Replication of Cai et al. (2015)

Social Networks and the Decision to Insure

Emma Clarke, Isabelle Feldhaus, Jinyi Zhu

May 3, 2018

1 Introduction

Individuals make financial decisions based on their own understanding, education, and experience, which can be influenced by their friends' choices and experiences as well. In their paper, "Social Networks and the Decision to Insure", Cai et al. (2015) [1] investigate the role of social networks in farmers' decisions to purchase weather insurance in rural China. They conducted a randomized experiment in which the insurance product is introduced through four sessions in each village, in two rounds three days apart, with one simple session and one intensive session in each round.

Households were randomly assigned to one of these four sessions, and the social network variable is defined as the fraction of one's friends who were invited to a first round intensive session. This randomization procedure generated exogenous variation among second round participants in the proportion of their friends exposed to first round intensive sessions and theoretically establishes grounds for causal interpretations of detected effects. At the village-level, authors also randomized the default option to buy the insurance product in the first round sessions, meaning that if the default was "buy," a farmer need to sign off if he did

not want to purchase insurance. This randomization created exogenous variation in the first round insurance take-up across villages for use as an instrumental variable (IV) for first round purchase decisions.

Using a linear probability model, Cai et al. find that providing intensive information about how insurance works and its expected benefits to a subset of farmers has a large and positive spillover effect on their friends' understanding of insurance benefits and take-up of insurance. Using an instrumental variables (IV) approach, they find that this spillover is driven by the diffusion of knowledge about how insurance works rather than by the influence of friends' purchase decisions.

Replication of study results is an important part of the scientific method to re-investigate and reaffirm key findings or experimental work upon which evidence continues to build. The objective of this project was to replicate the main results from Cai et al. (2015) as well as explore the impact of social networks on weather insurance adoption further through extensions of their analysis. To this end, we reexamine key assumptions of their methods and extend the IV analysis to test heterogeneous effects of the treatment. Extending the original analysis in this way adds to the robustness of study findings and may provide additional insights into the mechanisms of weather insurance take-up, which has particularly low take-up in most countries.

2 Methods

2.1 Replication

In the original article, empirical analysis was conducted using administrative data of insurance purchase from the People's Insurance Company of China and data collected from two surveys: a social network survey and a household survey. The social network survey asked household heads to list five close friends within or outside the village with whom

they most frequently discuss rice production or financial issues [1]. The household survey included questions on demographics, rice production, income, natural disasters experienced and losses incurred, experience in purchasing any kind of insurance, risk attitudes, and perceptions about future disasters.

Analyses using these data were replicated for the primary outcomes of interest described in the Cai et al. (2015) study. Replication was conducted according to their reported methods, which have been briefly summarized below. The social network effect on insurance take-up was determined for the sample of farmers assigned to second round groups, who did not receive first round take-up information by estimating:

$$Takeup_{ij} = \tau_0 + \tau_1 Network_{ij} + \tau_2 X_{ij} + \tau_3 NetSize_{ij} + \eta_j + \epsilon_{ij}, \quad (1)$$

where $Network_{ij}$ is the fraction of friends named by a household in the network survey who were invited to a first round intensive session, X_{ij} is five main household characteristics of household i in village j (i.e., gender, age, literacy of household head, household size, and area of rice production), and $NetSize_{ij}$ is a set of five dummy variables indicating the number of friends listed [1]. To determine whether take-up in social networks affect friends' adoption of insurance, Cai et al. used the randomly assigned default option and type of information session as IVs to determine how overall take-up rate in the first round sessions influence second round participants' behavior.

2.2 Missing Data Imputation

In this study, the treatment and outcome variables were fully observed. However, some of the covariate data, which come from survey questions and administrative data, were incomplete. The original analysis was conducted using a complete case approach, or list-wise deletion, allowing observations to be dropped if any data were missing. If the missing data are missing

completely at random (MCAR), a complete case approach will be unbiased. However, if the missing data are missing at random (MAR) conditional on observed covariates, this approach can generate biased estimates. Even in an experimental setting, it is possible for missing covariate data to induce bias, e.g. if missingness is associated with effect size, effectively leaving us with a measured effect that is not representative of the true effect in the population, as described in Hernan et al. (2004) [4] . Even if there is no bias, complete case analysis reduces the effective sample size and therefore reduces efficiency.

We imputed the missing data in order to evaluate the extent to which this imputation would change the effect size and standard errors. We began by describing and visualizing the missingness. We then implemented an Expectation Maximization missing data algorithm (using the Amelia package in R [5]) to impute missing data. This approach makes an assumption of multivariate normality of the missing variables. We repeated the authors' main analyses using the imputed datasets, combining point estimates and variances according to Rubin's rules [8]. This approach estimates standard errors that reflect both within- and between-imputation uncertainty.

2.3 Reexamining standard errors

The observations in the experiment are grouped at the village level and Cai et al. use robust cluster standard errors (RCSE) to capture the potential correlation between observations within groups. Although RCSE is the most frequently used method to account for clustering, studies have shown that RCSE can be biased downward even with a large number of clusters, and that using a non-parametric bootstrap approach - "bootstrap cluster standard errors" (BCSE) - consistently outperforms RCSE under a variety of conditions as they provide larger and unbiased estimates of the standard errors [2, 3]. Therefore, we reexamined the main results in Cai et al. using the bootstrap method and compared the BCSE we obtained against the RCSE in the original article.

2.4 Evaluating IV assumptions

Cai et al. use an instrumental variables approach to measure the impact of insurance uptake in "Round 1" on insurance uptake (in the same villages) in "Round 2." Their instrument was the randomly assigned default mechanism: in some villages, Round 1 participants were randomized to a default "opt-out" insurance mechanism and, in others, Round 1 participants were randomized to a default "opt-in" insurance mechanism. This instrument generated exogenous variation in Round 1 take-up of insurance, since farmers in villages assigned to "opt-out' mechanism had higher take-up rates. We conducted several analyses to evaluate the assumptions implicit in this IV analysis: relevance of the instrument, the exclusion restriction, monotonicity, and independence.

Relevance is assessed through the first-stage regression. We replicate the first-stage regression from the original study to assess whether this assumption holds.

A number of falsification tests of the instrument were conducted to examine the reduced-form relationship between the default option and insurance uptake in Round 2 [7]. Variables indicating insurance knowledge, loss in yield due to disaster last year, having received payouts from other insurance products before, and trust of the government were selected as placebo outcomes to be tested, as variables that were not expected to be affected by the instrumental variable. In the event that any of these tests indicate unexpected effects, it may suggest that the exclusion restriction does not hold.

The monotonicity assumption cannot fully be tested, but evaluation of the first-stage regression in population subgroups can point to potential violations, e.g. if the first-stage coefficient changes signs in some subgroups.

The independence assumption, which requires that the instrument be independent of potential outcomes, is not a significant concern in this study given that the instrument was randomly assigned.

2.5 Evaluating characteristics of the "compliers"

If the assumptions above hold, the IV analysis estimates the local average treatment effect (LATE). This is the treatment effect among the "compliers" (those who are induced to take up the treatment by the instrument). This analysis does not tell us the effect among "always takers" (those who would take up the treatment regardless of the instrument value to which they are assigned) or among the "never takers" (those who would not take up the treatment, regardless of the instrument value to which they are assigned). While it is impossible to identify which individuals in a dataset are compliers, Kowalski (2016) [6] describe an approach to estimate the size of the complier group and describe their average characteristics. This type of analysis is informative about external validity because it gives more information about the group for whom we are estimating a causal effect, and indicates the extent to which they are similar to the rest of the population.

Following Kowalski, we estimate the size of the complier population by calculating the proportion of the study population with the instrument "off" who are treated (the always takers), and the proportion of the study population with the instrument "on" who are not treated (the never takers). In a setting with a randomized instrument, these proportions can be taken as estimates of the proportion of always takers and never takers in the overall population. By the monotonicity assumption, we assume that the rest of the population is in the complier group.

We then examine average complier characteristics by comparing the strength of the first-stage regression among a particular population sub-group to the strength of the first-stage regression in the overall population. If the first-stage regression is stronger in the sub-group (meaning that the instrument has a larger effect on uptake of treatment in that group), then members of that sub-group are, on average, more likely to be compliers.

This type of analysis requires that the instrument is randomized (as it is in this study), and that the endogenous treatment of interest (in this case, the insurance take-up rate in the

first round) is dichotomous. Therefore, for this analysis we changed the treatment variable (previously “rate of insurance uptake in the first round") into a binary variable for whether or not insurance uptake was more than 50 percent in the first round. After redefining treatment in this way, “always takers” can be defined as those villages that would have majority uptake of insurance in the first round regardless of whether they were in the default opt-in group or the default opt-out group. “Never takers” can be defined as villages that would not have majority uptake of insurance, regardless of their default.

2.6 Evaluating heterogeneous treatment effects

While IV only tells us the estimated treatment effect among the compliers, in an experimental setting it is possible to also estimate bounds of the effects among always takers and never takers. We follow an approach described in Kowalski (2016)[6] for analyzing heterogenous treatment effects in this setting. Firstly, we estimate the average "untreated" outcome among the never takers. This can be observed because we can identify "never takers" as those who have the instrument on and do not take up treatment. Secondly, we estimate the average "treated" outcome among "always takers" by estimating the average outcome among those who have the instrument off and do take up treatment. We can estimate both the untreated and treated outcomes for the complier population, using the estimated 2SLS model to predict outcomes (setting all other covariates at estimated value for their average among the compliers). If we assume that selection into treatment is weakly monotonic, then in some scenarios we can estimate bounds on the treatment effects for the always takers and never takers. The weak monotonicity assumption states that, if we observe that the untreated outcome among never takers is higher than the untreated outcome among compliers, then we expect the untreated outcome among the always takers to be weakly higher than the untreated outcome among the compliers.

3 Results

3.1 Replication

All OLS and IV results relevant to the primary outcomes of interest reported in Cai et al. (2015) were successfully replicated using the procedures outlined in their paper.

3.2 Missing data imputation

We found a low level of missing data in the covariates included in the main regression models implemented by Cai et al. Figure 1 shows the proportion of observations with missing data for all variables with any missingness in the dataset, apart from those that are mechanically expected to be missing because certain questions were only posed to certain experimental groups within the study. There are missing data in several covariates (e.g. level of understanding of insurance) that may be relevant for the outcome. All of the network-related variables (e.g. the number of friends) are missing at the same rate (6 percent), and so only one of these variables is included in the figure. The figure demonstrates that there is no clear pattern to the missingness, as we do not observe monotone missingness. 4,370 observations have no missing variables at all.

After implementing multiple imputation with the Amelia package in R, we find that the results do not change significantly. The table below shows findings from a replication of Table 2 from Cai et al, using multiple imputation instead of complete case analysis. The table only shows the coefficients for the main treatment covariates, though the analysis included the same covariates as the original paper. While some of the point estimates change slightly, the direction and significance of the effects remain largely unchanged after imputation. One exception is the coefficient on the interaction term between "network rate pre-intensive" (a measure of takeup in a study subject's network among those in the "intensive" group) and intensive (an indicator of whether a study subject was in the "intensive" group), which

changes from -0.329 to 0.363 in the imputed data. This change does not influence the conclusions of the paper.

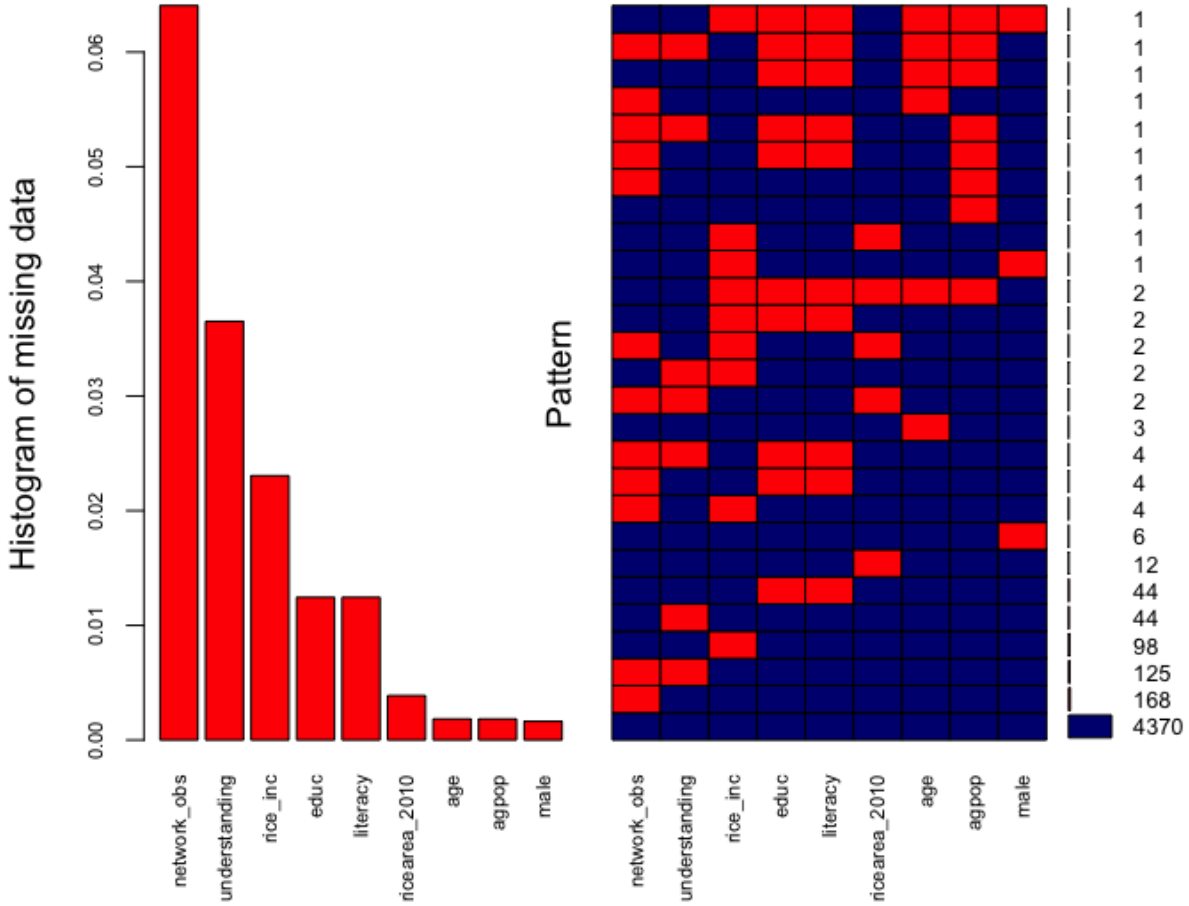

Figure 1: Missing data visualization

Table 1: Replication of main findings with data from multiple imputation

	takeup_survey					
	(1)	(2)	(3)	(4)	(5)	(6)
network_rate_preintensive		0.218	0.185	0.363		
		(0.080)	(0.084)	(0.104)		
network_rate_presimple			−0.099			
			(0.084)			
network_onlyone					0.155	
					(0.105)	
network_onlytwo					0.198	
					(0.090)	
network_twomore					0.137	
					(0.092)	
intensive	0.148	0.027	0.024	0.073	0.073	0.330
	(0.020)	(0.27)	(0.028)	(0.035)	(0.039)	(0.157)
delay						0.045
						(0.031)
risk_averse		0.221		0.216	0.214	0.070
		(0.045)		(0.45)	(0.045)	(0.028)
disaster_prob		0.002		0.002	0.001	0.000
		(0.001)		(0.001)	(0.001)	(0.001)
network_rate_preintensive*intensive				0.363		
				(0.104)		
network_onlyone*intensive					−0.054	
					(0.060)	
network_onlytwo*intensive					−0.058	
					(0.016)	
network_twomore*intensive					−0.145	
					(0.143)	
intensive*delay						−0.073
						(0.042)

3.3 Bootstrap cluster standard errors

We replicated the main results (Table 2) from Cai et al. using a bootstrap method to account for clustering. Table 2 shows the replication results along with the bootstrap cluster standard errors. We found no significant difference in the BCSE as compared to the RCSE in the original article. All noticeable differences are in the third decimal place and none of them affect the statistical significance of the point estimates.

Table 2: Replication of main findings using bootstrap cluster standard errors

	takeup_survey					
	(1)	(2)	(3)	(4)	(5)	(6)
network_rate_preintensive		0.291***	0.278***	0.444***		
		(0.082)	(0.086)	(0.102)		
network_rate_presimple			−0.108			
			(0.086)			
network_onlyone					0.097**	
					(0.041)	
network_onlytwo					0.177	
					(0.115)	
network_twomore					0.137	
					(0.094)	
intensive	0.141***	0.030	0.026	0.081**	0.094**	0.140***
	(0.024)	(0.032)	(0.032)	(0.040)	(0.040)	(0.026)
delay						0.032
						(0.036)
risk_averse		0.109**		0.105**	0.104**	0.072**
		(0.049)		(0.048)	(0.051)	(0.031)
disaster_prob		0.002**		0.002***	0.002**	0.000
		(0.001)		(0.001)	(0.001)	(0.001)
male	0.041	0.016		0.019	0.029	0.048
	(0.045)	(0.067)		(0.069)	(0.076)	(0.042)
age	0.002*	0.005***		0.005***	0.005***	0.002**
	(0.001)	(0.001)		(0.001)	(0.001)	(0.001)
agpop	−0.003	−0.010		−0.010	−0.008	−0.004
	(0.005)	(0.007)		(0.006)	(0.006)	(0.005)
ricearea_2010	0.001	0.004***		0.004***	0.004***	0.001
	(0.001)	(0.001)		(0.001)	(0.001)	(0.001)
literacy	0.083***	0.087***		0.087***	0.090***	0.063***
	(0.026)	(0.032)		(0.030)	(0.032)	(0.023)
network_rate_preintensive*intensive				−0.329**		
				(0.162)		
network_onlyone*intensive					−0.087	
					(0.055)	
network_onlytwo*intensive					−0.091	
					(0.189)	
network_twomore*intensive					−0.141	
					(0.170)	
intensive*delay						−0.053
						(0.047)
Constant	0.363***	−0.310*	0.132	−0.342**	−0.380**	0.333***
	(0.080)	(0.163)	(0.125)	(0.157)	(0.158)	(0.079)
Observations	2,137	1,255	1,274	1,255	1,255	2,756
R^2	0.125	0.119	0.091	0.123	0.129	0.107

Notes: ***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

3.4 Instrumental variables assumptions

The relevance assumption needed for IV analysis was supported based on the results of the first-stage regression, testing whether the default option designation had an effect on take-up of insurance. The default options in the first round sessions yielded significant variations in the first round session take-up rates with the default "buy" sessions resulting in a 12.4 percentage point increase in take-up rate compared to default "not buy" sessions.

The exclusion restriction necessary for IV analysis cannot be conclusively tested and proved to hold. However, the case that the exclusion restriction may be supported by the results of the falsification tests. All variables selected for falsification testing indicated that the instrumental variable used in the analysis, the default option of insurance purchase, did not have unexpected effects and supports the exclusion restriction assumption. Additionally, the randomization procedure for the default option outlined in the authors' study design suggests that the exclusion restriction should hold as all other covariates among the groups would in theory be balanced.

Evaluation of the first-stage regression in population subgroups suggested that there may be monotonicity violations. The first-stage regression coefficient is expected to be weakly positive (meaning that randomization into the 'default' group is expected to weakly increase the probability of insurance uptake in the first round of the experiment). However, we find that, among a sub-group of villages with below average age, the first-stage regression coefficient is negative (but statistically insignificant at the alpha 0.05 level). This indicates a potential violation of the monotonicity assumption.

3.5 Characteristics of compliers

In our analysis of complier characteristics, we collapse the data by village and use village-level insurance take-up in the first-round as the outcome variable of the first-stage regression.

[1]

We first estimated the size of the complier population relative to other groups. We found that, out of the 173 villages in the study, compliers made up 26 percent, always takers made up 29 percent, and never takers made up 44 percent.

Table 3 includes findings from our analysis of the strength of the first-stage regression in different sub-groups of villages. We found that the strength of the first-stage varied according to the farmers' perception of the probability of disaster. In villages with a high average perceived probability of disaster, the first-stage coefficient was 0.38 while, in villages with low averaged perceived probability of disaster, the first-stage coefficient was 0.02. This suggests that farmers in villages with low average perceived disaster probability are more influenced by the default mechanism, perhaps suggesting that this mechanisms operates more in areas where people are less concerned about potential disaster. The first stage was also stronger among villages with above-average mean age than in villages with below-average mean age. In fact, as mentioned above, in villages with average age below the mean, the first-stage coefficient was negative. This is important not only because it violates the monotonicity assumption (as discussed above), but also because it suggests that this instrument (default mechanisms) may not work as expected in all populations.

[1] While Cai et al. do not state this explicitly, there are two interpretations of "compliers" in that are relevant for their IV analysis: an individual-level interpretation and a group-level interpretation. The individual-level interpretation (which is implied by the first-stage regression analysis in the paper) is that the compliers are the individuals attending Round 1 insurance information sessions who are induced to take up insurance because they have a default opt-in. These individuals would not take up insurance if they happened to be in a session with a default opt-out. The group-level interpretation of compliers is that the compliers are the villages (called "natural villages" or "addresses" in the paper) that would have higher insurance take-up if they were assigned to the default opt-in group rather than the default opt-out group. This group-level interpretation of compliers is the one used in the 2SLS regression models used in the paper's main analysis. It would not be possible to run the 2SLS using the individual interpretation of compliers, since the treatment of interest (living in a place with higher uptake in the first round) is defined at the village level.

Table 3: First-stage regression in different sub-groups of villages

	pre_takeup_maj						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
default	0.178**	0.018	0.432***	0.152	0.104	0.388**	−0.091
	(0.085)	(0.153)	(0.164)	(0.153)	(0.164)	(0.153)	(0.164)
portion.male	0.555	1.377**	−0.272	0.244	1.002	−1.123*	1.002
	(0.509)	(0.636)	(1.258)	(0.636)	(1.258)	(0.636)	(1.258)
meanage	0.021**	0.019	0.035	0.054***	0.006	0.030**	0.025
	(0.010)	(0.014)	(0.026)	(0.014)	(0.026)	(0.014)	(0.026)
meanagpop	−0.135**	−0.238**	0.094	−0.189*	0.049	−0.114	−0.028
	(0.055)	(0.104)	(0.137)	(0.104)	(0.137)	(0.104)	(0.137)
meanricearea_2010	0.008	0.021**	0.000	0.020**	0.003	0.020**	0.007
	(0.008)	(0.010)	(0.020)	(0.010)	(0.020)	(0.010)	(0.020)
meanliteracy	0.085	−0.380	1.059	0.937*	−0.715	−0.267	0.096
	(0.296)	(0.549)	(0.831)	(0.549)	(0.831)	(0.549)	(0.831)
meanrisk_averse	−0.759**	−0.957*	0.228	−1.026**	−2.067*	−1.105**	−0.612
	(0.353)	(0.521)	(1.087)	(0.521)	(1.087)	(0.521)	(1.087)
meandisaster_prob	−0.005	0.008	−0.011	0.007	−0.010	−0.032	0.005
	(0.008)	(0.020)	(0.029)	(0.020)	(0.029)	(0.020)	(0.029)
Constant	−0.278	−1.690	−1.982	−3.070*	0.325	1.272	−2.454
	(0.896)	(1.606)	(2.425)	(1.606)	(2.425)	(1.606)	(2.425)
Observations	173	94	79	88	85	87	86
R^2	0.469	0.674	0.635	0.656	0.636	0.671	0.657
Adjusted R^2	0.225	0.277	0.138	0.233	0.195	0.326	0.213
Residual Std. Error	0.437	0.416	0.467	0.431	0.450	0.412	0.436
F Statistic	1.927***	1.699**	1.278	1.551*	1.443	1.946**	1.479

Notes:

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.
Column 1 shows the first-stage in all of the study villages.
Column 2 shows the first-stage in villages with above-average mean disaster probability.
Column 3 shows the first-stage in villages with below-average mean disaster probability.
Column 4 shows the first-stage in villages with above-average mean disaster risk aversion.
Column 5 shows the first-stage in villages with below-average mean disaster risk aversion.
Column 6 shows the first-stage in villages with above-average mean age .
Column 7 shows the first-stage in villages with below-average mean age.

3.6 Heterogeneous treatment effects

The analysis above demonstrates that compliers are different from the overall population in observed covariates. We can also evaluate whether the treatment effect(the effect of Round 1 uptake on Round 2 insurance uptake) among compliers differs from the average treatment effect in the population. If the treatment effect is different, it suggests that social networks have different levels of influence in different villages.

The table below shows the treated and untreated outcomes for compliers, the untreated outcome for never takers, and the treated outcome for always takers. In this table, the "outcome" is the insurance uptake in the second round of the experiment, and the "treatment" is a binary indicator for whether or not a majority of farmers took up insurance in the first round of the experiment. If we assume monotonic selection into treatment, we would expect the untreated outcome in the "always takers" group to be lower than that in the complier group. Even if the untreated outcome in the "always takers" group is 0 (meaning that expected takeup in Round 2 would be zero if a minority of villages took up insurance in Round 1), the treatment effect among always takers is less than the treatment effect among compliers. This suggests that the villages in which farmers are responsive to defaults (the compliers) may also have stronger social network effects. We cannot draw conclusions about the effects among never takers because this analysis does not provide clear bounds for that group.

Treatment	Never takers	Compliers	Always takers
1 (Majority take-up in R1)	Unknown	0.85	.56
0 (Minority take-up in R1)	0.41	0.23	Unknown

4 Discussion

We replicated Cai et al. (2015) following the methodology described in the paper. The study findings did not change significantly when we incorporated bootstrapped standard errors or when we imputed missing data through multiple imputation.

We used several different approaches to evaluate the assumptions underlying the study's instrumental variable analysis. We used placebo tests to evaluate the exclusion restriction, and found that the IV results were robust to these tests. Our estimation of the first-stage regression in population sub-groups suggest that the monotonicity assumption may not hold, as the first stage coefficient was negative for one sub-group. However, the coefficient was not statistically significant so it is possible that monotonicity still holds.

We examined complier characteristics by comparing the strength of the first-stage regression in the overall population to the population sub-groups. We found that, compared with the overall study population, compliers are on average older, less risk averse, and less concerned about the probability of disaster. This is not surprising, as we would expect that risk averse populations that are concerned about a high probability of disaster are likely to take up insurance regardless of the default mechanism. We also find that "always takers" have a smaller treatment effect than compliers. This is also unsurprising: in villages where a majority of Round 1 farmers tend to take up treatment regardless of the default, we would not expect social influence to play an important role in influencing takeup among Round 2 farmers.

Overall, replication of this study highlights the quality of the study design with respect to the applied analyses as well as the robustness of findings of the Cai et al. (2015) paper. Our findings support their conclusion that providing intensive information about how insurance works and the expected benefits of the product to a subset of farmers has a positive spillover effect on others driven by diffusion of relevant knowledge.

References

[1] J. Cai, A. De Janvry, and E. Sadoulet. Social Networks and the Decision to Insure. *American Economic Journal: Applied Economics*, 7(2):81–108, Apr. 2015.

[2] A. C. Cameron, J. B. Gelbach, and D. L. Miller. Bootstrap-Based Improvements for Inference with Clustered Errors. *The Review of Economics and Statistics*, 90(3):414–427, July 2008.

[3] J. J. Harden. A Bootstrap Method for Conducting Statistical Inference with Clustered Data. *State Politics & Policy Quarterly*, 11(2):223–246, June 2011.

[4] M. A. Hernan, S. Hernandez-Diaz, and J. M. Robins. A Structural Approach to Selection Bias:. *Epidemiology*, 15(5):615–625, Sept. 2004.

[5] J. Honaker, G. King, and M. Blackwell. Amelia II: A program for missing data. *Journal of Statistical Software*, 45(7):1–47, 2011.

[6] A. E. Kowalski. Doing More When You're Running LATE: Applying Marginal Treatment Effect Methods to Examine Treatment Effect Heterogeneity in Experiments. Working Paper 22363, National Bureau of Economic Research, June 2016.

[7] N. Nunn and L. Wantchekon. The slave trade and the origins of mistrust in africa. *American Economic Review*, 101(7):3221–52, 2011.

[8] D. B. Rubin. *Multiple imputation for nonresponse in surveys*. Wiley series in probability and mathematical statistics. Applied probability and statistics. Wiley, New York ;, 1987. HOLLIS number: 001654131.