard. Rainfall index insurance creates a contract based on rainfall at a local weather station, and daily rainfall readings from this station are used to calculate the insurance payout. Our policy is written based upon rainfall at a weather station administered by the Indian Meteorological Department located near the airport in Ahmedabad. All of our subjects lived within 30km of this rainfall station, so the rainfall at the station and on their farms should be similar.

The policy provides insurance coverage for three phases of the monsoon, and each phase provides coverage for excess or deficit rainfall. For deficit (excess) policies, payouts are made if the cumulative payout is below (above) a certain threshold. The specifics of the product used in this paper are outlined in Table A3.1. The policy begins when 50mm of rain have accumulated during the month of June, but starts on July 1st if this threshold is not met in June. The first phase lasts 35 days, and offers a payout of Rs 2 for each millimeter of rain below the "strike" of 150mm that accumulates during the phase. Phase 2 has similar conditions, though it also has a lower threshold known as an "exit". When rainfall falls below the exit, the payout for Phase 2 jumps to the policy limit of Rs 200. Phase 3 provides coverage for excess rainfall during the harvest period, starting payouts when rainfall is above 500mm.

The policy offers a maximum payout of Rs 500 per unit, and was priced by ICICI-Lombard at Rs 66. Based on historical data from the Indian Meteorological Department from 1965-2002, the policy would have paid out an average of Rs 22.[16] The offers of pure insurance gave the participants three units of coverage, while the 2/3 insurance + 1/3 savings bundles gave them two units, and the 1/3 insurance + 2/3 savings bundles gave them one unit.

The policies used in this experiment provided coverage for the monsoon season in 2010. The 2010 monsoon around Ahmedabad was one of above-average rains, resulting in good

[16]This is a normal amount of loading for market priced insurance.

crop outcomes for most farmers, and therefore the insurance policy did not give a payout.

3.7.2 Discount Factors

In order to calculate discount factors we asked customers a sequence of questions about whether they would like to receive Rs 80 now or a certain amount of money in the future. These questions were all hypothetical, with no actual money being dispensed. In the first question the future amount of money was Rs 60, and it increased in intervals of Rs 20 to Rs 280. Subjects with consistent time preferences who prefer present consumption to future consumption would be expected to initially prefer Rs 80 now, but at some point would switch to receiving money in the future. This switching point can provide bounds on the discount factor. Sample question wording (translated into English) is as follows:

Would you rather receive Rs 80 today or Rs 120 guaranteed in November?

- Rs 80 today
- Rs 120 in November
- Don't Know

The below table shows the implied discount factors generated by certain switching points as well as the percentage of respondents in each category. When performing regressions using the discount factor, we used the mean of the discount factor bounds where they were well defined. For people whose discount rates were unbounded from below we used a rate of .28, which implies that people would have switched to preferring a future payment of Rs 300. For people who were unbounded from above we used a discount factor of 2, implying that people would be indifferent between Rs 80 now or Rs 40 in the future. People who had multiple switching points do not have a well defined discount factor, and are therefore dropped from regressions requiring discount factors.

		Implied Discount Factor			
Money Offerered Now	Switching Point	Lower Bound	Upper Bound	Used for Regressions	Percentage Of Respondents
Rs 80	Always Prefer Future	1.33	∞	2	13.04
Rs 80	Rs 80	1	1.33	1.165	4.66
Rs 80	Rs 100	0.8	1	0.9	36.96
Rs 80	Rs 120	0.67	0.8	0.735	9.32
Rs 80	Rs 140	0.57	0.67	0.62	3.11
Rs 80	Rs 160	0.5	0.57	0.535	4.04
Rs 80	Rs 180	0.44	0.5	0.47	1.55
Rs 80	Rs 200	0.4	0.44	0.42	1.55
Rs 80	Rs 220	0.36	0.4	0.38	0.062
Rs 80	Rs 240	0.33	0.36	0.345	0.062
Rs 80	Rs 260	0.31	0.33	0.32	0.031
Rs 80	Rs 280	0.29	0.31	0.3	1.55
Rs 80	Always Prefer Present	-∞	0.29	0.28	17.7
Rs 80	Multiple Switches				4.97

Appendix Table 2: Elicitation of Discount Rates

3.7.3 Risk Attitudes

We elicit risk attitudes with gambles played for real money using the exact same question set used in Binswanger (1980). The exercise consists of a list of 8 lotteries where the subject has a 50% chance of gaining each possible outcome. The first lottery offers the subject Rs 50 with probability 1. As you move down the list, the lotteries increase in both expected value and variance. The exact text of the question (translated into English) is as follows:

> In this question you will be presented with a number of possible gambles to take. In each there is a coin flip, and you get a certain amount of money if it lands on heads and a different amount if it lands on tails. Note that this game will be played FOR REAL MONEY, so think carefully! If you choose 'I don't know', you won't play the game and will not have the opportunity to win any extra money.
>
> Please note that the 'coin flip' will be done on the computer, which will randomly show you either Heads or Tails. The flip will be done at the end of the session.
>
> Which of the following gambles would you prefer?

- Rs 50 for Heads, Rs 50 for Tails
- Rs 45 for Heads, Rs 95 for Tails
- Rs 40 for Heads, Rs 120 for Tails
- Rs 35 for Heads, Rs 125 for Tails
- Rs 30 for Heads, Rs 150 for Tails
- Rs 20 for Heads, Rs 160 for Tails
- Rs 10 for Heads, Rs 190 for Tails
- Rs 0 for Heads, Rs 200 for Tails
- I don't know

In order to assist with understanding, the enumerators showed each participant a coin and explained that they would receive the money at the end of the session based on a virtual coin flip. Based on the chosen lottery, we can classify the level or risk aversion of each subject. We adopt the partial risk aversion coefficient as the measure of risk aversion. The risk aversion coefficients and the number of subjects in each group are presented in Table A3.2.

Table A3.2: Elicitation of Risk Aversion

				Approximate Partial Risk Aversion Coefficient			
Gamble	Head Payoff	Tails Payoff	Risk Level	Upper Bound	Lower Bound	Coeff Used for Regressions	Percentage of Respondents
1	50	50	Extreme	∞	7.51	8	17.08
2	45	90	Severe	7.51	1.74	4.625	9.63
3	40	120	Intermediate	1.74	0.812	1.276	9.01
4	35	125	Inefficient			Observation Dropped	7.76
5	30	150	Moderate	0.812	0.316	0.564	12.42
6	20	160	Inefficient			Observation Dropped	5.9
7	10	190	Slight-to-Neutral	0.316	0	0.158	18.32
8	0	200	Neutral-to-Negative	0 to -∞		0	19.25
9	"I don't Know" selected						0.062

3.7.4 BDM Game

Valuations of insurance and guaranteed payments were elicited using a BDM game. In this section participants were asked to give their valuations of four separate objects:

- A "large" insurance policy with maximum sum insured of Rs 1500
- A "medium" insurance policy with maximum sum insured of Rs 1000 plus a guaranteed payment of Rs 60
- A "small" insurance policy with maximum sum insured of Rs 500 plus a guaranteed payment of Rs 120
- A guaranteed payment of Rs 180 after the monsoon

After giving their valuations, one of the games was played at random for real money, vouchers, and insurance policies. Before giving their valuations, subjects are shown the following text on the computer, which is read aloud by the numerator. Translated into English, it is:

> In the next section you will again need to make decisions about how much you would need to be paid in order to give up certain objects. However, in this case the objects will not be physical things. Instead, they will either be rainfall insurance policies, coupons for payment in the future or both. I'll call each gift a "bundle" because it may be multiple things.
>
> For each question you are to consider that you have been given one of these bundles. You then have to state what is the minimum amount of money you would need to give up this bundle. For each bundle you have to state how much money you would need to give up the bundle if the money was paid to you right now, and also if the money was paid to you in November.
>
> After you tell how much money you would need to give up each of these bundles, one of the games will randomly be played for real. One of the bundles will be picked at random and given to you. You will then be randomly offered an offer to buy back the bundle. If this offer price is greater than the minimum price you said you were willing to accept to give up the bundle, you will be paid

this offer price. If the offer price is less than your minimum, you will keep the bundle. The offer price will be somewhere in between Rs 10 and Rs 250.

It is in your best interests to think about each question thoroughly and give the actual minimum price you would accept for each one!

In the question text, participants are randomly shown one of three frames:

1. Bundle Frame - In this frame the savings/insurance mixtures are explained as a separate insurance policy and a voucher for guaranteed money.

2. Insurance Frame - In this frame the savings/insurance mixtures are explained as an insurance policy with a minimum payout.

3. ICICI Bundle Frame- In this frame we mention that the policy could be purchased from the ICICI-Lombard insurance company. The savings/insurance mixtures are described in the same way as the bundle frame.

As an example of the wording in the three frames, here is the text for the bundle of one insurance policy with maximum payout of Rs 1000 and a voucher for Rs 60. Note that this question appears after asking for a valuation for a pure insurance product, so there is a line clarifying the difference between this offer and the last one.

Bundle Frame:

> Consider you have been given a gift of 1 medium rainfall insurance policy and a voucher for Rs 60 that can be redeemed for cash in November. The policy would normally cost Rs 120 and pays out a maximum of Rs 1000 in the event of bad rainfall. The Rs 60 voucher is just a piece of paper that you can exchange for Rs 60 cash in November. This gift might be especially useful if there is a poor harvest. What is the minimum amount of immediate payment you would accept to give up the insurance policy and coupon? Our offer to purchase this bundle from you will be between Rs 10 and Rs 250. You would receive the payout at the end of today's session.
>
> The difference between this and the previous questions is that now the insurance policy you are offered has a maximum payout of Rs 1000 instead of Rs 1500. However, this time you also will get a gift of Rs 60 paid in November regardless of rainfall.

Insurance Frame:

> Consider you have been given a gift of a special "payout guaranteed" rainfall insurance policy. As before, this policy will pay out in the event of poor rainfall, but it will pay out at least Rs 60 regardless of rainfall. This policy would normally cost Rs 180 and will pay out a maximum of Rs 1060 in the event of bad rainfall. What is the minimum amount you would be willing to accept to

give up the insurance policy? Our offer to purchase this policy from you will be between Rs 10 and Rs 250. You would receive the payment at the end of today's session.

The main difference between this and the previous questions is that now the rainfall insurance policy has a maximum payout of Rs 1060 instead of Rs 1500, but it will pay out a minimum or Rs 60 instead of zero.

ICICI Bundle Frame:

Consider you have been given a gift of 1 medium rainfall insurance policy and a voucher for Rs 60 that can be redeemed for cash in November. The policy can be purchased from the ICICI-LOMBARD insurance company for Rs 120 and pays out a maximum of Rs 1000 in the event of bad rainfall. The Rs 60 voucher is just a piece of paper that you can exchange for Rs 60 cash in November. This gift might be especially useful if there is a poor harvest. What is the minimum amount of immediate payment you would accept to give up the insurance policy and coupon? Our offer to purchase this bundle from you will be between Rs 10 and Rs 250. You would receive the payout at the end of today's session.

The difference between this and the previous questions is that now the insurance policy you are offered has a maximum payout of Rs 1000 instead of Rs 1500. However, this time you also will get a gift of Rs 60 paid in November regardless of rainfall.

Conclusion

This thesis has explored a variety of topics related to the rainfall index insurance market in India. Chapter 1 looked at two methodologies to measure the willingness-to-pay (WTP) for index insurance. We found that a standard neoclassical model of insurance demand did a poor job of predicting WTP compared to using a Becker-DeGroot-Marschak (BDM) mechansism to empirically estimate WTP. Specifically, a neoclassical model heavily overestimates WTP compared with actual insurance purchasing at fixed prices or elicitations using BDM. But why is actual demand so much lower than theoretical demand? One plausible theory is that people may be unfamiliar with insurance and untrusting of insurance companies, which are factors not included in our model. If this were the case, then insurance demand should increase as people gain experience with insurance.

I explored this idea in Chapter 2, specifically asking whether insurance customers who received a payout were more likely to purchase insurance again the following year. Studying a panel dataset of insurance purchasers, I found that this is indeed the case, with people who received insurance payouts being 9-22% more likely to purchase insurance the following year. However, the results do not provide evidence for increased learning about insurance and trust in insurance companies, as villages with significant insurance payouts did not have more new purchasers in the following year. Most theories of technology adoption state that a new product is likely to catch on via spillover effects from early adopters. Our data does not support the hypothesis that this is occurring within the rainfall insurance market.

Based on the struggles of rainfall insurance documented above, in Chapter 3 we proposed a new type of financial product that combines rainfall index insurance with savings, which we call a WISA. We developed a simple theoretical model that shows that each individual will have an optimal WISA type, and then attempt to measure this type in the laboratory. Contrary to our theoretical predictions, we find that most of our laboratory participants preferred both pure savings and pure insurance to any combination of the two. While it is not immediately clear what is driving this behavior, the results certainly do not support the development of a WISA product.

In Chapter 1 we showed that demand for rainfall index insurance is lower than theory would predict, and in the following two chapters we unsuccessfully looked for clues on how to make the product more attractive to consumers. Where does this leave the future of the rainfall insurance market in India? One can look at this situation as either an optimist or a pessimist. An optimist would say that the product and marketplace are still young and simply need to evolve. New products could be developed with lower cost and lower basis risk that may be much more attractive to customers. Furthermore, as poor farmers become more familiar with the concept of insurance their demand may rise.

However, I personally would fall into the camp of the pessimists. Rainfall index insurance has been sold commercially in India for six years, and simply does not seem to be attractive for poor farmers. Customers frequently purchase the smallest amount of insurance possible in order to test out the product, and often do not renew, even after receiving a payout. While the index insurance market may have potential to evolve into something more viable, my personal opinion is that governments and donors can find more productive uses for their funds.

Bibliography

Arellano, Manuel and Stephen Bond, "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *The Review of Economic Studies*, 1991, *58* (2), 277–297.

Becker, G. M., M. H. DeGroot, and J. Marschak, "Measuring utility by a single-response sequential method.," *Behav Sci*, July 1964, *9* (3), 226–232.

Bentvelsen, C.L. and V. Branscheid, ""Yield Response to Water"," *FAO Irrigation and Drainage Paper*, 1986, *33.*

Berry, Jim, Greg Fischer, and Raymond Gutieras, "Eliciting and Using Willingness to Pay: Evidence from Field Trials in Northern Ghana," 2011. Technical Report, mimeo.

Binswanger, Hans P., "Attitudes Toward Risk: Theoretical Implications of an Experiment in Rural India," *The Economic Journal*, 1981, *91* (364), pp. 867–890.

Bohm, Peter, Johan Lindén, and Joakin Sonnegård, "Eliciting Reservation Prices: Becker-DeGroot-Marschak Mechanisms vs. Markets," *The Economic Journal*, 1997, *107* (443), pp. 1079–1089.

Borch, Karl, *Economics of Insurance*, Vol. 29 of *Advanced Textbooks in Economics*, Amsterdam/New York/Oxford/Tokyo: North-Holland, 1990.

Bryan, Gharad, "Ambiguity and Insurance," Nov 2010. Technical Report, mimeo.

Carroll, Christopher D. and Andrew A. Samwick, "The nature of precautionary wealth," *Journal of Monetary Economics*, 1997, *40* (1), 41 – 71.

Clarke, Daniel, "A Theory of Rational Insurance Demand," 2011. Technical Report, mimeo.

Cole, Shawn A., Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery, "Barriers to Household Risk Management: Evidence from India," *Harvard Business School Working Paper*, 2010, *09* (116).

Cole, Shawn and Peter Tufano, "BASIX," *Harvard Business School Case Study*, 2007, (207-099).

Connor, Robert A., "More than risk reduction: The investment appeal of insurance," *Journal of Economic Psychology*, 1996, *17* (1), 39 – 54.

deNicola, Francesca, "The Impact of Weather Insurance on Consumption, Investment, and Welfare," April 2011. Technical Report, mimeo.

Fafchamps, Marcel and John Pender, "Precautionary Saving, Credit Constraints, and Irreversible Investment: Theory and Evidence from Semiarid India," *Journal of Business And Economic Statistics*, 1997, *15* (2), pp. 180–194.

Giné, Xavier and Dean Yang, "Insurance, credit, and technology adoption: Field experimental evidencefrom Malawi," *Journal of Development Economics*, 2009, *89* (1), 1 – 11.

_ , Lev Menand, Robert M. Townsend, and James I. Vickery, "Microinsurance: A Case Study of the Indian Rainfall Index Insurance Market," *World Bank Policy Research Working Paper Series*, Oct 2010.

_ , Robert Townsend, and James Vickery, "Patterns of Rainfall Insurance Participation in Rural India," *World Bank Economic Review*, October 2008, *22* (3), 539–566.

Gneezy, Uri, Muriel Niederle, and Aldo Rustichini, "Performance in Competitive Environments: Gender Differences*," *Quarterly Journal of Economics*, 2003, *118* (3), 1049–1074.

Gomes, Francisco J., "Portfolio Choice and Trading Volume with Loss-Averse Investors," *Journal of Business*, March 2005, *78* (2), 675–706.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese, "Earnings Uncertainty and Precautionary Saving," CEPR Discussion Papers 699, C.E.P.R. Discussion Papers June 1992.

Halek, Martin and Joseph G. Eisenhauer, "Demography of Risk Aversion," *The Journal of Risk and Insurance*, 2001, *68* (1), pp. 1–24.

Hansen, Lars Peter and Kenneth J. Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 1983, *91* (2), pp. 249–265.

Hess, Ulrich, "Innovative financial services for India, monsoon indexed lending and insurance for smallholders," *ARD Working Paper*, 2004, *09.*

Hill, Ruth and Miguel Robles, "Flexible Insurance for Heterogeneous Farmers: Results from a Small Scale Pilot in Ethiopia," 2010. Technical Report, mimeo.

Horowitz, John K., "The Becker-DeGroot-Marschak mechanism is not necessarily incentive compatible, even for non-random goods," *Economics Letters*, 2006, *93* (1), 6 – 11.

_ , "What do we know about the performance of the Becker-DeGroot-Marschak mechanism?," May 2006.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 1979, *47* (2), pp. 263–292.

Karlan, Dean, Issac Osei-Akoto, Robert Osie, and Chris Udry, "Examining Underinvestment in Agriculture: Returns to Capital and Insurance in Ghana," November 2010. Presentation given at the 6th Annual Microinsurance Conference in Manila, Philippines.

Karni, Edi and Zvi Safra, ""Preference Reversal" and the Observability of Preferences by Experimental Methods," *Econometrica*, 1987, *55* (3), pp. 675–685.

Köszegi, Botond and Matthew Rabin, "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 2006, *121* (4), 1133–1165.

_ **and** _ , "Reference-Dependent Risk Attitudes," *The American Economic Review*, 2007, *97*, 1047–1073(27).

Kunreuther, Howard, Warren Sanderson, and Rudolf Vetschera, "A behavioral model of the adoption of protective activities," *Journal of Economic Behavior and Organization*, 1985, *6* (1), 1 – 15.

Lusardi, Annamaria, "On the Importance of the Precautionary Saving Motive," *The American Economic Review*, 1998, *88* (2), pp. 449–453.

Mazar, Nina, Botond Köszegi, and Dan Ariely, "Price-Sensitive Preferences," February 2010. Technical Report, mimeo.

Morduch, Jonathan, "Micro-Insurance: The Next Revolution?," in Dilip Mookherjee Abhijit Banerjee, Roland Benabou, ed., *What Have We Learned About Poverty?*, Oxford University Press, 2006.

Platteau, Jean-Phillipe, "Mutual Insurance as an Elusive Concept in Traditional Rural Communities," *Journal of Development Studies*, 1997, *33* (6), 764–796.

Pratt, John W., "Risk Aversion in the Small and in the Large," *Econometrica*, 1964, *32* (1/2), pp. 122–136.

Rosenzweig, MR, "Savings behaviour in low-income countries," *Oxford Review of Economic Policy*, 2001, *17* (1), 40–54.

Schlesinger, Harris, "The Theory of Insurance Demand," in Dani Rodrik and Mark Rosenzweig, eds., *The Handbook of Insurance*, Boston: Kluwer Academic Publishers, 2000, chapter 5.

Schoemaker, Paul J. H. and Howard C. Kunreuther, "An Experimental Study of Insurance Decisions," *The Journal of Risk and Insurance*, 1979, *46* (4), pp. 603–618.

Skees, Jerry, Stephanie Gober, Panos Varangis, Rodney Lester, and Vijay Kalavakonda, "Developing rainfall-based index insurance in Morocco," Policy Research Working Paper Series 2577, The World Bank April 2001.

Slovic, Paul, Baruch Fischhoff, Sarah Lichtenstein, Bernard Corrigan, and Barbara Combs, "Preference for Insuring against Probable Small Losses: Insurance Implications," *The Journal of Risk and Insurance*, 1977, *44* (2), pp. 237–258.

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

Townsend, Robert M., "Risk and Insurance in Village India," *Econometrica*, 1994, *62* (3), pp. 539–591.