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

Shawn Cole, Daniel Stein, and Jeremy Tobacman*

Preliminary Draft September 7, 2011

Abstract

Rainfall index insurance is a theoretically attractive financial product that has achieved only limited adoption. This paper seeks to understand the structure of demand for rainfall index insurance in India. We develop two approaches to estimating households' valuation of rainfall insurance and evaluate them against an experiment in which fixed prices are randomly assigned. The 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 survey data from members of an insurance pilot in Gujarat, India to fit the model and estimate the willingness to pay (WTP) for rainfall insurance coverage. Relative to the choices 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 points 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 learning dynamics may be a promising direction for future analyses.

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

*Harvard Business School (scole@hbs.edu), London School of Economics (d.stein@lse.ac.uk), and Wharton School at University of Pennsylvania (tobacman@wharton.upenn.edu). We would like to thank the many team members who helped makes this study a success, including research assistants Nilesh Fernando, William Talbott, Sangita Vyas, and Carina Wendel. We would also like to thank Greg Fischer for providing input on the BDM implementation and commenting on drafts of the paper.

(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 environment. 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 2 we give an overview of the insurance products and data used in the experiment. In Section 3 we develop our model of insurance demand, and Section 4 presents benchmark tests of its predictions against insurance decisions at fixed prices. In Section 5 we discuss the implementation of BDM in the field, and test the predictions of BDM against insurance decisions at fixed prices. In Section 6 we directly compare WTP estimates from our model to those of BDM. Section 7 assesses reasons for the discrepancies between the model and the empirical measures of WTP. Section 8 concludes, and offers policy prescriptions based on the results.

2 Product and Data Description

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

Table 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%

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.

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 shows the payout structure for the first phase of 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

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

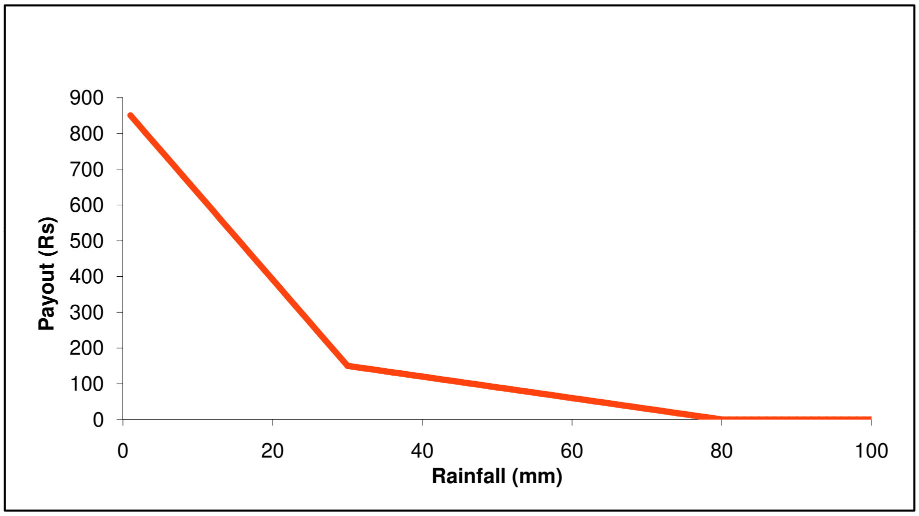

informal secondary market for the insurance policies could technically exist, we have never witnessed any evidence of this.

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 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 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.

Table 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

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 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

Table 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

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 3 outlines the various discounts and games offered.

3 Structural Approach

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.

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.

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 sets the expected utility from purchasing insurance equal to the expected utility from not purchasing insurance, representing the maximum WTP.

$$E[u(Y - P^* - S + M)] = E[u(Y - S)] \tag{1}$$

Timing is as follows:

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

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

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[1] 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:

$$S = \mathbf{1}(R < R_{max})K_y(1 - \frac{R}{R_{max}})Y_n \tag{2}$$

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.[2] 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 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 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

[1] 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.

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

Figure 2: Actual versus Fitted Rainfall over first three phases of the Monsoon. Anand District 1965-2003, fitted to lognormal distribution.

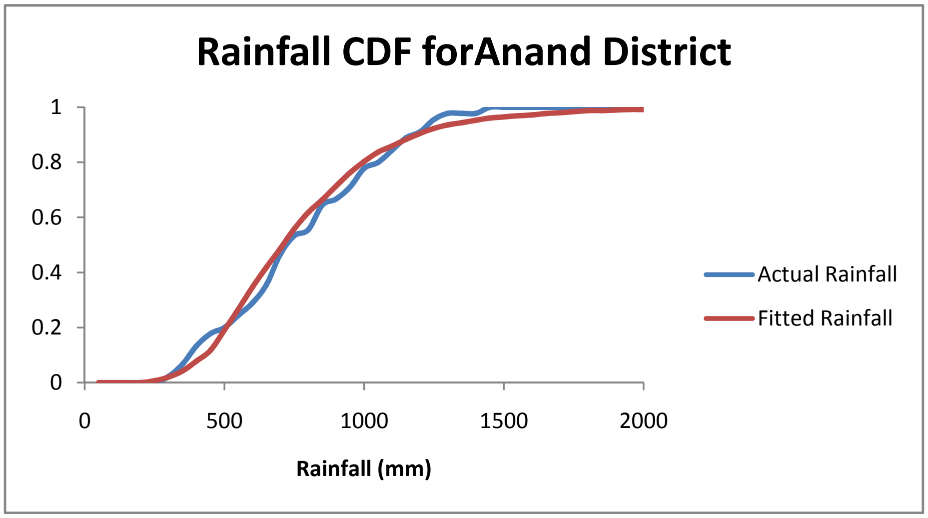

were roughly actuarially fair, we set the payout amounts such that in expectation the payout equals the premium of Rs 150.[3] This corresponds to a payout of Rs 1.11 for each millimeter below the first threshold and Rs 7.77 for each millimeter below the second threshold.

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).

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 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

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

Table 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

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:

$$U(c) = \frac{c^{1-\phi}}{1-\phi} \tag{3}$$

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 is 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 4 outlines the calibration of various constants in the model.

To calculate the willingness to pay for insurance, we numerically solve Equation 1 for P^* for each person who played the BDM game. In the next section we present the results of this model.

4 Test of Model Predictions against Fixed Prices

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.[4] 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 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.[5] The 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 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 differences are driving a wide gap in WTP between the model and observed behavior. We therefore conclude that the model severely overestimates WTP.

[4] 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.

[5] 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 indistinguishable from that of the full sample.

Figure 3: Comparison of Model Predictions to Purchasing at Fixed Prices

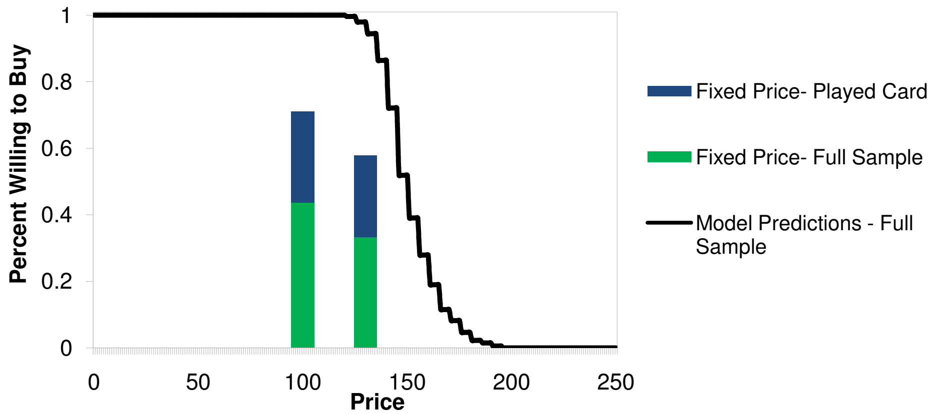

Table 5: Model Predictions Compared to Fixed Prices

	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		

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 A.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.

5 BDM and Fixed Discounts in the Field

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.[6] 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.[7]

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.

[6] 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.

[7] 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.

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 A.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 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.

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 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 4 shows that demand predicted by BDM is close to fixed price demand at a price of RS 100, but is much lower at a fixed price of Rs 130. In Table 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

Figure 4: Comparing BDM Bids to Fixed Price Decisions for Non-Surveyed Population

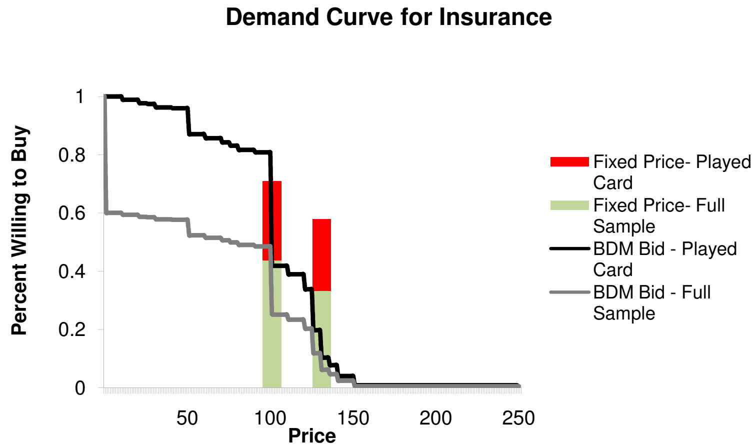

Table 6: Comparison of BDM to Fixed Discounts

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

Table 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

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 7.

The results in Table 7 confirm the comparisons outlined in Table 6. We see that 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 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.

Figure 5: Histogram of BDM bids for 1 Policy

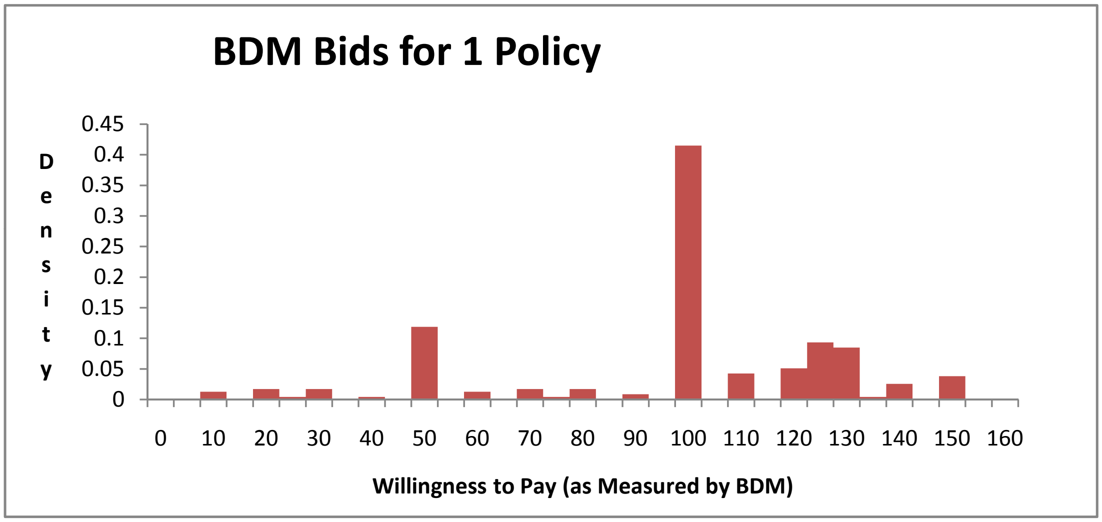

6 Test of BDM against the Theoretical Model

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.

6.1 Demand Curves

In Figure 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 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

Figure 6: Demand Curves for 1 Policy

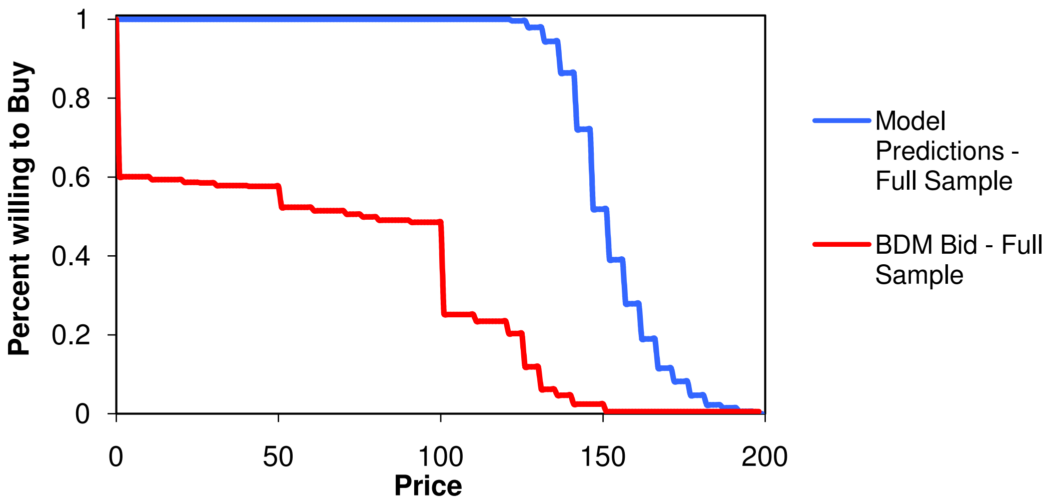

predicts higher WTP at all levels compared to BDM.

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 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 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.

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.

Figure 7: Demand Curves for 4 Policies

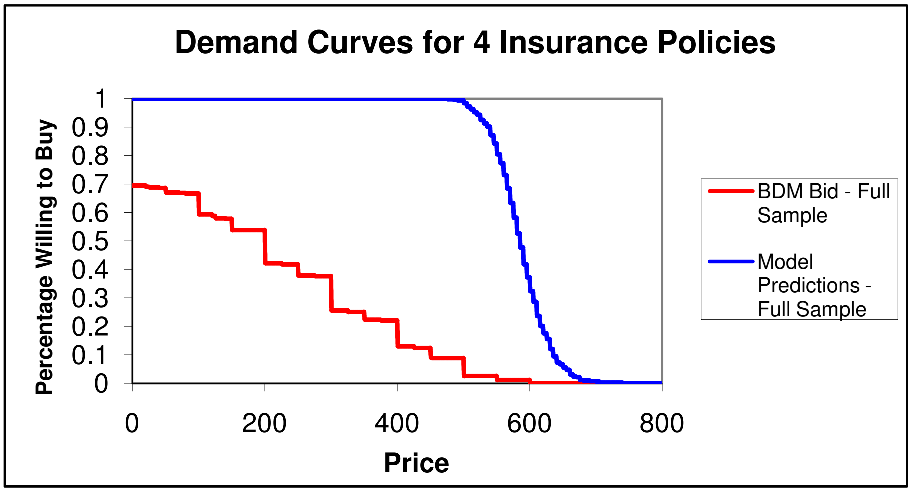

Table 8: Individual Comparison of BDM Bid and Predicted WTP

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

Table 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					

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 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 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 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.

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 10: Correlates of WTP

Dependent Variable:	Univariate BDM bid / Total price	Univariate BDM bid (zeroed)	Multivariate BDM bid / Total price	Multivariate BDM bid / Total price	Multivariate BDM bid (zeroed)	Multivariate 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

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. However, there were large insurance payouts in some areas in 2010. In a future paper, we will look at how recent experiences with insurance can affect insurance demand, and hope this can improve on the static approach developed here.

References

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.

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).

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

Giné, Xavier, 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.

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.

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

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

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.

Table A1.1: Risk Factors

Gamble	Head Payoff	Tails Payoff	Risk Level	Approximate Partial Risk Aversion Coefficient 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

A Appendix

A.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.

A.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. 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

Figure A1.1: Sensitivity to Premium Loading

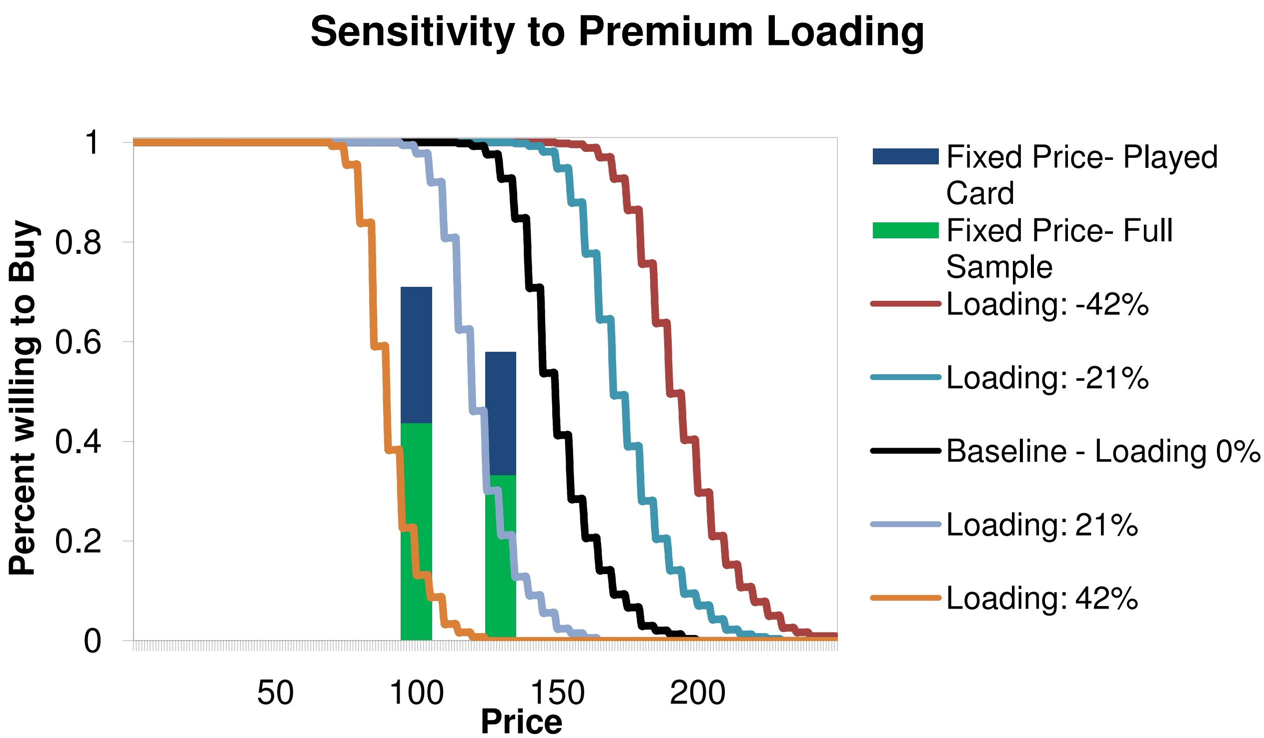

Figure A1.2: Sensitivity to Wealth

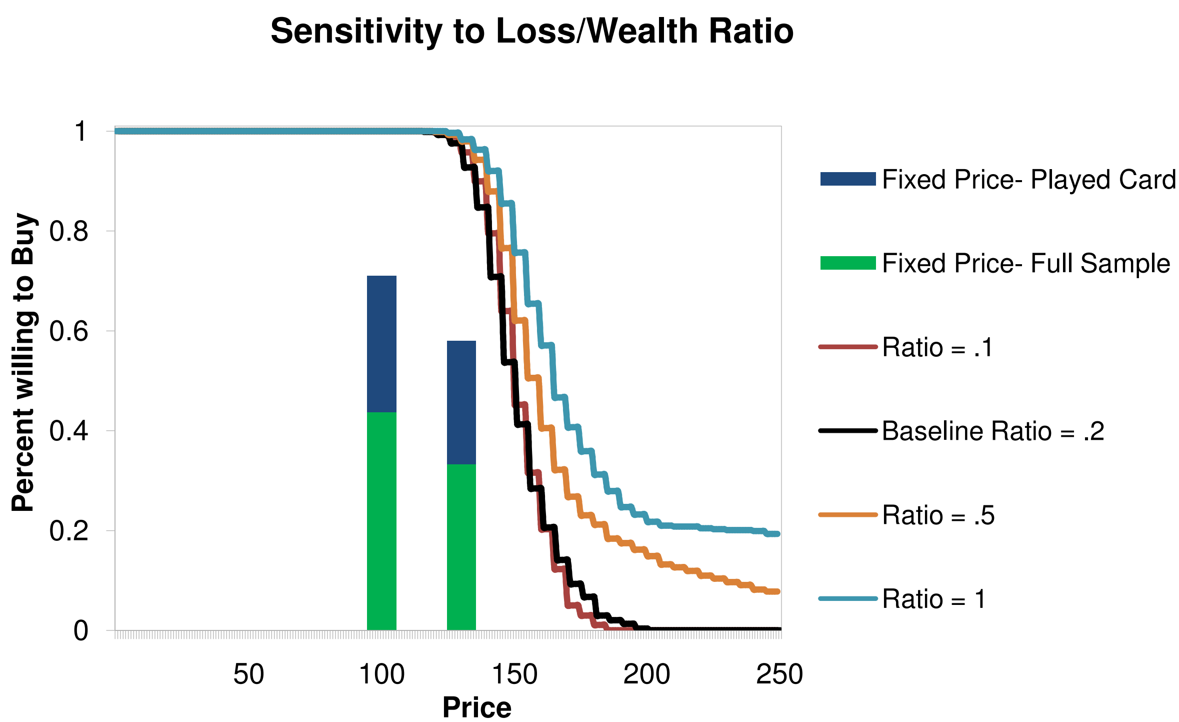

Figure A1.3: Sensitivity to Basis Risk

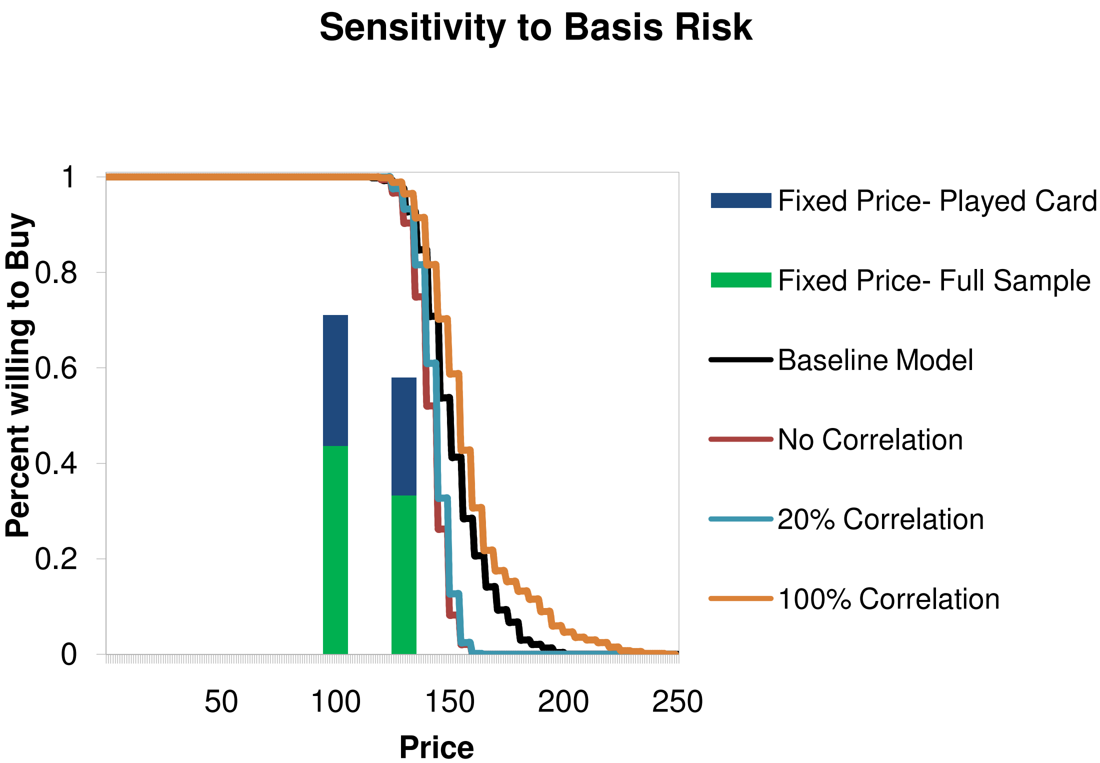

income shocks and payouts. As expected, higher basis risk leads to lower demand, but the effect is not extremely strong. Even at zero correlation, predicted demand is above observed demand.

A.3 BDM Implementation

A.3.1 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.

A.3.2 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

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

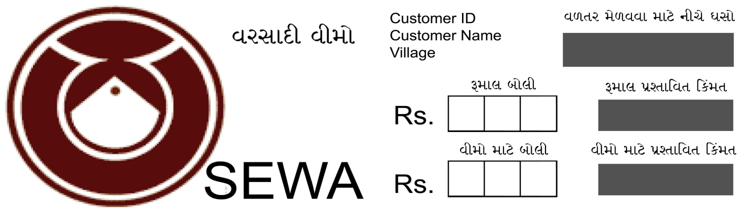

Figure A1.5: Distribution of BDM offers. Left is offers seen in the field. Right is offers generated on all the cards.

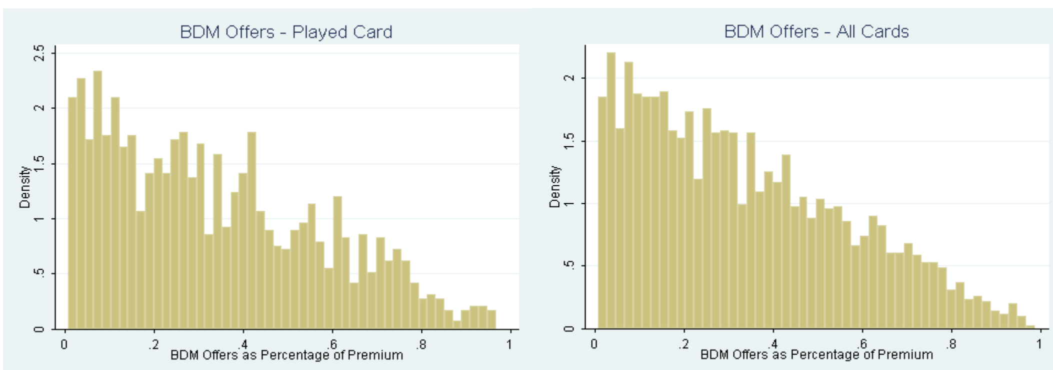

and therefore were allowed to purchase at a discounted price. But when participants "lost" the game, meaning they did not purchase the policy, we worry that the BDM bid may have not been reported. If this was true, then we would expect the cards that were filled out in the field to have the distribution of their offer prices skewed downward.

Figure A1.5 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.

A.3.3 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 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.

Table A1.2: Correlation of BDM Bids and BDM Offers

Dependent Variable is BDM Bid as % of premium	All Districts	Patan	Anand	Ahmedabad
	(1)	(2)	(3)	(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

A.3.4 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.

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			