Social Learning in Experimental Games: Evidence from Rwanda

Alexander Coutts[*,†]
Nova School of Business and Economics

March 22, 2018

Preliminary Draft: Please do *not* circulate without author's permission.

Abstract

The use of lab experiments outside of traditional experimental economics laboratories, or "lab in the field" experiments, is rapidly increasing, given the recognition among researchers of their value in studying the preferences and behavior of theoretically relevant populations. While purporting to maintain the paradigm of control that a standard laboratory experiment offers, lab in the field experiments are often conducted in natural settings where researchers inherently must relinquish some of this control. In this paper I present evidence that extemporaneous communication about a public goods game across rural villages in Rwanda led to unanticipated increases in cooperative behavior. To recover causal estimates of the effect of communication, I utilize a matching strategy which matches similar villages on all observables available at planning stages, and compares villages that had opportunities to communicate with past participants, with villages that had no such opportunities. I conclude that exposure to previous participants increases overall cooperativeness by 8-14%. As these results are theoretically unanticipated, they suggest caution regarding the design and interpretation of lab in the field studies.

JEL classification: C92, C93, D83, H41, O10, O12.

Keywords: Experimental Games, Public Goods Games, Field Experiments, Development, Rwanda, East Africa, Spillovers, Information Transmission, Social Learning, Learning.

[*]Nova School of Business and Economics, Faculdade de Economia da Universidade Nova de Lisboa, Campus de Campolide, 1099-032 Lisbon, Portugal; alexander.coutts@novasbe.edu

[†]Acknowledgements: I would like to thank Hunt Allcott and David Cesarini for providing constant guidance and expert advice for this paper. I am additionally grateful for helpful comments from Isaac Baley, Canh Thien Dang, Deborah Goldschmidt, Nicole Hildebrandt, Elliot Lipnowski, Molly Lipscomb, David Low, Joseph Mullins, Yaw Nyarko, Giorgia Romagnoli, Andrew Schotter, Emilia Soldani, Tobias Salz, Pedro Vicente, Christopher Woolnough, and seminar participants at New York University. I am indebted to Zachary Clemence and Kris Cox at Innovations for Poverty Action for their help with management and implementation, and to Lea Stuff for assistance running the experiments. All errors are my own.

1 Introduction

Laboratory experiments to measure preferences and behavior have become a standard component of the economist's toolkit, and now a significant number of experiments are conducted in "field contexts", which I refer to as those outside of a university classroom or computer lab. In developing countries, the utilization of such lab in the field experiments has increased dramatically in recent years, as researchers recognize the critical importance of understanding relationships between preferences and development outcomes.[1] Such work is increasingly finding outlets in leading journals of economics and political science.[2]

The internal validity of such studies and their conclusions hinges on unbiased identification of preferences, which among other things, requires thoughtful experimental design. There is significant awareness of the effects of individual learning both within and across games, however there has been significantly less study of social learning in these contexts.[3] Failure to account for such learning in experiments can bias the identification of preferences.

To a certain extent, the fact that individuals may communicate socially and learn from one another has always been possible in standard lab experiments. A typical participant in an economics lab experiment is often a student, and they may have friends or colleagues who previously participated in a particular experiment. Thus, it is conceivable that these students may discuss outcomes or strategies of a particular experiment with friends who are about to participate.

However, I argue that many experimental labs at universities, where most lab experiments are conducted, in fact should be less susceptible to this issue that corresponding field settings, such as across rural villages in Rwanda. First, many experimental labs have multiple experiments occurring in any given week, and often subjects are not aware which experiment they will be participating in. Second, participating in economics experiments typically is not considered a highly noteworthy event in a subject's semester of study.[4] Combined, these reasons hint that communication may be a bigger issue in field contexts.

[1] See Gneezy and Imas (2017) for a discussion of the definition of "lab in the field".

[2] Recent examples include Jakiela and Ozier (2015), Kosfeld and Rustagi (2015), and Avdeenko and Gilligan (2015).

[3] One can many discussions of learning from repeated play in various games, a good starting point is Kagel and Roth's *Handbook of Experimental Economics*. See Bednar et al. (2012) for learning across games.

[4] For researchers who have conducted experiments in US or European university contexts versus rural Sub-Saharan African contexts, it should be clear that the latter are often perceived as significant (rather strange) events that are not typical. Further, the stakes tend to be substantially higher. In the case of the current experiment average earnings were greater than a typical full day of earnings.

In this paper, I study one such large implementation of a lab in the field experiment in Rwanda involving the participation of 150 rural villages and implemented over a period of three months. Contrary to both expectations and theoretical predictions, I find significant evidence of social learning in public goods experiments. Using GPS data I am able to characterize how cooperation in the public goods games increased in villages that were located nearby other villages which previously participated.

The key difficulty in the analysis is identifying whether differences in cooperation in villages located near past participants are attributable to selection on existing characteristics that are correlated with the order that villages were visited. Because the order of village visits was not randomized, this presents a clear problem for identification. To circumvent this problem, I use a matching strategy whereby I match villages using all village characteristics available to the planner who chose the order of village visits. The planner's

Critically, the planner had access to very few characteristics about villages, with the most important being distance to the home base. Thus after matching similar villages, one observes significant variation in the number of neighboring villages who previously participated. This variation is precisely the treatment of interest - villages that appear ex-ante identical to the planner, but for idiosyncratic reasons some were "treated", i.e. they had neighboring villages which previously participated, and others were not.

In certain instances it was readily apparent that communication had occurred between past and future participants. In one anecdote, a few weeks into the study, the survey team visited a village that appeared similar to others in the region. Standard protocol was followed, however in this specific village, all (12) participants contributed the maximum possible amount in the public goods game. Because this was so exceptional, the team stayed behind to ask the villagers what led to such high levels of cooperation. A woman explained that she was friends with some of the women in a neighboring village, and one of her friends had participated in the same game only two days prior. Her friend told her that she should contribute the maximum amount, and she had shared this information amongst these villagers before the team had arrived.

Anecdotes like the one above suggest that social learning may impact behavior in important ways. This phenomenon can be tested more rigorously in the data using the methodology discussed above. I find that communication increased contribution rates in the entire sample by 8-14%. The bulk of this paper is dedicated to estimating the effects of social learning on behavior and understanding the channel through which communication facilitated this effect.

To my knowledge this is the first study to find evidence for and document unstructured social learning across sessions in experimental games - either in the lab or field. This finding has important implications for the design of experiments as well as the broader interpretations of their results. The finding is especially significant, since these considerations exist in a context where standard theory makes the clear prediction that social information should not change behavior in public goods games. Randomized control trials that recognize the importance of spillovers through social learning may nonetheless fail to account for such spillovers when ex-ante theoretical predictions rule it out. If researchers fail to account for social learning this may weaken the external validity of the study.[5]

While standard theoretical predictions suggest such social learning should not lead to behavioral change, social learning has been shown to alter behavior in controlled lab settings. For example, Chaudhuri et al. (2006) find that providing future participants with advice from previous participants in a public goods experiment increases average contributions.[6] Similarly pre-game communication has been found to change behavior in the lab, for example, Isaac and Walker (1988a) find that such communication leads to higher contributions in public goods games.[7]

It is thus surprising that with lab and lab in the field experiments researchers have been less concerned with social learning or spillovers. While researchers conducting field

[5]Consider an example of a randomized controlled trial, where a new product (e.g. bed net) is being sold at a discount to households. Suppose individuals in the study area happen to communicate extensively. Those that have purchased these nets may communicate with others, and this may alter these individuals' willingness to purchase. For example they might give advice, "I slept under this net, and it was (un)comfortable. I think you should (not) purchase it." The external validity of such a study is compromised, unless information environments are similar. If part of the success (or failure) of a treatment reflects the role of advice, then in a different setting where communication is less prevalent, these effects may be substantially different. For example, Dupas (2009) conducts a field experiment with the purchase and usage of insecticide treated nets (ITNs) in Kenya. Randomly selected households were visited and given a voucher to purchase an ITN, over a period from April to October 2007. One possibility is that households visited early in the project gave advice to friends or family that could be visited later. In fact, in the data from Dupas (2009), being visited one month later in the study increases probability of purchase by approximately 5% (author's calculations). In a follow-up study Dupas (2014) does in fact find that being in an area where there was a high density of treated households had a positive effect on the likelihood of purchasing an ITN.

[6]They find that when when advice is common knowledge to future participants it increases average contributions of these participants. Other authors have also looked at the effects of advice for a range of experimental games. Relevant examples include Chaudhuri et al. (2009), Celen et al. (2010), and Ding and Schotter (2015).

[7]There is a sizeable literature on the effects of pre-game communication or cheap-talk in experiments. Sally (1995) conducts a large meta-analysis of experiments with face to face communication. Crawford (1998) additionally surveys a number of experiments with cheap talk. More recently Bochet and Putterman (2009) and Brosig et al. (2003) have examined communication in public goods games.

experiments in development are attuned to anticipating and accounting for spillovers, there has been less diligence for lab experiments. The results of this paper suggest that this needs to change.

To summarize the remainder of this paper, the next section outlines details of the public goods games and the data. This is followed by an empirical analysis that demonstrates the effects of social learning on behavior. I next spend time examining potential alternative explanations for the findings, followed by a concluding discussion.

2 Data and Design

The experimental games were conducted as a component of a broader evaluation of community health programs. In parallel with household surveys for that evaluation, public goods experiments were conducted in 150 villages in the Rusizi district in Rwanda. These villages were chosen randomly from a total of 598 villages in the district. For the purposes of this paper, the evaluation of the community health programs is not relevant.

The experimental games were conducted over a three month period, from May to July 2013. All 150 villages that were part of the larger evaluation participated in these games. 12 individuals were randomly selected from the household survey list, and given a ticket to participate in the games the following working day. At the time of the games, the 12 individuals were checked-in by the survey team and completed a brief questionnaire.[8] Local survey staff then explained the game in the local language of Kinyarwanda. A significant amount of time was spent explaining the game, providing a demonstration, and conducting a full practice session. It was important that individual decisions were completely private and anonymous; at no time were individual contributions revealed, a fact emphasized to participants.

The experimental design followed a standard public goods game format. Individuals were given an endowment of 4 x 100 RWF coins.[9] They were given real money to ensure that the stakes were salient, and to minimize confusion. One by one participants were instructed to leave the room, go to a completely private area, and decide how much to

[8] 8 individuals were randomly selected for a wait list, in case individuals did not show up at the specified time.

[9] At the time of the study 400 RWF was approximately 0.60 USD. From the Integrated Household Living Conditions Survey 2010-2011, 400 RWF comprises of more than an average day's income for 45% of the district population. Earnings in the experiment were larger than 1800 RWF on average, which greatly exceeds a day's income for the majority of the population.

contribute to the common fund by depositing this in a small change purse, henceforth referred to as the contribution purse. The remainder of their endowment was kept on their person. This had the added advantage of making it clear for individuals that the money they kept was theirs. Since all individuals carry money on their person, there was no reason for concern about having decisions accidentally revealed.

After allocating their money, the participant would then place the contribution purse in a designated location. After all 12 participants had made their decisions, each individual amount was recorded, using anonymous ID numbers located inside the contribution purse, to prevent identification of individuals by the survey team. After recording, all the purses were emptied publicly, and in a transparent manner the coins were counted, tripled, and divided equally among all 12 participants.[10]

Individual payoffs v_i were thus

$$\begin{aligned} v_i &= 400 - c_i + \frac{3}{12} \cdot \sum_{j=1}^{12} c_j \\ &= 400 - \frac{3}{4} c_i + \frac{1}{4} \cdot \sum_{j \neq i} c_j \end{aligned} \tag{1}$$

Notice that contributing $c_i = 0$ is the unique optimal strategy, with the Marginal Per Capita Return (MPCR) to cooperation set at 0.25. A long history of public goods experiments have shown that individuals tend to contribute non-negligible positive amounts in these type of games, despite the dominant incentives to free-ride. Individuals were given complete information on how the experimental payoffs were calculated, and were provided with a number of examples so that they understood the tradeoff between contributing to oneself or to the group. In addition, a full practice round was conducted to further ensure that everyone understood.

Subjects played two rounds of the public goods game with real stakes, receiving income directly after each round. The second round consisted of one of four different versions of the game.[11] Subjects were aware that there would be a second round, but were not given any information about the specific variation that would be used, ensuring comparability

[10] When necessary, the amount was rounded to the nearest 50 RWF.

[11] The four versions included one game with the ability to punish, one with the ability to reward, and a unique game meant to measure uncertainty in public goods investment. Subjects always knew the outcome of the first round before playing.

across villages. One of the four treatments was the baseline (first round) game repeated. Because of the differences in treatments for the second round, in the primary analysis of social learning I use only the first round of data. However, the second round of data is useful for classifying different behavioral types, and is utilized in a secondary analysis.

3 Results

3.1 Overview of the Games

Table 1 presents summary statistics of individual level variables. The sample of participants was randomly selected from the population in the larger survey.[12] Women were over-represented in the larger survey, and subsequently are 75% of the participants in the experimental games.[13] The average years of completed education is 4.46, which corresponds to partially completing primary school.[14] No participant was under the age of 18, while the average age was 35.

Subjects were asked whether they currently had a mobile number they could be reached at.[15] Additionally, subjects were asked how many of the other 12 participants they knew, the average number known was approximately 2.5.[16]

Of the 150 villages visited, in only 2 villages were we unable to find the full 12 participants. These villages have been dropped from the analysis. The remaining sample is of 1728 individuals that were matched from the broader baseline survey, from 147 villages.[17]

[12]This population sample was a person from *every* household in a selected village with at least one child under the age of 5.

[13]Women were over-represented in this larger survey because the respondent was not randomly selected from the household. Of note is that the gender ratio in Rusizi district is highly skewed, so the over-representation is not as large as it appears. According to the 2013 Rusizi District Gender Statistics Report, the percentage of females aged 25-64 is 55.4%. This age group represents 89% of those participating in this study.

[14]In the region, depending on a subject's age, primary school requires either 6 or 8 years in the majority of cases.

[15]The actual question did not explicitly require the subject to own the mobile phone. Hence it is possible that some subjects used numbers of friends or family members.

[16]We defined knowing someone as knowing both their surname and given name. Piloting suggested this was a good means of eliminating superficial relationships.

[17]48 individuals were not able to be matched, which includes one village. Additionally, in some of the analysis, missing variable entries accounts for slightly lower sample sizes.

Table 1: Summary of Individual Level Variables

	Mean	Std. Dev.	Min.	Max.
Participant is Female	0.74	0.44	0.0	1.0
Years of Education	4.43	3.20	0.0	18.0
Age	35.42	10.12	16.0	90.0
Has Mobile Phone	0.33	0.47	0.0	1.0
Number of Participants Known	2.56	2.03	0.0	11.0
Observations	1711			

Figure 1 presents the distribution of contributions in the public goods games. The possible levels ranged from 0 to 400 RWF, in 100 RWF increments. The average level contributed to the group fund is slightly over 255 RWF, which is about 64% of the socially optimal level of contributing the maximum 400 RWF. The modal contribution was the maximum, 400 RWF, with approximately 36% of subjects contributing this amount. As is typical in public goods experiments, the standard theoretical prediction of contributing zero is rejected.

Figure 1: Distribution of Contributions (RWF)

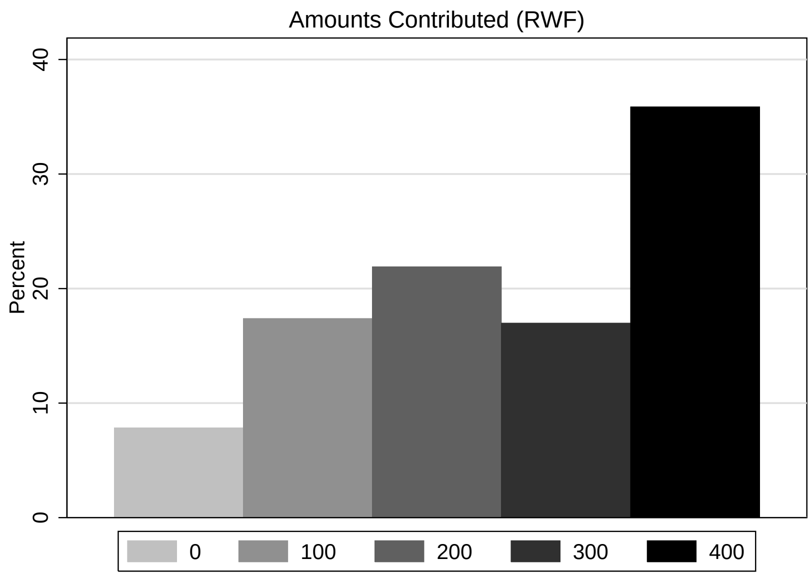

Contributions in the public goods experiments. Possible values ranged from 0 to 400 RWF, in 100 RWF increments. N = 1748.

The 64% contribution rate is on the higher end of contributions in experimental public goods games. Typically, contribution rates range between 40-60%, though with a range of different MPCR. Another difference between these results and previous experiments is that the proportion of "free-riders" or those contributing nothing, is lower than commonly found.[18]

3.2 Identification Strategy

The identification of the effects of social learning and communication on cooperation involves a matching strategy, which pairs otherwise similar villages which only differed on whether there existed opportunities for such social learning to occur.

In identifying the effects of social learning on contribution levels, the primary threat to identification is that the order of visits was not explicitly randomized. In practice, the

[18]See Ledyard (1995) and Chaudhuri (2011). A MPCR of 0.25 is on the lower end, suggesting that the observed contribution rates are indeed quite high, see Isaac and Walker (1988b).

study planner observed some available characteristics of the 150 villages that were in the study, and had to determine an ordering. While not random, there was also no explicit strategy conditioned on these observables. The primary concern of the planner was logistical convenience, as well as ensuring "difficult" villages were spread evenly throughout the study. Difficult villages were those that were located far from the study's base location, and/or those that had large numbers of households. This in fact helps with the identification strategy, since it helps to re-balance these characteristics among villages visited earlier and later, which will correlate with opportunities for potential social learning.

The key strategy is to exploit the exogenous variation in the planner's decision making, to find otherwise identical looking villages (based on all observables available to the planner), but by chance some had neighbors who previously participated in the public goods games, while others had no such neighbors. Those with previously participating neighbors thus had potential opportunities to communicate with past participants, while those with no previously participating neighbors had no opportunities.

For this estimation, it is most useful to consider an exercise where one frames this in terms of the treatment effects literature. In particular, certain villages are "treated" with exposure to previous participants of the public goods games, while others are not.

In identifying the effects of information on contribution levels, the primary threat to identification is that the order of visits was not explicitly randomized. However this setting is particularly amenable to propensity score matching techniques, see Rosenbaum and Rubin (1983), in order to recover causal effects of information on behavior. The reason is that the order of visit could only be conditioned on observables known to the planner, *before* the games were conducted. Thus one can make use of this full set of observables to generate propensity scores for the "treatment": having neighbors who previously participated. By matching treatment and control villages with similar values of the propensity score, i.e. villages who are similar based on all observables available to the planner, one is able to recover causal estimates of treatment effects.

3.3 Treatment

Treatment in this study is defined as whether a village had any neighboring villages that previously participated in the public goods games. Two villages are defined as neighbors if they are within two kilometers of one another. Critically, the treatment of interest for this study is whether a village had opportunities for communication with past participating village neighbors. The treatment is not the effect of communication directly, since this is not observed. The distance of two kilometers was selected as the first choice for a distance which most reasonably captured the possibility of communication between villages. However, since this research question (and distance) was *not* identified prior to the study, it is important to demonstrate that the results are consistent for other distances, and that two kilometers was not chosen after an ex-post comparison of different distances. To demonstrate this, Appendix Section 5.1 shows robustness checks which verify that the results are consistent for other distances.[19]

There is a further dimension through which treatment definition may be interpreted differently, and that is the threshold for the number of neighboring villages. The primary concern is in identifying reasonable opportunities for cooperation. Using one neighbor is the intuitive starting point. Moreover, it is important to note that choosing the threshold incorrectly should only work to bias the estimates downwards *against* finding effects. Choosing too low a threshold runs the risk of downward bias due to the inclusion of villages which didn't have opportunities to communicate in the treatment group. On the other hand, choosing too high a threshold runs the risk of downward bias due to control villages having opportunities to communicate.

According to the classification of treatment of having at least one neighboring villages that previously participated in the games, 58 out of 147 or 39% of villages are not-treated, while the remaining 89 (61%) of villages are treated. Figure 2 shows the number of villages in control and treatment groups, broken down into exactly how many past participating neighboring villages they had. Of note is that the vast majority (91%) of villages have 2 or fewer neighbours that participated. The maximum number of neighbours was 7, this occurred for only one village in the study.

[19]There is also an element of timing. For example, one could further restrict the definition of treatment to only apply to villages that had neighboring participants in the last week or month. It is not clear what the relationship should be between information and time. On one hand, more time might allow information to be disseminated across villages. On the other hand, more time could allow information to deteriorate. As the relationship may be non-monotonic, I try to remain agnostic by not conditioning on time.

Figure 2: Defining the Treatment

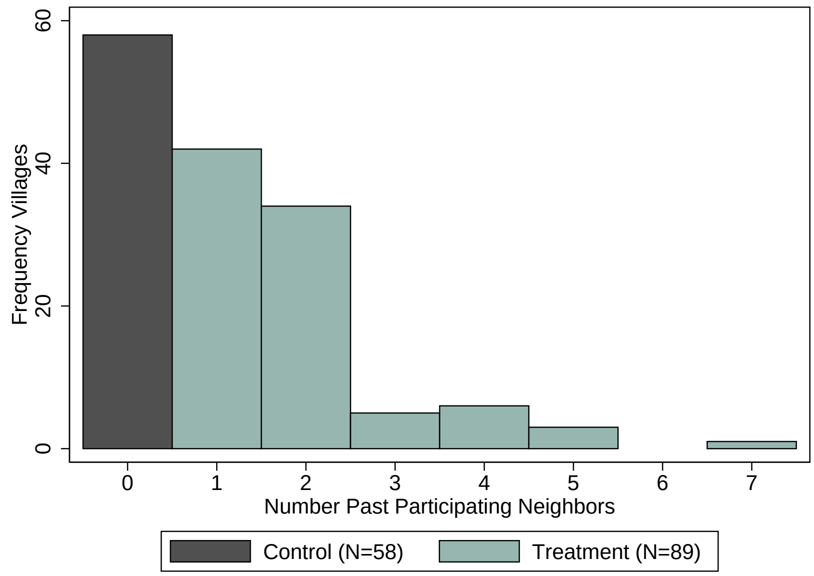

Figure shows the frequency of observing villages in the sample with given number of past participating villages within two kilometers as neighbors.

3.4 Determinants of Treatment and Balance

Table 2 examines a logit regression village characteristics on the treatment indicator for potential communication, i.e. 1 if the village had at least one neighboring village within 2km who previously participated in the public goods games, and 0 otherwise.

From Table 2 it is possible to see that among the variables that the planner potentially had access to, only the number of neighboring villages within 2km is significant. This is not surprising, since villages with more neighbors that participated in the sample (at any point in time) are mechanically more likely to have neighbours at an earlier point in time.

An issue could potentially arise if villages without past participating neighbors (i.e. control villages), happen to be villages that have no neighbors in general. There are a number of ways I will address this possible issue with identification. First, one thing to note is that only 9% of villages had no neighbors within 2km, and dropping these villages does not alter the results. For the most part, matching will solve this issue as only similar villages across treated and control groups will be matched. Beyond this, in robustness checks I will also examine exact matching of villages by the number of neighbors in the study, and show that the results continue to strongly hold. Finally, I will also show that there is no significant relationship between contributions in the public goods games and the number of sampled villages within 2km.

One reason this is less of an issue than first appearances may suggest is that the variable is defined as neighboring villages *in the study*. Since only 150 villages out of 598 participated in the study, the variable itself is only correlated with the actual number of neighbors. This is important, as one may be worried that there are unobservables correlated with the true number of neighbors within 2km.

Table 2: Logit regression for treatment: potential communication

Logit Regression	**Treatment**
	(1)
Variables available to planner:	
Distance to base (km)	-0.054
	(0.118)
Population	0.001
	(0.006)
# Villages $\leq$ 2 km	0.858***
	(0.194)
Observations	143
Sector Fixed Effects	Yes

Note: *** p<0.01, ** p<0.05, * p<0.1. Robust standard errors are reported in parenthesis.

Next Table 3 examines the balance across treated villages: those with at least one neighboring villager who previously participated in the public goods games, and control villages. Intuitively, statistically significant imbalances are found for variables which the planner had available, namely the distance to the study's home base, and the number of study villages nearby. Regarding the variables unavailable to the planner, the only statistically significant variable is in fact the main outcome of interest of this paper, village average contributions in the public goods game. In particular contributions are 29 RWF *greater* in treatment villages, i.e. villages which had neighbors that were past participants.

Since treatment and control villages are similar on observable characteristics related to demographics and preferences, one might in fact interpret the difference in contributions as the unbiased treatment effect, that is the true impact of opportunities for communication on cooperation. However, one might worry that unobservables may vary across the two groups. To account for this possibility, I now continue in the next section to the matching strategy.

3.5 Propensity Score Matching

As stated, this setup is well suited to matching strategies, due to the fact that all observables that the treatment could have been conditioned on are available in the data. While control and treatment villages may have differed on average, matching enables one to compare similar groups of villages who either received or did not receive the treatment respectively.

Here I follow the notation of Imbens and Rubin (2015), with some slight adaptations. Let $C_i(1)$ denote the outcome of interest, contributions, if the individual i was a member of a village who had neighbors who previously participated in the games (treated, i.e. $W_i = 1$), and $C_i(0)$ be the contribution of an individual in a village with no previous neighbours (untreated, i.e. $W_i = 0$). In an ideal world, one could observe both outcomes (treated and untreated) for the same individual, and hence could calculate average treatment effects τ. In the real world, the classic problem is that one cannot obtain an unbiased estimate of the treatment effect by naive comparison of the average outcomes of the two groups ($\tau = \bar{C}(1) - \bar{C}(0)$) because these may have different characteristics.

In practice, randomization can solves this problem, by creating comparable treatment and control groups. Here, randomization did not occur. Instead, following Rosenbaum and Rubin (1983) and a number of others, the strategy is to find a set of observable covariates X, which are known to be not affected by the treatment, such that:

$$W_i \perp C_i(1), C_i(0) | X_i. \tag{2}$$

This assumption is referred to as unconfoundedness or selection on observables. It means that the outcomes are uncorrelated with treatment, conditional on covariates X_i.

In the current context this assumption is likely to be satisfied. The reason is that, unlike most observational studies, the treatment W_i (being exposed to villages who previously participated) could only have been conditioned on observables. This is because, as stated earlier, the planner determined the order of visits, in advance, with a limited number of pre-visit observables. In particular, it would be impossible for the planner to condition the treatment on unobservables or on the games outcomes, since these were not available at

Table 3: Balance

	Control	Treatment	Difference
Available to planner			
Distance to base (km)	15.02	11.91	
			3.10*
Population	132.24	131.32	
			0.92
# Villages $\leq$ 2 km	1.84	3.44	
			-1.59***
Unavailable to planner			
Contributions	237.15	266.50	
			-29.35**
Risk Aversion Index	3.95	3.90	
			0.05
Female	0.74	0.74	
			-0.00
Age	35.56	35.27	
			0.29
Education	4.44	4.49	
			-0.05
Believes cooperates	1.37	1.37	
			-0.01
Believes exerts effort	1.41	1.42	
			-0.01
Number of Participants Known	2.64	2.51	
			0.14
Believes supportive	1.36	1.35	
			0.01
Believes trustworthy	1.27	1.27	
			-0.00
Observations	58	91	149

the planning stage.

The variables available to the planner were the following.

1. Village location.

2. Distance to village (from base location).

3. Sector (political region, below the district level).

4. Number of households in the village.

Denote the propensity score, $e(x)$ by:

$$e(x) = Pr(W_i = 1|X_i = x), \tag{3}$$

i.e. the probability that an individual receives the treatment conditional on having characteristics $X_i = x$. This is also equivalent to the expectation of the treatment, $\mathbb{E}[W_i = 1|X_i = x]$.

We can thus define the average treatment effect as:

$$\tau = \mathbb{E}[\mathbb{E}[C_i|W_i = 1, X_i] - \mathbb{E}[C_i|W_i = 0|X_i]] \tag{4}$$

As Imbens and Rubin (2015) note, there are two key assumptions that will be required for the analysis. The first, typically most challenging assumption is unconfoundedness, i.e. Equation 2. Unconfoundedness implies that after controlling for observable characteristics, X_i, assignment to treatment is not correlated with outcomes. In the current setting, unconfoundedness is likely to be satisfied, as previously discussed. The second assumption involves overlap in the distribution of covariates across treatment and control villages. Intuitively speaking, one needs to be able to find similar villages in control and treatment groups, in order to make valid comparisons.

Regarding the second assumption, overlap, as noted earlier, in Table 3, control villages were located farther away from the base location and had lower numbers of neighboring villages within two kilometers. For the most part, there is substantive overlap on covariates between treated and untreated villages.

3.6 Estimating the Propensity Score

The propensity score needs to be estimated from the covariates which are assumed to have an impact on which villages received the "treatment" (having neighbors that previously participated). In the case of this study, these variables can only come from the set of all observables available to the planner at the time the order of visits was determined. Including all of these variables is essential. I additionally include the geographical variables of latitude and longitude, since the planner also knew the geographic location of these villages. It will be important later to add political sectors, I only refrain from doing this early to get an initial look at the data focusing on a smaller number of initial variables without increasing the dimensionality.

Beyond this I include higher orders of these variables and interactions, as one cannot assume that the planner used the available variables in a linear way. To determine the precise specification, I follow the algorithm outlined in Imbens and Rubin (2015), which involves selection of these higher order terms based on their added value in terms of predicting treatment assignment. The algorithm involves step-wise regression estimation of the propensity score, and involves the possibility of dropping covariates should they add no value in determining treatment status. The algorithm does not select any higher order terms, hence the final terms selected for estimation of the propensity score are the four variables corresponding to distance from base, village population, village density within 2km.

Note that political sectors may also be important. As such I will also examine specifications which involve different types of matching on sectors. First I will include political sectors in the propensity score estimation. Next, I will also consider specifications which require exact matching on sectors. That is, I require that in addition to villages being similar across treatment and control according to available variables to the planner, I also require that these villages be located in the same sector.

Figure 3 presents the distribution of the propensity score by treatment status. While there is overlap over the sample, it is important to note that overlap is problematic, particularly at lower ranges. In the analysis I will impose restrictions that matching must occur in regions with common support.

Figure 3: Propensity Scores by Treatment Status

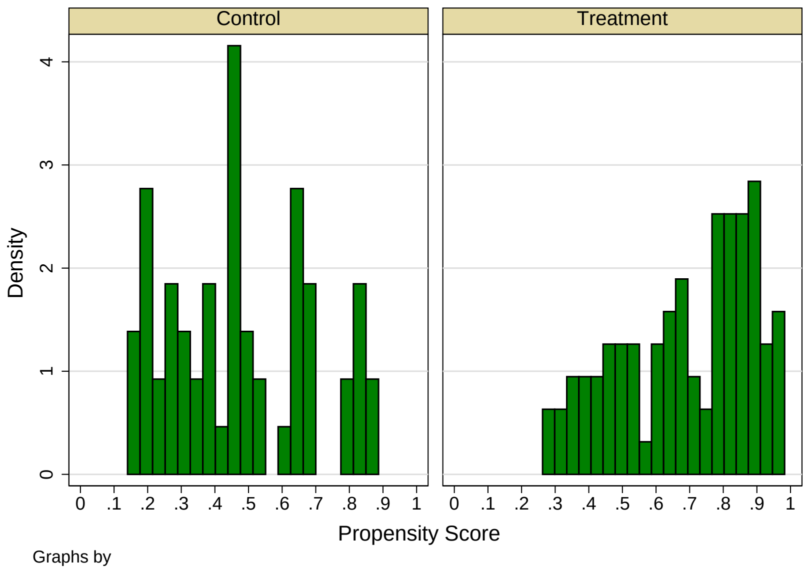

Histogram of the propensity score by treatment status.

3.7 Results on Average Treatment Effects

3.7.1 Initial Matching Estimates

Table 4 presents the first empirical specification. The average treatment effect estimated varies depending on the number of neighbors used in the matching strategy. The first row does not include political sector dummies in the calculation of the propensity score, while the second row includes these variables.

The average treatment effects range from 21.28 to 32.95 RWF. In the first row which does not include sectors in the propensity score, the effects are always significant ranging from the 1% to 10% level depending on the specification. In the second row which includes sectors, the average treatment effect is only significant at the 5% level when examining one neighbor matches.[20] This corresponds to a 8-13% increase in contributions over the entire sample.

I now turn to the main empirical analysis where I strengthen the matching by requiring

[20]In the other two columns it approaches significance, at the 15% level.

that villages be exactly matched on particular important attributes.

Table 4: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Sectors Excluded from Propensity Score			
Contribution	26.864*	32.953***	29.301**
	(13.713)	(11.927)	(12.000)
Observations	118	118	118
Sectors Included in Propensity Score			
Contribution	31.277**	21.667	21.275
	(14.473)	(13.265)	(13.595)
Observations	132	132	132

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

3.7.2 Exact Matching on Sector and Village Density

In this section I improve the matching estimates by ensuring exact matching on two important attributes. The first attribute is the number of villages (in the study) that are within 2km of a given village, referred to henceforth as village density. By construction, the treatment variable, which is 1 whenever a village had neighbors located within 2km that previously participated, is highly correlated with the number of total neighbors within 2km in the study. Further, the initial balance checks indeed revealed that village density was not balanced across treatment and control villages.

To account for the possibility that village density may be correlated with unobserved variables, and that the earlier propensity score matching may have been unable to adequately control for this, I conduct a matching strategy where I require that matched treatment and control villages *must* have exactly the same number of neighbors in the study. Village density has a minimum value of 0 and a maximum of 7.

The first two rows in Table 5 presents this matching analysis. The first row includes the whole usable sample[21], while the second row further restricts this to exclude a category of village density which contained only 2 control villages. The motivation for this restriction is to reduce noisiness in the matching which may result when a small number of control villages are re-matched to multiple treatment villages.

The results in rows 1 and 2 of Table 5 suggest treatment effects may be more substantial than initially estimated. The effects range from 19.73 to 37.41, a 8% to 15% increase in average contributions.[22]

The next part of Table 5 examines exact matching on a different dimension, political sectors. The Rusizi district of Rwanda is divided into 18 political sectors. Rows 3 and 4 consider exact matches on these sectors: requiring that treatment and control villages may only be matched should they be located within the same sector. It is intuitive to create such pairings, since from the perspective of the planner, villages within the same sector would appear nearly identical in terms of location. Further, there is the possibility that villages located within a particular sector vary by some unobservable variables correlated with the treatment assignment.

The number of villages per sector ranges from a minimum of 3 to a maximum of 13. Row 3 shows that while the magnitude of the estimated average treatment effect is large and comparable to previous results, ranging from 27.74 to 30.87, the effects are not significant due to large standard errors. This is not surprising, as here the matching process involves much more noise, since villages now need to be matched among a much smaller pool.

To eliminate some of the noise from this matching procedure, row 4 considers only sectors that have at least 2 villages per treatment or per control group.[23] One can see that this is successful in reducing some of the noise, as the standard errors are reduced substantially. From row 4, one can see that the estimated average treatment effects range from 32.61 to 36.5, significant at the 5% to 10% level, and corresponding to a range of percentage increases of 13-14%.

[21] Less 21 villages which either had no treatment villages or no control villages for a specific village density, and thus could not be matched.

[22] The estimate of 19.73 should be interpreted with some caution as well, since it involves some villages which do not have 3 neighbors in the opposite treatment group, and thus are in fact matched with fewer than 3 neighbors in certain instances.

[23] Previously, in row 3, there were sectors with only 1 control village or 1 treatment village that was being matched with replacement to multiple villages with the opposite treatment status.

Table 5: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Exact Matching on Village Density			
Contribution	31.289*	23.164*	19.734
	(17.376)	(14.047)	(13.053)
Observations	121	121	121
Exact Matching on Village Density (Excluding group with 2 villages in C)			
Contribution	37.406**	28.909*	24.715*
	(18.919)	(14.977)	(13.627)
Observations	108	108	108
Exact Matching on Political Sectors			
Contribution	30.867	27.742	29.172
	(21.793)	(23.745)	(24.548)
Observations	114	114	114
Exact Matching on Political Sectors (Excluding sectors with one village in T or C)			
Contribution	36.497**	32.606*	34.425*
	(17.410)	(17.813)	(17.825)
Observations	102	102	102

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

3.7.3 Robustness Checks

In the Appendix, Section 5.1 examines these specifications for different treatment definitions that involve different distances. There is evidence that shorter distances are associated with large treatment effects, though there also arise issues of sample size, due to few treatment villages when distances are shorter.

4 Conclusion

Experiments conducted in lab and field settings provide valuable insight on regional variation in preferences, and how these preferences relate to economic behavior. When individual or social learning occurs outside of the design of the experiment, the external validity of these insights is weakened. In this paper I studied a large implementation of public goods experiments across villages in rural Rwanda. Because of normal logistics, villages participated at different dates, and because of natural variation in geography this led to variation in opportunities for communication between past and future participants.

Theoretically, such communication is expected to have no effects. However, previous laboratory experiments have shown that both advice, as well as cheap-talk have increased observed contributions in public goods games. Chaudhuri et al. (2006) find that advice alters beliefs about average contributions, which leads to higher contributions from conditional cooperators. Isaac and Walker (1988a) discusses the importance of communication in helping to build credibility around expected decisions of group members. Such credibility is only relevant in theory, insofar as individuals behave in a conditionally cooperative manner.

Using a matching strategy which matched on a propensity score estimated using solely variables available to the planner, I find substantial effects of communication. The estimates point to a 8-14% average treatment effect on contributions. These results have important implications for the design of field experiments, and the interpretation of results. Social learning may occur if past participants communicate with future ones, and this communication can change behavior in meaningful ways.

5 Appendix

5.1 Robustness Checks for Different Distances

The following presents the matching analysis for difference distances. Most of the patterns present in the earlier analysis can be seen across these tables. It is important to note that by changing the distance the interpretation of the treatment also changes.

First with short distances, such as 1km, there are very few villages that have neighbors who previously participated. This is the reason why the sample size is so low, and the standard errors so large. Second, with longer distances, such as 3km, there is the opposite issue - the number of treatment villages is high. Moreover, it is not clear whether the definition of treatment (at least one neighbor that previously participated within the specified radius) is suitable. One could argue that the control group should be expanded to include villages with either 0 or 1 neighbors that previously participated.

Of further note is that the propensity score estimation uses the village density variable for the appropriate distance. That is it examines the number of total neighbors within the distances of 1km, 1.5km, 2.5km, and 3km, respectively.

5.1.1 1km

Table 6: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Sectors Excluded from Propensity Score			
Contribution	63.118***	51.239*	57.299**
	(24.075)	(28.203)	(25.017)
Observations	31	31	31
Sectors Included in Propensity Score			
Contribution	54.681*	53.504**	55.339**
	(29.983)	(26.742)	(27.136)
Observations	28	28	28

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

Table 7: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Exact Matching on Village Density			
Contribution	77.273	48.380	44.753
	(52.907)	(52.570)	(52.129)
Observations	18	18	18
Exact Matching on Village Density (Excluding groups with 2 villages in T/C)			
Contribution	64.015	31.510	27.431
	(45.630)	(39.375)	(35.334)
Observations	16	16	16
Exact Matching on Political Sectors			
Contribution	56.970	52.652	50.101
	(60.704)	(69.530)	(73.679)
Observations	25	25	25
Exact Matching on Political Sectors (Excluding sectors with one village in T or C)			
Contribution	61.023	54.110	50.922
	(65.527)	(74.865)	(79.306)
Observations	20	20	20

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

5.1.2 1.5km

Table 8: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Sectors Excluded from Propensity Score			
Contribution	21.026 (16.830)	9.115 (16.121)	16.572 (15.087)
Observations	88	88	88
Sectors Included in Propensity Score			
Contribution	19.949 (14.131)	22.419 (15.293)	19.320 (15.750)
Observations	76	76	76

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

Table 9: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Exact Matching on Village Density			
Contribution	43.658***	30.659**	28.552*
	(14.355)	(14.349)	(15.321)
Observations	84	84	84
Exact Matching on Village Density (Excluding groups with 2 villages in T/C)			
Contribution	35.788**	25.404*	23.719
	(14.776)	(14.848)	(15.882)
Observations	74	74	74
Exact Matching on Political Sectors			
Contribution	39.602*	33.208	34.344
	(22.616)	(25.719)	(27.260)
Observations	80	80	80
Exact Matching on Political Sectors (Excluding sectors with one village in T or C)			
Contribution	48.244**	43.251*	42.746
	(22.216)	(25.117)	(26.739)
Observations	63	63	63

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

5.1.3 2.5km

Table 10: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Sectors Excluded from Propensity Score			
Contribution	24.123	22.218	24.216*
	(17.280)	(14.647)	(13.378)
Observations	119	119	119
Sectors Included in Propensity Score			
Contribution	4.674	12.006	16.344
	(14.686)	(13.538)	(13.205)
Observations	82	82	82

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

Table 11: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Exact Matching on Village Density			
Contribution	18.815 (17.003)	22.596 (15.849)	17.443 (15.553)
Observations	101	101	101
Exact Matching on Village Density (Excluding groups with 2 villages in T/C)			
Contribution	24.085 (16.638)	23.880 (15.112)	18.448 (14.305)
Observations	89	89	89
Exact Matching on Political Sectors			
Contribution	33.158 (21.397)	26.595 (23.639)	23.210 (24.804)
Observations	104	104	104
Exact Matching on Political Sectors (Excluding sectors with one village in T or C)			
Contribution	31.434 (19.548)	25.127 (21.190)	21.480 (22.124)
Observations	95	95	95

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

5.1.4 3km

Table 12: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Sectors Excluded from Propensity Score			
Contribution	17.120 (15.351)	16.198 (14.433)	18.939 (15.392)
Observations	125	125	125
Sectors Included in Propensity Score			
Contribution	7.468 (12.781)	2.930 (13.307)	7.905 (13.016)
Observations	83	83	83

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

Table 13: Average Effect of Presence of Past Participating Neighbors

	(1) 1 Neighbor	(2) 2 Neighbors	(3) 3 Neighbors
Exact Matching on Village Density			
Contribution	18.811	15.794	15.198
	(14.554)	(14.136)	(13.899)
Observations	92	92	92
Exact Matching on Village Density (Excluding group with 2 villages in C)			
Contribution	6.320	5.159	4.641
	(19.330)	(17.943)	(17.104)
Observations	60	60	60
Exact Matching on Political Sectors			
Contribution	28.262	27.923	23.174
	(23.853)	(26.837)	(28.397)
Observations	97	97	97
Exact Matching on Political Sectors (Excluding sectors with one village in T or C)			
Contribution	25.798	25.990	20.876
	(21.868)	(24.374)	(25.773)
Observations	89	89	89

Analysis uses nearest neighbor propensity score matching, with replacement. Significantly different from zero at * 0.1; ** 0.05; *** 0.01. Abadie-Imbens Robust Standard Errors. Values of propensity score outside common support range are dropped.

5.2 Alternative Estimation Strategies

5.2.1 OLS

In this section I examine an alternative OLS regression specification for a different dependent variable, the order of visit within a sector. The notation is as follows. $C_{i,j,s}$ is individual contributions, $O_{j,s}$ is within sector visit order. It ranges from 0 to 7 depending on whether the village was the first or the last to be visited, within the political sector. $X_{i,j,s}$ is a vector of individual characteristics, $K_{j,s}$ is a vector of village characteristics, Γ_s is a vector of sector level fixed effects, and $\epsilon_{i,j,s}$ is a residual error term.

Because $C_{i,j,s} \in \{0, 100, 200, 300, 400\}$ is discrete, OLS is misspecified. Nonetheless, it is a useful comparison exercise to compute OLS estimates, which have the advantage of being more readily interpretable. As in the previous analysis, robust standard errors are clustered at the village level. Unlike the ordered probit analysis, $\epsilon_{i,j,s}$ is not assumed to be a standard normal residual.

$$C_{i,j,s} = \alpha O_{j,s} + \beta X_{i,j,s} + \delta K_{j,s} + \gamma \Gamma_s + \epsilon_{i,j,s} \tag{5}$$

Table 14: Effects of Social Learning on Contributions (OLS)

Dependent Variable: Contribution		
$O_{j,s}$	10.878***	17.650***
	(3.477)	(4.518)
Date of Visit	0.311	−2.731**
	(0.479)	(1.377)
Participant is Female	−7.609	−8.071
	(8.504)	(8.576)
Years of Education	−0.327	−0.535
	(1.354)	(1.346)
Age	6.030***	6.147***
	(1.956)	(1.925)
Age^2	−0.059**	−0.062***
	(0.023)	(0.023)
Number of Participants Known	−2.387	−2.305
	(2.067)	(2.087)
Has Mobile Phone	4.604	4.379
	(8.357)	(8.330)
Female Coordinator	12.978	70.782***
	(22.492)	(21.446)
Average Female	68.328*	72.623*
	(39.018)	(39.694)
Average Number Known	−8.263	−10.087*
	(5.744)	(5.462)
Village Characteristics	Yes	Yes
Sector Fixed Effects	No	Yes
R^2	0.06	0.09
Observations	1692	1692

Analysis uses OLS regression. Difference is significant at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at village level.

5.3 Alternative Empirical Specifications

In this section I examine a number of other empirical specifications that look at alternatives to proxy for social learning.

I next construct an index that is designed to be correlated with social distance to previous participants.

The index I create consists of a measure of how many previously participating neighbors an individual has, and how close these neighbors are. In order to reduce the correlation with time, I examine only neighbors that participated in the previous 7 days.

I define the index, denoted by (I_i), as follows. $I_i = (45 - \mu_d) \cdot n_i$ where n_i is the number of individuals within a 10km radius of i that participated in the previous 7 days, and μ_d is the average distance (in km) of i to those individuals. This index uses the intuition that an individual is more likely to have communicated with a previous participant when that participant is near, and is more likely to have had communication when there are a large number of previous participants. 45 km is chosen as the maximum distance - no two individuals in the data are more than 45km apart. Finally, the index is scaled to be in between 0 and 1.

Table 15 presents the empirical analysis, using both OLS and an ordered probit specification. Examining the OLS regression, the coefficient of interest is approximately 89 RWF and significant at the 1% level, meaning an individual at the very high end of the index (many neighbors that previously participated and are close) is expected to contribute 89 RWF more than an individual at the other end (no recent participating neighbors). This represents 35% of average contributions, an extremely large effect. The average of I_i is 0.61, indicating an average effect of an increase of 54 RWF (21% increase above average).

Table 15: Effect of Information Index on Contributions

Dependent Variable: Contribution	OLS	Ordered Probit
I_i	89.686***	0.291***
	(32.717)	(0.106)
Date of Visit	−1.297	−0.005
	(1.295)	(0.004)
Participant is Female	−8.476	−0.017
	(8.502)	(0.026)
Years of Education	−0.594	−0.001
	(1.340)	(0.004)
Age	5.904***	0.019***
	(1.908)	(0.006)
Age^2	−0.059***	−0.000***
	(0.023)	(0.000)
Number of Participants Known	−2.274	−0.007
	(2.080)	(0.006)
Has Mobile Phone	4.564	0.013
	(8.344)	(0.026)
Female Coordinator	76.190***	0.265***
	(22.242)	(0.081)
Average Female	84.441**	0.268**
	(40.876)	(0.130)
Average Number Known	−7.031	−0.021
	(5.701)	(0.018)
Village Characteristics	Yes	Yes
Sector Fixed Effects	Yes	Yes
†(Pseudo) R^2	0.09	0.03†
Observations	1692	1692

Analysis uses OLS regression in column 1 and ordered probit in column 2. Marginal effects on the probability of contributing the maximum of 400 RWF are calcuated in column 2. Difference is significant at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at village level.

A further specification is to examine the number of villages visited previously within the same sector. Table 16 presents this analysis. As a robustness check, Table 17 examines the total number of villages visited in the sector, regardless of whether they were visited previously or not. That this variable is non-significant provides some reassurance that the result in Table 16 is not being driven by sectors that simply have more villages.

Table 16: Effect of Number Visited Before on Contributions

Dependent Variable: Contribution	OLS	Ordered Probit
Number Visited Before (in Sector)	6.355***	0.020***
	(2.143)	(0.007)
Date of Visit	−1.668	−0.006
	(1.270)	(0.004)
Participant is Female	−8.289	−0.016
	(8.562)	(0.026)
Years of Education	−0.556	−0.001
	(1.349)	(0.004)
Age	6.125***	0.019***
	(1.925)	(0.006)
Age^2	−0.062***	−0.000***
	(0.023)	(0.000)
Number of Participants Known	−2.281	−0.007
	(2.086)	(0.006)
Has Mobile Phone	3.891	0.011
	(8.320)	(0.026)
Female Coordinator	68.059***	0.243***
	(21.322)	(0.079)
Average Female	68.750*	0.214*
	(40.412)	(0.128)
Average Number Known	−9.810*	−0.030*
	(5.514)	(0.017)
Village Characteristics	Yes	Yes
Sector Fixed Effects	Yes	Yes
†(Pseudo) R^2	0.08	0.03†
Observations	1692	1692

Analysis uses OLS regression in column 1 and ordered probit in column 2. Marginal effects on the probability of contributing the maximum of 400 RWF are calcuated in column 2. Difference is significant at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at village level.

As a robustness check I present in Table 17 the total number of villages within the same sector, visited at any date. This variable should not have any strong relationship with contributions, since neighbors that are yet to participate should clearly have no impact on current contributions. That the variable is not significant provides some evidence that the previous results are not driven by having more neighbors generally. That it is positive is to be expected, since it is highly correlated with the number visited before.

Table 17: Effect of Number of Villages in Sector (Robustness Check)

Dependent Variable: Contribution	OLS	Ordered Probit
Total Number in Sector	11.981	0.012
	(31.418)	(0.098)
Date of Visit	0.155	0.000
	(1.507)	(0.005)
Participant is Female	−8.366	−0.016
	(8.529)	(0.026)
Years of Education	−0.714	−0.002
	(1.353)	(0.004)
Age	5.960***	0.019***
	(1.912)	(0.006)
Age^2	−0.060***	−0.000***
	(0.023)	(0.000)
Number of Participants Known	−2.244	−0.007
	(2.079)	(0.006)
Has Mobile Phone	4.720	0.013
	(8.408)	(0.026)
Female Coordinator	60.238***	0.210**
	(23.051)	(0.083)
Average Female	60.659	0.189
	(41.257)	(0.130)
Average Number Known	−8.580	−0.026
	(5.633)	(0.017)
Village Characteristics	Yes	Yes
Sector Fixed Effects	Yes	Yes
†(Pseudo) R^2	0.08	0.03†
Observations	1692	1692

Analysis uses OLS regression in column 1 and ordered probit in column 2. Marginal effects on the probability of contributing the maximum of 400 RWF are calcuated in column 2. Difference is significant at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at village level.

Finally, I examine the proportion of villages visited before, in Table 18. This variable is constructed by taking the number visited before within the sector and dividing it by the total number of villages in the sector.

Table 18: Effect of Proportion Visited Before on Contributions

Dependent Variable: Contribution	OLS	Ordered Probit
Proportion Visited Before (in Sector)	50.843**	0.157**
	(22.747)	(0.072)
Date of Visit	−1.543	−0.005
	(1.407)	(0.005)
Participant is Female	−8.318	−0.016
	(8.544)	(0.026)
Years of Education	−0.652	−0.002
	(1.355)	(0.004)
Age	6.029***	0.019***
	(1.922)	(0.006)
Age^2	−0.061***	−0.000***
	(0.023)	(0.000)
Number of Participants Known	−2.245	−0.007
	(2.083)	(0.006)
Has Mobile Phone	4.260	0.012
	(8.342)	(0.026)
Female Coordinator	62.395***	0.222***
	(22.363)	(0.082)
Average Female	63.692	0.197
	(40.787)	(0.129)
Average Number Known	−9.847*	−0.030*
	(5.453)	(0.017)
Village Characteristics	Yes	Yes
Sector Fixed Effects	Yes	Yes
†(Pseudo) R^2	0.08	0.03†
Observations	1692	1692

Analysis uses OLS regression in column 1 and ordered probit in column 2. Marginal effects on the probability of contributing the maximum of 400 RWF are calcuated in column 2. Difference is significant at * 0.1; ** 0.05; *** 0.01. Robust standard errors clustered at village level.

References

Avdeenko, Alexandra and Michael J Gilligan, "International Interventions to Build Social Capital: Evidence from a Field Experiment in Sudan," *American Political Science Review*, 2015, pp. 1–45.

Bednar, Jenna, Yan Chen, Tracy Xiao Liu, and Scott Page, "Behavioral spillovers and cognitive load in multiple games: An experimental study," *Games and Economic*

Behavior, jan 2012, *74* (1), 12–31.

Bochet, Olivier and Louis Putterman, "Not just babble: Opening the black box of communication in a voluntary contribution experiment," *European Economic Review*, 2009, *53* (3), 309–326.

Brosig, Jeannette, J Weimann, Axel Ockenfels, and Joachim Weinmann, "The effect of communication media on cooperation," *German Economic Review*, 2003, *4* (2), 217–241.

Celen, B., S. Kariv, and a. Schotter, "An Experimental Test of Advice and Social Learning," *Management Science*, 2010, *56* (10), 1687–1701.

Chaudhuri, Ananish, "Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature," *Experimental Economics*, 2011, *14*, 47–83.

_ , Andrew Schotter, and Barry Sopher, "Talking ourselves to efficiency: Coordination in inter-generational minimum effort games with private, Almost Common and Common Knowledge of Advice," *Economic Journal*, 2009, *119* (534), 91–122.

_ , Sara Graziano, and Pushkar Maitra, "Social Learning and Norms in a Public Goods Experiment with Inter-Generational Advice1," *Review of Economic Studies*, apr 2006, *73* (2), 357–380.

Crawford, Vincent, "A Survey of Experiments on Communication via Cheap Talk," *Journal of Economic Theory*, 1998, *78* (2), 286–298.

Ding, Tingting and Andrew Schotter, "Intergenerational Advice and Matching : An Experimental Study," *mimeo*, 2015.

Dupas, Pascaline, "What Matters (and What Does Not) in Households' Decision to Invest in Malaria Prevention?," *American Economic Review*, apr 2009, *99* (2), 224–230.

_ , "Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence From a Field Experiment," *Econometrica*, 2014, *82* (1), 197–228.

Gneezy, Uri and Alex Imas, "Lab in the field: Measuring preferences in the wild," *Handbook of Field Experiments*, 2017.

Imbens, Guido W. and Donald B. Rubin, *Causal inference: For statistics, social, and biomedical sciences an introduction* 2015.

Isaac, R. Mark and James M. Walker, "Communication and free-riding behavior: the voluntary contributions mechanism," *Economic Inquiry*, 1988, *26* (4), 585–608.

_ **and James M Walker**, "Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism," *The Quarterly Journal of Economics*, 1988, *103* (1), 179–199.

Jakiela, Pamela and Owen Ozier, "Does Africa Need a Rotten Kin Theorem ? Experimental Evidence from Village Economies," *Review of Economic Studies*, 2015.

Kosfeld, Michael and Devesh Rustagi, "Leader Punishment and Cooperation in Groups: Experimental Field Evidence from Commons Management in Ethiopia †," *American Economic Review*, 2015, *105* (2), 747–783.

Ledyard, John O., "Public Goods: A Survey of Experimental Research," *Social Science*, jan 1995, *35* (12), 111–194.

Rosenbaum, Paul R and Donald B Rubin, "Biometrika Trust The Central Role of the Propensity Score in Observational Studies for Causal Effects The central role of the propensity score in observational studies for causal effects," *Source: Biometrika Biometrika*, 1983, *70* (1), 41–55.

Sally, David, "Conversation and Cooperation in Social Dilemmas: A Meta-Analysis of Experiments from 1958 to 1992," *Rationality and Society*, jan 1995, *7* (1), 58–92.