Development Assistance and Collective Action Capacity: Results from a Field Experiment in Post-Conflict Liberia

James D. Fearon* **Macartan Humphreys†** **Jeremy M. Weinstein‡**

December 12, 2013

Abstract

Social cooperation is critical to a wide variety of political and economic outcomes. For this reason, international donors have embraced interventions designed to strengthen the ability of communities to solve collective-action problems, especially in post-conflict settings. We exploit the random assignment of a development program in Liberia to assess the effects of such international interventions. Using a matching funds experiment we find evidence that these external interventions can alter cooperation capacity. However, we also find that effects are only observed in communities where, by design, both men and women face collective action dilemmas. Focusing on mechanisms, we find that effects likely worked through improving mobilization capacity which may have enhanced the ability of communities to coordinate. These gains did not operate however in areas where only women took part and where coordination dilemmas appear reduced. The combined evidence suggests that the impact of such external interventions can depend critically on the kinds of social dilemmas that communities face, and the flexibility they have in determining who should solve them.

*Stanford University. We thank the International Rescue Committee (IRC) for their partnership in undertaking this study; Jodi Nelson and Liz McBride played a key role in designing the instruments while IRC field staff, both in Monrovia and Lofa County, provided support on the ground. We are particularly grateful to Nicholai Lidow and Gwendolyn Taylor for leading the game and follow-up survey data collection teams in the field, and to Andrea Abel, Jessica Gottlieb, and Amanda Robinson for their fieldwork on the public goods games. We thank the National Ex-Combatant Peace-building Initiative for their research support in Liberia, in particular K. Johnson Borh and Morlee Zawoo, and Brian Coyne for work on data checking. We acknowledge the support of a number of funders that made this study possible. DFID provided the bulk of funding for the panel survey as part of its initial grant to IRC. The Center for Global Development provided some additional funding for the second wave survey. AUSAID, through a grant to the Center for Global Development, provided the resources to implement the behavioral protocol for the measurement of social cohesion. The Center for Democracy, Development, and the Rule of Law provided funding for field work expenses and the International Growth Center provided support for final data compilation and analysis.

†Columbia University

‡Stanford University

1 Introduction

The ability of political units to generate and sustain cooperation is part of what distinguishes successful states from those that fail, communities with high-levels of service provision from those that lack essential services, societies with well-functioning democratic institutions from those that underperform, and political groups that achieve power and influence from those that find themselves stuck on the sidelines. Prior research suggests that cooperative behavior is a relatively stable characteristic of a political unit – reflective of demographic, economic, and political factors that have deep historical roots – and thus unlikely to respond to short-term interventions. For example, in some accounts, the social capital that supports well-functioning local governments in Northern Italy derives from the 14th century, modern anti-semitism and intra-group tensions have their roots in the medieval period, and contemporary levels of distrust in Africa relate to historical exposure to the slave trade (Putnam (1994); Voigtländer and Voth (2012); Nunn and Wantchekon (2011)).

Yet, over the past decade, donor agencies have embraced a variety of participatory development strategies on the belief that externally-funded programs can enhance the prospects for local cooperation and effective governance and that these changes can be brought about relatively quickly (Mansuri and Rao, 2012). One of the largest such aid models is known as "Community-Driven Development" (CDD), or "Community-Driven Reconstruction" (CDR) when applied in post-conflict settings. In CDD programs, the delivery of block grants is combined with efforts to build local governance capabilities. Technical assistance is provided to help communities organize elections to community development councils (CDCs) that then select specific projects in a public process and help to implement and maintain them. Reflecting a broader movement in international development toward local empowerment, donor agencies believe that community involvement will make projects more successful in terms of their impact on material welfare. Advocates also promise that CDD will improve local governance and the capacity of communities for collective action to provide and manage public goods. In particular, CDD projects are designed to promote more democratic institutions and practices in communities where, donors presume, traditional political structures are more autocratic

and the possibilities for cooperation more limited. Depite the substantial investment in these programs, until recently, there has been little empirical evidence to support any of these claims (Mansuri and Rao, 2004).

In this article, we present the results of a field experiment evaluating a CDR program implemented in two districts of northern Liberia, roughly from November 2006 to March 2008. The intervention, funded by the UK government and implemented by the International Rescue Committee (IRC), worked in 42 communities that were randomly sampled in the fall of 2006 from a pool of 83 eligible communities. Outcomes were measured in two main ways: by baseline and follow-up household surveys that asked about material well-being and attitudes regarding community governance, and by a "matching funds experiment" conducted in all 83 communities in the summer of 2008 after the CDR project had come to end. In the matching funds experiment, treatment and control communities were invited to organize to receive up to $420 (U.S.) for a new local development project. In order to participate, communities needed to decide how the funds would be spent, select three community representatives to handle the funds, and participate collectively in a contribution game to determine the total amount of the grant. In the contribution game, 24 randomly selected adults from each community were given 300 Liberian dollars (about $5) and asked to make a private decision about how much to keep and how much to contribute to a community fund for a new development project. They had been instructed that their contributions would be matched (at two different, known "interest rates") by us at a community meeting held after the game.

The article makes five main contributions. First, we provide evidence of the causal impact of a prominent model of development aid designed to introduce new governance practices in local communities. While data from surveys asking individuals about material welfare and opinions on community governance can be suggestive, the matching funds experiment allows us to assess whether the CDR program affected behavior in a real-world, community-wide collective action problem. We observe both individual- and community-level behavioral outcomes and decision-making processes. Random assignment of the CDR intervention allows us to assess whether patterns of social cooperation respond to external intervention, even when

underlying demographic, economic, and political factors remain unchanged.

We observe a significant impact of the CDR program on the collective action capacity of treatment communities. Communities that received the program contributed significantly more in our matching funds experiment, collecting on average 82.1% of the total possible payout versus 75.8% in control communities. At the individual level, the impact of exposure to the CDR program was approximately the same as the impact of increasing the rate of return on contributions from 100% to 400% (the interest rate manipulation mentioned above).

Second, by varying the types of collective action problem that communities face, we document striking variation in effects. Surprisingly, we find that the CDR impact is concentrated entirely in one of two treatment arms that we introduced at the measurement stage: in 42 of the 83 communities, 12 men and 12 women were randomly selected to play the contribution game, whereas in the other 41 communities, 24 women played (and communities knew that this would occur). The estimated CDR impact was very large in the mixed group communities, raising average contributions from 67 to 82% of the total possible, but nonexistent in the communities where only women could contribute. In the "all women" communities, total contributions averaged about 84% of the total in both CDR and no-CDR cases. Importantly, this pattern cannot be accounted for by simple heterogeneity of treatment effects, in particular the difference in outcomes is not explained by women contributing more regardless. In the mixed groups, women and men contributed similar amounts on average and responded similarly to the CDR treatment.

Third, we explore systematically the mechanisms that might account for the estimated causal effect of CDR. This focus on mechanisms is essential if we are to generalize the results of any given field experiment and project them to other settings. This paper is one of a small number of experimental studies of CDR programs begun in the period 2006-2008, and already it is evident that the results are not consistent across studies. Casey, Glennerster, and Miguel (2012), who examined a broadly similar program in Sierra Leone, find no effect on community decision-making processes or collective action capacity, although they find positive impacts on some measures of material welfare.[1] Beath, Christia, and Enikolopov (2013) find that

[1]By contrast, we find no significant impact of the CDR program in Liberia on measures of material

the combination of CDCs created by a CDR program in Afghanistan and the assignment of responsibility for food aid distribution to CDCs improved targeting of food aid to marginalized community members. Humphreys, Sanchez de la Sierra, and van der Windt (2013) study a CDR program in eastern Democratic Republic of the Congo that, like the Liberia project assessed here, was administered by the IRC; they find no impact on indicators of governance performance or material welfare. Avdeenko and Gilligan (2013) study a related program in Sudan and find no effects on social capital using behavioral games but do observe gains on their survey measures.

Do these results differ due to differences in program design, program implementation, variation in social or political contexts, or due to complex interactions between these factors? This question is made more difficult by the fact that in these studies and in most field experiments examining foreign aid programs, the "treatments" are complex bundles of actions, events, and resources. Thus the randomization does not by itself allow inference about what aspects of the complex treatment caused the observed effects. The methodological problem of inferring the mechanisms through which an experimental manipulation worked is extremely difficult, as shown by recent work by Imai, Keele, and Tingley (2010) and Green, Ha, and Bullock (2010). But some understanding of mechanisms is important in order to gain useful social scientific knowledge from the results — that is knowledge of effects that plausibly extend beyond the study area, or perhaps the population from which the study was drawn. Thus field experiments generally include speculation about what mechanisms might account for causal effects (or non-effects) identified by the experimental designs.[2]

This article employs a simple model to help distinguish among a number of possible causal paths by which CDR might have an impact, and draws out testable implications of each

welfare; these results are from the post-program surveys implemented several months before we returned to carry of the matching funds experiment. We examine those survey results in detail in another paper [Ref Ommited].

[2]This point holds even for experiments that find no effects of the treatment. Why did the treatment fail to "work"? Was it something about this particular setting? For a related argument about the value of designing experiments to test mechanisms rather than to evaluate broad policies, see Ludwig, Kling, and Mullainathan (2011). Habyarimana, Humphreys, Posner, and Weinstein (2007) use "in the field" decision-making experiments to evaluate different mechanisms that have been proposed to explain greater cooperation among co-ethnics.

mechanism. We then use survey and behavioral measures, taking advantage of the fact that our expectations differ for mixed gender communities where CDR "worked" and women only communities where CDR had no impact, to assess the plausibility of different causal pathways. We find fairly strong evidence that CDR did not increase contributions in the mixed communities by directly increasing individuals' value for public goods, trust in local leadership or foreign NGOs, or fear of punishment for not contributing. Instead, it appears that in CDR communities where it was known that both men and women could be chosen for the contribution game, community leaders engaged in greater mobilization and information-sharing efforts in the week prior to the play of the contribution game.[3] In the all-women communities, mobilization occurred in both CDR and non-CDR communities. Our assessment is that that the prior experience of the CDR program improved the ability of communities to "solve" a non-traditional, mixed gender collective action problem, whereas this problem was easier in the "all women" communities either because traditional women's networks could be used or because there were half as many people to mobilize and inform, or both.

Fourth, our analysis makes a contribution to methods. Over the last decade there has been growing interest in the use of experimental games for the measurement of social outcomes, including in post conflict areas (Whitt and Wilson, 2007; Coleman and Lopez, 2012; Voors, Nillesen, Verwimp, Bulte, Lensink, and Soest, 2012). These studies use games as measurement devices to estimate effects but generally do not enjoy experimental variation to claim that measured effects are identified. In our study, by layering experimental games over an experimental treatment, we benefit from both controlled measurement and strong identification. Our results provide a proof of concept for the use of games for this purpose and demonstrate that they can pick up effects of interest. But our results also provide grounds for caution in the use and interpretation of measures of this kind. We used two very similar games to assess the effects of an international intervention; both were intended to measure the same concept of interest, social cohesion, and the details of the games differed only in the gender of partners playing. Yet these two versions produced starkly different estimates of treatment effects. Had

[3]Note that communities were informed that the contribution game would be played (and what the rules were) a week in advance, but game players were chosen randomly on the day of play, so that communities could not lobby or pressure specific individuals in intervening week.

we played just one version of the game our overall conclusions regarding the effects of the intervention would have depended entirely on the version that we would have implemented. Sensitivity of results from experimental games to fine details of game design has also been seen in lab settings (Haley and Fessler, 2005; Levitt and List, 2007). Our analysis shows that this sensitivity is inherited in field settings. We draw two conclusions from this sensitivity. First, for research design this sensitivity suggests that if games are to be used for this purpose the game should be tightly, rather than broadly, linked to questions of interest. Moreover researchers should seek to assess the robustness of results to seemingly irrelevant features of game design. Second these results suggest that it is possible that similar-looking studies can produce sharply divergent results across contexts, even if underlying effects are constant. High variability in the outcomes from any one study emphasize the importance of accumulation of knowledge across studies.

Finally, the results speak to larger debates and questions about the role of external actors in changing domestic institutions, governance practices, and patterns of social cooperation. International donors have made major commitments in recent years in these areas. Indeed, Mansuri and Rao (2012) point out that the World Bank alone has invested more than $85 billion in local participatory development alone over the past decade. Critics of these programs have begun to illuminate the limits of these efforts. One major concern is that governments will adopt new practices without actually changing how they operate or what they deliver. Andrews (2013) warns of this kind of "isomorphic mimicry" in which, under pressure from donors, elites adopt forms of governance that are being promoted, with very few meaningful changes in how they function. Others have underscored the risks of elite capture as existing power holders shape new institutions or governance practices to reinforce their status, or as new opportunities for participation attract the most capable and politically connected, increasing the marginalization of the poor (Bardhan (2002); Mansuri and Rao (2012)).

The results in this article point to a third challenge for interventions of this form: CDR appears to improve the efficacy of community leadership and mobilization capability, but only when community members had little choice about how to solve the problem. When community

members were required to work across gender lines in the matching funds experiment, there is strong evidence that the prior experience of organizing to decide on and implement a public good project yielded significant returns in the community's ability to act collectively. But these gains from CDR were not in evidence when the community could rely on more traditional institutional structures, such as single gender networks, to mobilize participation. In this context, existing structures facilitate high levels of collective action regardless of prior exposure to CDR. This underscores the risk of designing interventions that foster new institutions and practices which are no more effective than existing approaches, and not likely to be used given the way that communities themselves approach problems of social cooperation.

In section 2, we provide background on the setting of the CDR program in Liberia, and then describe the intervention and the experimental assignment. Section 3 details the matching funds experiment and the associated contribution game that was implemented in all 83 treatment and control communities, six to seven months after the completion of the CDR intervention. It also presents the main effects of exposure to CDR, and for two additional experimental manipulations we introduced into the game concerning the gender composition of the game players and the interest rate individual players faced on their contributions.

Section 4 turns to the problem of discerning the mechanism(s) by which the CDR program increased collective action capacity, beginning with a model of the decision problem faced by game players and by community elites considering how much to mobilize in preparation for play of the game. The model yields a set of hypotheses about how exposure to the CDR intervention might have had an effect. Section 5 then considers these hypotheses using our evidence from surveys of the game players, community representatives, and chiefs. The conclusion summarizes our results and expands on policy implications.

2 Community-Driven Reconstruction in Northern Liberia

Brutal and inept rule by Sergeant Samuel K. Doe in the 1980s had, by the end of the decade, attracted a variety of armed challengers for power in Liberia, the most successful of whom was Charles Taylor. Taylor's National Patriotic Front of Liberia (NPFL) fought the government,

a rival rebel group (ULIMO), and international intervention forces (ECOMOG) until a peace agreement brought him to power by elections in 1996. Something of a peace then prevailed until new rebel groups, and chiefly the Liberians United for Reconciliation and Democracy (LURD; in some ways a successor to ULIMO), began trying to unseat Taylor in 1999. With the help of international pressure, they succeeded in 2003. A major United Nations peacekeeping operation (UNMIL) and the election of Ellen Sirleaf Johnson as president, followed. In 2006, the United Kingdom's Department for International Development (DFID) funded a $1.6 million project by the IRC for a CDR project in two districts of northern Liberia, Voinjama and Zorzor.

2.1 The CDR Program

The IRC project sought to support CDR programs in 42 "communities," where "communities" were constructed as groupings of a relatively large "hub village" (which in practice ranged from roughly 30 to 600 households) and smaller neighboring "satellite" villages (usually small clusters of households).[4] The authors participated with the IRC in the identification of a set of 83 potential communities, and in designing public lotteries held in September 2006 to select 42 to receive an IRC CDR program.

Voinjama and Zorzor districts were hard hit by the war, due in large part to several episodes when control of the area shifted between Taylor's NPFL and the rebel groups ULIMO and LURD. People of the majority ethnic group in the two districts, the Loma, mainly supported Taylor and the NPFL; members of the largest minority, the Mandingo mainly supported, or at least were identified with, ULIMO and LURD. Based on our household surveys, Voinjama district is about 59% Loma and 30% Mandingo; in Zorzor, the proportions are 92% and 4%. Most communities, however, are relatively ethnically homogeneous; for example, about 50 of the 83 communities are 90% or more from one group, and in only 16 is there a minority

[4]In the region, these settlements are generally called towns, not villages. The size and number of the CDR communities was determined in part by funder requirements on the number of people to be served by the project, in part by the logistical capacity of IRC in northern Liberia, and in part by distribution of villages and people in these two districts. In practice, our impression is that the IRC projects focused on the hub villages (and town quarters), which are natural communities in the sense that they have a traditional authority structure.

population of at least 15%.

Our baseline data record information on almost 6,000 household members living in the region in 1989. Of these, over 4% are reported to have died directly from war related violence and a further 6% suffered injury or maiming. 5% took active part in the fighting, with three fifths of these reporting that they were abducted. A similar share (4.9%) of approximately 1500 subjects we interviewed in our follow-up survey self-reported as ex-combatants. The most widespread impact, however, was one that could bear on communities' ability to cooperate (Richards, 2005): 85% of these individuals were displaced during the conflict and many were displaced multiple times, often to refugee camps in Guinea.

The IRC's CDR program adapts for a post-conflict context the Community-Driven Development model that is now widely supported by the World Bank and other donors for aid programs aimed at poverty reduction. The goals are to "improve material welfare, build institutions and promote community cohesion ... [and to facilitate] the creation of sustainable community (governance) structures and communities participating within those structures through a system responsive to community rights and needs – paying particular attention to the most vulnerable and those most impacted by war (women, youth, excombatants and vulnerables)."[5] A premise of the project is that past conflict increases tensions and distrust within communities, thus creating a need for interventions that will promote reconciliation. Implicit in the design are hypotheses about the impact of CDR on: (i) material well-being, through investments that improve access to services and greater opportunities for income generation (ii) local governance, through process interventions that change attitudes and patterns of community participation; and (iii) community cohesion, through experiences that facilitate a reduction in tensions and a greater ability to organize. In this paper, we focus primarily on community cohesion and in particular the effects of the intervention on capacity for collective action.

The program in Liberia had the following core components (for convenience Table 6 in the web appendix provides a timeline summarizing the major steps from baseline survey to

[5]This is from the IRC's final proposal to DFID for the project. See Mansuri and Rao (2012) for an extended presentation and analysis of the philosophy behind CDD, and for a systematic review of evidence on its effects to date.

treatment to measurement). After the communities were identified and the treatment set selected in September 2006, the IRC undertook initial activities to explain the program to local communities, including meetings with chiefs and elders to solicit their cooperation on an advisory board. In each community, the IRC then oversaw the election of community development councils (CDCs), with 5 to 15 representatives (the average was 9). All adults in the community were eligible to vote, and the IRC staff encouraged though did not require the CDC to include female members (in practice, all communities had at least one and in the median case, one-third of members were women). CDCs were then empowered to oversee a community-wide process to select and implement a "quick impact" project (median value of $2,700), followed by a larger development project (median value of $12,000). Communities were also encouraged to consider using part of the total block grant (median value of about $13,000) for a "marginalized project" intended to address needs of vulnerable groups, although in practice these projects, when undertaken, were similar to the quick-impact and larger projects.

All three types of projects tended to involve construction of community facilities, such as community meeting houses and guest houses (approximately 35%), latrines (30%), and hand dug wells (15%). Very few projects (less than 5%) focused on school or health clinic construction, and almost none in agriculture, skills training and small business development, and other income-generating activities. The IRC staff helped to conduct a needs assessment with the CDC and in community meetings, but, subject to a few constraints, the "community driven" philosophy deliberately leaves project selection to the community.[6] For all projects, communities were supposed to supply labor or in-kind contributions worth 10% of project value. IRC staff also assisted the CDCs with project design and tendering bids from local contractors. CDCs managed the implementation process and continue to have responsibility for project maintenance over time.

By March 2008, construction had been completed on 55 of 131 projects in all 42 treatment

[6]Projects must be for community-wide rather than private or narrowly targeted benefit, and it seems that purchase of capital equipment for income-generating projects (such as a rice mill) was also not allowed in this case.

communities, and had at least begun on almost all of them.[7] Delays were ascribed mainly to an initial overestimation of the capacity of the local construction sector, and perhaps also to to what may have been an unusual level of IRC staff turnover.

2.2 Treatment Assignment and Covariate Balance

In September 2006, IRC staff randomly assigned collections of villages to treatment. The method used was block randomization, with 21 of 40 clusters of villages selected with equal probability in Zorzor and 21 of 43 communities selected with equal probability in Voinjama. Selection was implemented by IRC staff by drawing lots during public lotteries with participants from the community clusters. Reports from the IRC suggest that representatives of communities were in general positive about the process of random allocation on the grounds that it seemed both transparent and fair relative to the standard approach of selection by NGOs and government officials.

In March and April 2006, before the community boundaries were decided, we implemented a baseline household survey of 1,702 households in the two districts. Of these, 1,606 lived in communities ultimately assigned to treatment or control status. The baseline data allow us to assess whether the treatment and control communities are similar on various dimensions such as material well being, conflict experience, ethnic composition, as well as a large set of indicators of attitudes about governance.

In online Appendix B, we provide the distribution in treatment and control communities of a core set of variables that are plausibly associated with collective action capacity: basic population data (number of households, persons per household), a set of three wealth indicators (two composite measures of material wellbeing and percent with primary school education), exposure to conflict (percent household members injured or killed in conflict since 1989 and share that are former combatants), a measure of ethnic heterogeneity (percent Mandingo), and a measure of rurality (percent of communities that are "quarters" of a larger town). With one exception, balance is very good; our many attitudinal indicators show excellent balance as

[7]Note that some communities, particularly the larger ones, which received larger block grants, pursued multiple projects.

well. An F-test for the hypothesis that these variables are jointly uncorrelated with treatment has an associated p value of 0.56 indicating that we cannot reject the null hypothesis that the randomization was faithfully implemented by IRC field agents.

The variable for which balance is poor is "quarters." Twenty-eight communities in the five largest towns, are classed as quarters — an administrative level within a town that has a chief or sub-chief, and more or less well-delimited boundaries. Chance allocated 10 quarters to CDR treatment and 18 to control, a somewhat skewed distribution. This difference corresponds to a 0.40 standard deviation difference between treatment and control group values. In the contribution game discussed below, we found that the quarters generated markedly lower contributions, an outcome consistent with other observations suggesting that these communities were less well organized on average than more rural communities. There is disagreement on the merits of trying to "control" for variables on which there is imbalance of this form. Introducing controls does not reduce bias since bias does not depend on the realization of the randomization. Moreover, it may *introduce* bias if controls are selected using a 'conservative' approach in which controls are introduced precisely because they lead to smaller estimated effect sizes or larger standard errors. Introducing controls may improve efficiency, although the efficiency rationale for introducing controls is weakened, not strengthened, by the failure of a balance test (Mutz and Pemantle, 2011). Nevertheless, many researchers view invariance of estimated effects to the introduction of controls as evidence of robustness (cites). For this reason, in most analyses we report results with and without a control for "quarters" although we emphasize that our preferred specification, the unconditional estimate of treatment effects, provides an unbiased estimate.

2.3 Estimation of Effects

Unless otherwise noted, we report estimates of the average treatment effect. These estimates take account of the blocked randomization by using district as strata; in addition, strata are used to account for other treatment arms where relevant and, where noted, to account for 'quarters' as a potential confound. Estimates of standard treatment effects and their

associated standard errors are calculated as follows:

$$\hat{\tau}_{ATE} = \sum_{S \in \mathcal{S}} \frac{n_S}{n} \left(\frac{1}{n_{S1}} \sum_{S \cap T} y_i - \frac{1}{n_{S0}} \sum_{S \cap C} y_i \right) \tag{1}$$

$$\hat{\sigma}_{ATE} = \sqrt{\sum_{S \in \mathcal{S}} \left(\frac{n_S}{n} \right)^2 \left(\frac{\hat{\sigma}^2_{S1}}{n_{S1}} + \frac{\hat{\sigma}^2_{S0}}{n_{S0}} \right)} \tag{2}$$

where y_i is the observed outcome of interest in unit i; $\mathcal{S}$ is a set of strata with typical element S; T and C are the collections of units in treatment and control; n and and n_S denote the number of all units and units in stratum S respectively; n_{S1} and n_{S0} denote the number of treated and untreated units in stratum S; and $\hat{\sigma}^2_{Sj}$ is the estimated variance of potential outcomes under treatment condition j in stratum S. All analyses of CDR effects use the community as the unit of analysis since this was the level of treatment assignment (or in the analysis of heterogeneous effects, subsets of community responses are analyzed); in our analysis of the behavioral game, we also analyze a composition treatment at the level of communities and interest rate effects at the level of individual players.

Exact p values are estimated using randomization inference (Gerber and Green (2012)) and taking account of the structure of blocking in the randomization scheme. In general, results for the average treatment effect, or a t-test on the difference of means without matching, yield very similar results.

2.4 Multiple Comparisons

In the sections analyzing mechanisms, we often have many outcomes of interest based on responses to multiple related survey questions. This multiplicity of possible outcome measures gives rise to a well-known problem. With so many questions, an item-by-item analysis will find some differences between treatment and control groups to be "statistically significant" even if the null hypothesis of no impact is true.[8]

[8]Our preliminary analysis proceeded item by item, noting greater item-by-item "significance" for some groups of questions, and also a general pattern of positive CDR treatment impact that was unlikely to be explained by chance even though for most individual questions the CDR effect was not "statistically significant" (Fearon, Humphreys, and Weinstein, 2009).

When we have multiple measures for a construct, we address this problem following the approach of Kling, Liebman, and Katz (2007) and create a set of standardized indices of outcomes on related items. Within each set of variables, we first define items so that higher values imply a positive treatment effect, we then subtract the mean for the control group and divide by the control group standard deviation. The index is then constructed as the standardized average of the standardized variables for each community. (For details, see online Appendix C.)

3 CDR Impact on Collective Action Capacity

Although panel surveys of community members in treatment and control communities can be used to assess whether and how the CDR program changed self-reported attitudes and opinions about community governance and institutional performance, we were concerned that CDR could lead to a change in reported responses without changing capacity or inclination for collective action. For example, the NGO's intervention might influence people's understanding of what they are "supposed to say" but not their willingness or ability to act and coordinate in line with those expressed beliefs. For this reason, we designed a behavioral measurement strategy–a matching funds experiment–in which communities were confronted with a real-world problem of raising funds for a small-scale development project.

Starting about four months after the formal completion of the IRC project, an advance team visited each of the 83 hub towns and gained consent for a community meeting to describe a new opportunity for the community to receive funds for development. One week later, we ran a meeting in which community members were told that they could receive up to $420 to spend on a development project.[9] Receipt of funds would depend on whether the community completed a form indicating how the funds would be spent and the names of three community

[9]The initial community meeting and the game itself were administered by a Liberian NGO – National Excombatant Peacebuilding Initiative – working with a team of Stanford graduate students under our oversight. NEPI members did not know that we were studying the effects of the IRC CDR program, and our graduate students typically did not know which villages were treatment and control (although in some cases signs advertised IRC projects). Of course, the communities themselves did not know that there was any connection to the CDR program.

representatives to receive and handle the funds. The specific amount received would depend on the private contribution decisions of a random sample of 24 adults who would be given about $5 each by us – the more these individuals contributed, the more we would match their anonymous contributions at a public meeting, after which the total amount raised would be handed over to the three community representatives.

One week after this protocol was explained at community meeting, a team returned to the village, collected the form, sampled 24 households, played the contribution game, and publicly announced and provided the total payout to the village. Between these two visits, the community had time to select their community representatives and potential projects, and to spread information about the game and how it should be played. On game day, detailed surveys were completed with all 24 game players (after they made their private contribution decisions), the three community representatives, and the village chief.

3.1 Game Description

The contribution game itself was straightforward. Twenty-four randomly selected adults – from households selected at random from the entire village – were given three 100LD notes, worth in total about $5 US or close to a week's wages. They then chose, in private, how much to contribute to the community and how much to keep for themselves. It had been explained in the community meeting that half of the players would have their contributions multiplied by two, while the others would be multiplied by five, corresponding to interest rates of 100% and 400%. Thus each community had the opportunity to earn up to 25,200 LD. For this "interest rate treatment," players were randomly assigned to the high and low rate conditions (with blocking on gender and location). Players knew their interest rate when choosing how much to contribute.[10]

To be clear, note that the 24 game players were selected from the entire village, not from the set of people who attended the community meeting a week earlier. Attendance at that initial meeting varied greatly, averaging around one quarter of the adult population of the

[10]In the community meetings, our presenters stressed that the contribution decision was up to the game player and that there could be valid reasons to keep the money for private use. It was evident, however, that attendees immediately grasped the conflict between private and social good.

village, with the percentage varying negatively with village size. The village chief was almost always in attendance and assistant chiefs and elders were always present. In the week between the community meeting and "game day" the village had the opportunity to mobilize to inform members not present at the meeting about the project, and to meet (if they chose to) to decide who the community representatives would be and what to do with the money raised.

3.2 Gender Composition Treatment

In addition, we ran a cross-cutting experimental treatment. In half of the communities all 24 game players were women, while in the other half we selected 12 men and 12 women players. We implemented the gender composition assignment using a matched pair design in which units were matched based on estimated population size. This 'gender composition' treatment related only to the makeup of the players for the contribution game, and not necessarily the set of beneficiaries of the potential development project. In verbal instructions delivered at each community meeting it was made clear that, regardless of the gender of the game players, in all communities both women and men could participate in meetings to decide on projects, serve as community representatives, and be beneficiaries of the project itself.[11]

Table 1 gives the overall distribution of treatments and reports the number of communities and treatments in each condition.

3.3 Implementation of Games

Eighty-two communities successfully completed the behavioral game.[12] The average payout to villages was 20,020LD, or 79.4% of the total possible, with a standard deviation of 13.3%. Among individuals, fully two-thirds contributed the maximum amount (300LD), with the rest

[11]We added the gender treatment because it was simple to implement and could allow us to speak to donor and academic interest in the question of how women's involvement in community decision-making affects outcomes; we also expected, as noted above, that we would find no CDR treatment effect. We did not have the statistical power to have a third set of communities where all game players were men. Because our focus in this paper is on CDR, we do not provide a full analysis of the gender composition treatment.

[12]Play was halted prematurely in one community after a player changed her mind about her contribution decision and a public scene developed when she and her sister made this known. The community was later given an approximately average payout to avoid hard feelings.

Table 1: Distribution of Treatments

Gender composition	CDR Intervention Control	CDR Intervention Treatment	Total Communities (Participants)
Mixed groups (12 men, 12 women)	20	22	42 (1008)
Women only (24 women)	21	20	41 (984)
Total communities (Participants)	41 (982)	42 (1008)	83 (1992)

Notes: In all communities, 12 players were randomly assigned to have a high interest rate and 12 to low. In areas with mixed groups half the men and half the women were assigned to each interest rate condition. The CDR assignment was blocked on district. The gender composition assignment was blocked using a matched pair design with matching on village size.

almost evenly divided over giving 200 (10%), 100 (12%), or 0 (11%). The average contribution was about 235LD, which is 78.3% of 300.[13]

The contribution game has the structure of a public goods game, at least to the extent that players expected that funds raised would be spent on community projects they viewed as beneficial. Given that the communities had a week to mobilize and exhort individuals to play for community benefit and given the novelty of the situation, it is difficult to say what one should have expected in terms of average contributions. Arguably, though, contribution levels of almost 80% of the total possible represent an impressive amount of cooperation.[14]

3.4 Main Effects

Table 2 presents estimates of the main effects of our three randomized treatments on contributions in the game. The upper half of the table shows that all three treatments had substantial impact in terms of dollars contributed. For the share of total payout CDR raised the average contributed from 75.8% to 82.1% of the maximum possible, a difference that is close to half the standard deviation of the payouts in our sample. Another way to scale the magnitude of

[13]A handful of players disobeyed instructions and put amounts other than 0, 100, 200, or 300 into the envelope. We use their actual contributions in the individual-level analyses that follow.

[14]In lab experiments, contributions in the first play of analogous public goods games are typically around 50% of individual endowments (Ledyard, 1995). Lab experiments usually involve smaller stakes (relative to wealth and income) and smaller groups, both of which favor contributions relative to our case. On the other hand, our game involved actual communities that had a week to mobilize and exhort people to contribute if chosen to play.

the CDR effect is to note that it is about the same as the impact of raising the rate of return on individual contributions from 100 to 400%: both increased average individual contributions by about 17LD, or 5.6% of 300. The gender composition treatment ("all women") had a somewhat larger positive impact, with payouts at 84% of the maximum on average versus 75% in the mixed gender communities. The lower half of the table shows estimates of average treatment effects when quarter is included in the list of strata. We see that this substantially reduces the estimated average CDR treatment effect.

Table 2: Experimental Effects: Average Contributions (in Liberian dollars)

	CDR Treatment	Gender Composition Treatment	Interest Rate Treatment
Level in control group	226	223	226
Average treatment effect	17	27	18
p-value	0.03	0.003	0
Conditioning on quarter:			
Level in control group	231	224	226
Average treatment effect	8	23	18
p-value	0.3	0.002	0

Note: CDR and gender composition treatment effects, in Liberian dollars, are estimated at the community level; district and the other treatment form strata. For the interest rate, strata are formed by community gender groups. The p values are calculated using two tailed tests and randomization inference.

Tables 3 and 4 show the heterogeneity of the measured CDR treatment effect across strata. Table 3 compares the CDR effect in the mixed groups versus the communities where only women played the game. We find a striking difference. In the "all women" communities the estimated impact of CDR on contributions is statistically insignificant (and marginally negative). By contrast, in the mixed groups the CDR program is estimated to have increased contributions by 43 LD on average, a very large and highly statistically significant effect that is two and half times the estimate for all communities taken together (17LD). When we condition on quarter, the estimated impact in the mixed groups declines somewhat, to about 28 LD, but this remains a large and strongly statistically significant difference. Substantively, 28 LD is about 70% of the standard deviation of average contributions across all communities. Thus,

the positive estimated CDR effect in the mixed communities is *not* the result of lack of balance on quarters.

Why did the measured impact of the CDR program depend so strongly on the gender composition of the game players? Table 4 also gives contribution levels by player gender in the mixed groups. We see that men and women in mixed groups gave similar amounts in the control (non-CDR) communities, and responded to the CDR treatment in roughly the same way. By contrast, in the all-women communities women contributed at a high level with or without the CDR treatment. Thus, the lack of a CDR effect in the all-women communities is *not* explained as the result of CDR having a direct effect on men but not women.

The difference in CDR impact across mixed and all-women groups is also probably not explained by a ceiling effect. For some reason (to be discussed below), women contributed substantially more than men, closer to the maximum, *when they knew they were playing with other women*. So one might conjecture that we do not observe a CDR impact in the all-women communities simply because it was hard to drive contributions any higher. However, if contributions in the all women groups were close to a ceiling, then we would expect the effect of a higher interest rate in the all women groups to be lower than in the mixed groups, and this is not the case, as shown in Table 4.

In the next sections, we return to the puzzle of differential CDR impact across the gender composition treatments, a puzzle we cannot definitively resolve but that nonetheless proves useful for evaluating hypotheses about causal mechanisms.

Table 4 shows the impact of the interest rate manipulation on contributions in different strata. Here we do find a striking gender difference at the individual level: Interest rate effects are strong in both the mixed and women-only groups but surprisingly they appear only for women. Male players ignore the interest rate in both treatment and control conditions. Notice also that there is only weak evidence of an interaction between the interest rate and the CDR treatment.[15] We make use of this observation in sorting between mechanisms below.

Table 4 about here

[15]The interaction effect is positive but not close to significant (not shown).

Table 3: CDR effect by gender composition and player gender

	Mixed communities			Women only		
Group	Control level	CDR Effect	p value	Control level	CDR Effect	p value
All	200	43	0.000	253	-11	0.205
Women	195	48	0.005			
Men	204	41	0.001			

	Mixed communities			Women only		
Group	Control|Q	CDR Effect|Q	p value	Control|Q	CDR Effect|Q	p value
All	209	28	0.008	254	-14	0.135
Women	205	29	0.054			
Men	211	28	0.013			

Notes: The top table gives the CDR treatment effect on individual contributions in Liberian dollars; the bottom table is the same but conditions on Quarter. p values are from randomization inference.

Table 4: Interest rate effect by CDR status, gender composition, and player gender

	Mixed communities		Women only	
Group	High interest effect	p value	High interest effect	p value
All	11.2	0.076	24.5	0.000
CDR	15.5	0.057	32.1	0.001
no CDR	6.5	0.500	17.6	0.034
Women	21.1	0.021		
CDR	21.9	0.080		
no CDR	20.2	0.122		
Men	1.3	0.888		
CDR	9	0.461		
no CDR	-7.1	0.624		

Notes: Effect on average individual contributions (in Liberian dollars) of being assigned to the high interest versus the low interest group. p values are from randomization inference.

Despite statistical objections to the practice (Freedman, 2008), multiple regression is often used to check whether a treatment effect remains when covariates are considered. This is indeed the case for the CDR effect in the mixed gender groups when we 'control for' community size, whether the community is a quarter in a larger town, and percent Mandingo. Larger communities generated significantly lower payouts in the game, as did the more urban quarters

and predominantly Mandingo communities. The estimated CDR effect in the mixed gender groups diminishes some when we control for these factors (and especially quarter, which was not well balanced), but remains substantively and statistically significant (see Table 11, Web Appendix F).

4 Mechanisms

Like other CDD and CDR programs, the CDR program in northern Liberia was a short-lived NGO effort to improve communities' collective action capacity by introducing elected development councils and providing funds for community-chosen projects. Given much evidence suggesting that local-level social and political institutions are highly persistent, one would probably not expect a significant behavioral impact. Nonetheless, we found strong evidence that the CDR program substantially increased collective action in the communities where both men and women engaged in the contribution game. By what pathway or pathways did the CDR program change collective behavior in these communities?

This question cannot be answered definitively because of the difficulties randomly assigning many variations within a complex treatment in order to identify mechanisms. We can, however, use our three randomized manipulations, together with the surveys of game players, community representatives, and chiefs, to draw inferences about what mechanisms are more or less likely to have been important. We argue that such efforts are essential if experimental analyses of complex governance interventions are to produce results that can have broader social science and policy relevance (see also Acemoglu (2010) and Ludwig, Kling, and Mullainathan (2011)).

It is important to note at the outset that our question is what explains the CDR effect on contributions, and *not* the question of why did individuals contribute in the game in the first place. Individuals may have contributed because they put a high value on projects proposed; out of a sense of wanting to do the right thing to "bring development to their community"; from fear that despite our precautions their contribution decisions could be discovered; or due to a desire to please well-heeled foreign donors who might be expected to bring more

funds later. Our question instead is about the impact of the CDR program on the level of contributions.

The CDR program could have affected play in the public goods game by a number of causal paths. A simplified model of the contribution game is useful for distinguishing between possible mechanisms and for drawing out testable implications of particular mechanisms.

4.1 A Model

Consider a game with n players who simultaneously decide what share of a dollar to contribute for a community project. We let $y_i \in [0,1]$ denote the choice of individual i and $\hat{y}_{-i}$ the average contribution of others. We assume players may gain utility from contributions in three ways. First, they may gain value from money spent on the project. Second, they may see the contribution as either a cost or an obligation that they derive satisfaction (or avoidance of social sanctions) from fulfilling. Finally, they may value coordination and gain from contributing especially when others also contribute.[16] We allow for these possibilities in a simple way by representing individual preferences using a function with three components:

$$u_i(y) = \mu \left(\sum_{j \neq i} r y_j + r y_i \right) - \theta y_i + 2\kappa \sqrt{(y_i + 1)\hat{y}_{-i}} \tag{3}$$

The first term on the right hand side refers to the benefits from spending on the project. The term in large parentheses is the total amount contributed, where interest rates are given by r.[17] The key parameter of interest, $\mu \in \mathbb{R}$ captures the rate of transformation of money raised into public goods. The second term captures net costs or benefits associated with contributing, *independent* of the amount raised and spent on the project. The parameter of interest here, θ, may be positive or negative. The third term captures gains from coordination; here gains from one's own contribution are increasing in the contributions of others at a rate

[16] Such coordination preferences could reflect increasing marginal returns from contributions to the project, or possibly decreasing marginal costs of contributing when others also contribute, or they may reflect social preferences such as discomfort for deviating from what most others are doing or the presence of prefernces reflecting "strong reciprocity" (Bowles and Gintis, 2004).

[17] Or more precisely r denotes the multiplier which is equal to one plus the interest rate. In the real game these are player specific though we ignore that feature here.

governed by κ.

The first derivative of the individual's utility is given by:

$$\frac{\partial u}{\partial y_i} = \mu r - \theta + \kappa \sqrt{\frac{\hat{y}_{-i}}{y_i + 1}} \tag{4}$$

From this expression it is easy to see that individuals have a dominant strategy to contribute nothing if $\theta > \mu r + \kappa$. They have a dominant strategy to contribute everything if $\theta < \mu r$. A coordination dilemma arises if $\theta \in \{\mu r, \mu r + \kappa\}$. In this case equilibria exist in which all individuals would contribute everything if others did, and in which all would contribute nothing if others did not.

From these results, we see that gains in μ relative to θ induce shifts from a situation where nobody contributes to one where all do. Moreover, gains from μ are amplified by larger interest rates. Gains in κ can turn a game with dominant strategies not to contribute into a coordination dilemma.

Consider now the case of a coordination dilemma and assume that each person holds beliefs regarding how much others are likely to contribute.[18] Say in particular that a given individual has a belief that it is equally likely that either share $\lambda + \sigma$ will contribute or share $\lambda - \sigma$ will contribute. Assume $\lambda \in (0, 1)$ and note that admissible values of σ are constrained by λ.[19]

In this uncertain world, the individual maximizes:

$$u_i(y) = \mu \left(\sum_{j \neq i} r y_j + r y_i \right) - \theta y_i + \kappa \left(\sqrt{(y_i + 1)(\lambda + \sigma)} + \sqrt{(y_i + 1)(\lambda - \sigma)} \right) \tag{5}$$

This maximization problem is solved by:

[18]Note that in the cases with dominant strategies beliefs are determined given knowledge of the preferences and rationality of other players; knowledge of rationality and preferences do not however pin down beliefs in the case of a coordination dilemma.

[19]Of course if all players knew that all others believed that a subset would cooperate and if all believed others would choose to cooperate (or not) given these beliefs, then, generally, they should not in fact expect that only a subset would cooperate (Aumann, 1976). Here we do not assume that beliefs are common knowledge and, rather than focusing on equilibrium behavior, focus on choices by individuals conditional on possibly out-of-equilibrium beliefs.

$$y_i = \left(\frac{\kappa}{\theta - \mu r}\right)^2 \frac{\lambda + \sqrt{\lambda^2 - \sigma^2}}{4} - 1 \tag{6}$$

From this expression, we see again that in the coordination case, contributions are decreasing in θ and increasing in κ and μ. In addition, they are increasing in expectations regarding the contributions of others λ and *decreasing* in the uncertainty regarding those contributions, σ.

In the discussion so far we have focused on the direct effects of the intervention on parameters in the model. But the CDR program might also have indirectly affected individuals' disposition to contribute by increasing mobilization and information-sharing efforts of community leaders or community members as a whole.

In the terms of our model, suppose that mobilization effort by community elites could affect any of the parameters that determine individuals' disposition to contribute. Let $\gamma = (\mu, -\theta, \kappa, \lambda, -\sigma)$ denote these parameters (or the negative of these parameters for parameters that have a negative effect on contributions). Under the assumption that strategies are implemented with some error, we let $Y(\gamma)$ denote the expected total contributions given γ which we take to be smooth and increasing in each element of γ. To represent leadership investments in mobilization, say that there exists a baseline (no mobilization) value of the parameters, $\underline{\gamma}$, and that, at a cost, leaders can exert effort to choose $\gamma \geq \underline{\gamma}$ to maximize

$$u(\gamma) = Y(\gamma) - c(\gamma, \alpha|\underline{\gamma})$$

Here $c(\cdot)$ denotes a smooth convex cost function that is increasing in each element of γ; α denotes *organizational capacity* and we assume that marginal costs of organizing are decreasing in capacity in the sense that the cross partial $c_{\gamma\alpha}$ is negative. Leaders might want to maximize contributions net of effort costs either for the sake of the community or because they appropriate the money for themselves. It is then easy to show that the elites' optimal level of mobilization increases, as do total contributions, as mobilization capacity α, increases. [20]

[20] Let γ belong to a compact, convex parameter space. A solution to the problem exists because a

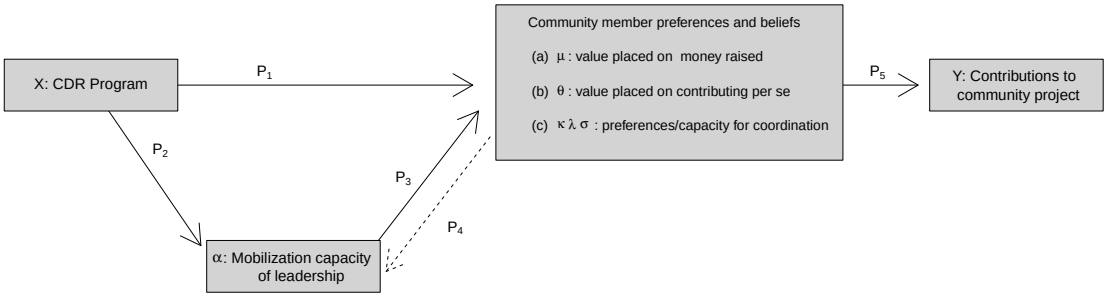

Figure 1: Summary of mechanisms

Figure 1 about here

In summary, this model focuses attention on six parameters of interest: the benefits of increased contributions (μ); the net costs of contributing per se (θ), preferences for coordination (κ), beliefs about the likely actions of others (λ, σ), and the organizational capacity of leaders (α). The CDR program may have operated by influencing any of these parameters; moreoever the effects of the program may have worked through either direct of indirect channels.

Figure 1 provides an illustration of the relationship between direct and indirect effects implied by the model. As shown in the figure, the CDR program might have directly affected the beliefs and preferences of community members in such a way as to increase their contributions (P_1-P_5). Or CDR might have affected the ability or willingness of community leaders to mobilize communities and alter their behavior (P_2-P_3-P_5).

It is also possible that CDR increased mobilization in part by changing some individual preferences and beliefs – whether leaders' or regular community members' – leading them to work harder at mobilization (the path P_1-P_4-P_3-P_5). For example, the more of the money

continuous function has a maximum on a compact set; it is unique by concavity of $u(\gamma)$. The (positive) incremental expected contributions attributable to mobilization is given by:

$$\Delta(\alpha, \underline{\gamma}) \equiv Y(\gamma^*(\alpha, \underline{\gamma})) - Y(\underline{\gamma})$$

We have then that a drop in the costs of mobilization, α, weakly increases total contributions, but may do so through changes in any of the parameters affecting individual preferences over the contribution decision. From the implicit function theorem, $\gamma^*_\alpha = -c_{\gamma\alpha}/c_{\gamma\gamma} > 0$.

that the elites can appropriate for themselves (smaller μ), the less incentive for individuals to contribute, but the greater the incentive for elites to mobilize and try to induce contributions.

4.2 Measurement of model parameters

With this simple characterization of the problem in hand, we turn to an empirical assessment of the channels of interest. In doing so, we note that the random assignment of the CDR program allows us to identify the effect of CDR on mobilization activity (P_2, or perhaps P_2 and P_1-P_4), as well as the total effect of CDR on our measures of community member preferences and beliefs (P_1 plus P_2-P_3).[21] We can also draw on partial correlations between our measures of mobilization, game-player preferences, beliefs, and contributions to provide suggestive evidence on some of the other paths. Finally, we can use results from our other randomizations, and in particular the fact that CDR "worked" in the communities where both men and women played the contribution game but not when only women played, to see which paths are mostly likely to be in operation.

In the following sections, we describe the ways that the CDR program might have affected each of these model parameters and describe measurement strategies to capture these effects.

4.2.1 Value for the public good: μ

There are multiple ways in which the CDR program might have affected the value individuals place on the public good produced by collective contributions. We highlight three.

Project selection effects. First, if more democratic methods were used to select projects for our game in CDR communities, this may have yielded projects more highly valued by the average community member and thus increased contributions. In the game player surveys, we asked respondents to rate how important they thought the projects selected were; whether they liked, didn't like, or were indifferent to the projects chosen; how many people had "different views" about which project should be chosen; and whether they thought most people in the community would benefit from the projects chosen. We constructed a mean effects index based

[21] Of course, we can also identify the total effect of CDR on contributions through all paths (P_1-P_4 plus P_2-P_3-P_4), which we presented in section 3 above.

on these questions to estimate individual and community level values for the projects.

Trust in leaders. Second, if CDR had the effect of reducing the scope for graft – for example, by rendering leaders more accountable to citizens – then game players could be more confident that their contributions would indeed go to the proposed project. This pathway is particularly important given the stress that donors put on improving democratic accountability as the core mechanism by which CDD and CDR programs are supposed to have good effects on governance and community well-being. Our game player survey asked how trustworthy respondents thought "community leaders" were in "this town" and relative to other towns in Lofa County; what share community leaders would keep for private use if they got access to funds intended for the community; whether the three community representatives would use the money raised to implement the project chosen; whether the chief and representatives would benefit more than others from the projects; and whether a concern that the money would be mishandled was a factor in the respondent's contribution decision. These seven questions were used to construct a mean effects index of community trust in leadership.

Trust in NGOs. Third, the experience of the CDR program with the IRC could have increased trust that we would actually match contributed funds as stated and hand the total over to the community representatives. In the community surveys carried out immediately after completion of the IRC program (March-April 2009), we asked respondents about twelve actions they might take "to try to change the situation ... if you had concerns about how things were going in your village." Two of the actions offered were "appeal to local NGOs for assistance" and "appeal to the international community for assistance," with possible responses being that this would "make things worse," "make no difference," "help a little," or "help a lot." We also asked each respondent to select from the list of actions the ones they thought would be most effective and second most effective. These questions allow us to create a measure of trust and beliefs in the efficacy of NGOs and "the international community." Our composite index comprises questions about whether appeals to national or international NGOs would help and whether this action would be either most or second-most effective.

4.2.2 Value for contributing independent of public goods outcome: θ

CDR could also have direct effects on individuals' value for contributing independent of anticipated private or collective benefits from the money raised. These are captured by θ in the model.

Income effects: If the CDR program significantly increased incomes in the community, this could lower game players' value for cash versus public projects.[22] We asked a battery of questions in the follow-up and game player surveys about household income, assets, quality of housing materials, and access to water and land. We use these to construct a composite measures of individual and community material welfare.

Sanctioning effects: The CDR program may increase cooperation by increasing individuals' expectation that they might be sanctioned for failing to contribute. For example, CDR could in principle have established stronger norms of cooperation and thus increased expectation of disapproval for noncompliance. Or CDR might have increased information flows about behavior. Alternatively, it is possible that CDR actually weakened the capacity to sanction, for example by weakening traditional authority structures. Two survey questions about whether game players thought others would find out what they contributed and whether this concern affected their contribution decision were used to construct a composite measure of fear of discovery, a precondition for fear of sanctioning.

Legitimacy effects from participation: If CDR increased the use of participatory procedures in community decision-making, this could lower individuals' costs of contributing independent of their value for the project and trust in leaders, by increasing the perceived legitimacy of the action.[23] Note that in contrast to some of the other "direct effects" listed above, this one

[22] In principle, higher incomes could also reduce individual's demand for community infrastructure as well.

[23] Three recent studies support a logic of this form. In lab experiments, Dal Bó, Foster, and Putterman (2010) find that use of democratic procedures independently increased the effect of a given policy on cooperation in a public goods game (that is, use of elections themselves increased cooperation, rather than only via the policy chosen). Hamman, Weber, and Woon (2011) find that electoral delegation to a leader increased contributions despite moral hazard temptations facing the leader and Baldassarri and Grossman (2011), deploying lab experiments in rural Uganda, find that participation in elections to select third-party enforcers increased contributions to public goods without direct effects on the characteristics of leaders.

requires at least some mobilization activity by community leaders: in order for democratic methods to be used to make community decisions, meetings must be organized and held.

We asked the game players a number of questions about the process used to select the projects and community representatives: were meetings to make these decisions organized by the chief or by community members; were the community representatives selected by a vote; were they selected in a public place; were the projects selected by a vote; were they chosen in a public meeting.[24]

4.3 Coordination versus dominant strategy preferences: κ, λ, σ

The CDR program might have increased collective action capacity by facilitating coordination. Changes in κ would arise if the program altered the values individuals placed on "doing one's part" conditional on others doing so. But the program could also have increased the community's ability to coordinate given such preferences either by altering beliefs about the likely behavior of others in a coordination problem or uncertainty about those beliefs. Such changes could arise if, for example, CDR increased information flows and as result facilitated the spread of information and common expectations about play.

In the game player survey, we asked respondents how they expected other players to choose. This allows us to construct a measure of both expectations and the accuracy of those expectations, which we use to proxy an individual's uncertainty over the behavior of others. In addition this measure allows us to assess the correlation between an individual's expectation of others' play and his or her own contribution, a possible indication of coordination preferences.

4.4 Indirect effects (mobilization): α

Finally, the CDR program could have increased community contributions by creating or improving the effectiveness of a cadre of community members with experience in mobilizing the broader community for collective action. Such indirect effects could produce overall ef-

[24]We also asked if there was competition and disagreement over the community representative positions. These could be considered indicators of democratic process but probably also tap other dimensions like community cohesion. Results are the same whether or not these are included in the "democratic process" measure.

fects whether or not CDR directly affected the average disposition of community members to contribute.

We have multiple measures to assess these mobilization effects. First, the game player, community representative, and chief surveys contain items that tap different aspects of mobilization effort by community leadership. Questions about whether additional meetings were held to discuss the project, whether the respondent attended, and estimates of the number attending such meetings were used to construct an index called *meetings*. Questions about whether respondent had been personally contacted about the game, about the project, about staying home on game day (so as to have a better chance of being selected), and about whether anyone asked the respondent to contribute were used to create a measure tapping efforts at *contact*. A measure of *game player knowledge* was based on questions about whether the respondent had heard of the matching funds experiment, knew the projects selected by the community and the names of the community representatives, and could answer the questions about who organized meetings and how many had attended. We also asked chiefs and the community representatives about meetings held and efforts to contact individuals, and constructed parallel "elite" measures from these.[25]

5 Evidence on Mechanisms

5.1 Effects on community preferences and beliefs: $\mu, \theta, \kappa, \lambda, \sigma$

We begin by examining the impact of CDR on a variety of measures designed to proxy for the beliefs and preferences of community members. Random assignment allows us to identify CDR's total causal impact via its direct *and* indirect paths on our measures of game player's beliefs and preferences relevant to their contribution decisions.

We note first that the model suggests that if there are effects operating through valuation of the project, μ, we would expect these effects to be stronger when interest rates are higher (or conversely that interest rate effects would be greater in CDR areas). If CDR increases the

[25] As noted above, we took the approach of Kling, Liebman, and Katz (2007) to construct mean effects indices. Details are given in Appendix C.

value of every penny that goes to the public good then this effect is enhanced when pounds are in play.[26] The results given in table 3 provide only weak and non-significant support for this proposition.

Figure 2 shows CDR impacts on the intermediate measures we identified above, providing separate estimates for mixed and all-women communities, along with (in the lower panel) the correlation between these measures and average contributions in the development project experiment (P_5).

We find that CDR did not have an overall impact on outcomes associated with μ: satisfaction with projects, trust in community leaders, or trust in NGOs. This is consistent with our finding on the interest rate / CDR interaction.

The results on total effects on θ are more mixed. We see no effects on two measures: material welfare and perceptions of the anonymity of one's contribution in the game, whether in the mixed or in the all-women communities. This reinforces the inference that direct effects of CDR on community member preferences regarding the value of public goods or the value of contributing are unlikely to be the pathway by which CDR had an impact on collective action. It also suggests that it is unlikely that greater mobilization due to CDR in the mixed communities worked to increase contributions through any of these intermediate variables, since we do not see the pattern of positive effects in the mixed communities and no effects in the all-women communities.

By contrast, we do see evidence, in the mixed but not all-women communities, of a total effect of CDR on our measures of whether democratic process was employed to choose projects and representatives.[27] While the CDR program did not increase trust in leaders or project satisfaction as reported by the game players, it did, in the mixed communities, cause greater use of ostensibly democratic procedures like elections to choose projects and representatives. This feature may have led to greater support for the collective endeavor. Moreover this feature is also correlated with significantly higher average contributions in the contribution game.

[26] More formally we have that if contributions can be written $Y = f(\mu r)$ and benefits of treatment X operate through μ then $\partial Y/\partial \mu = r f'(\partial \mu/\partial X)$.

[27] CDR's total effect on democratic process is still positive and significant when we condition on quarter.

To assess whether CDR mattered by increasing the prevalence of "coordination preferences" (κ), such that people want to contribute if they think others will also, we used information from a question asking game players what they thought other players were likely to do. Their answers allow us to look at the correlation between their beliefs about others' behavior and their own chosen actions. Figure 3 shows the estimated marginal effect of higher reported beliefs about others' actions on one's own contribution decision. The effects are estimated using ordinary least squares and controlling for village fixed effects, so that not only are features such as quarter accounted for, but so is the 'truth' about village level contributions. Thus the results assess the relationship between an individual's contribution and the extent to which he or she had unusually high or low beliefs about the actions of others. We find a strong positive relationship, particularly amongst men, which is consistent with coordination preferences. But we see that the size of the correlation is similar in treatment and control communities. This suggests that a coordination logic may have been in operation but that this logic was not a product of exposure to CDR.

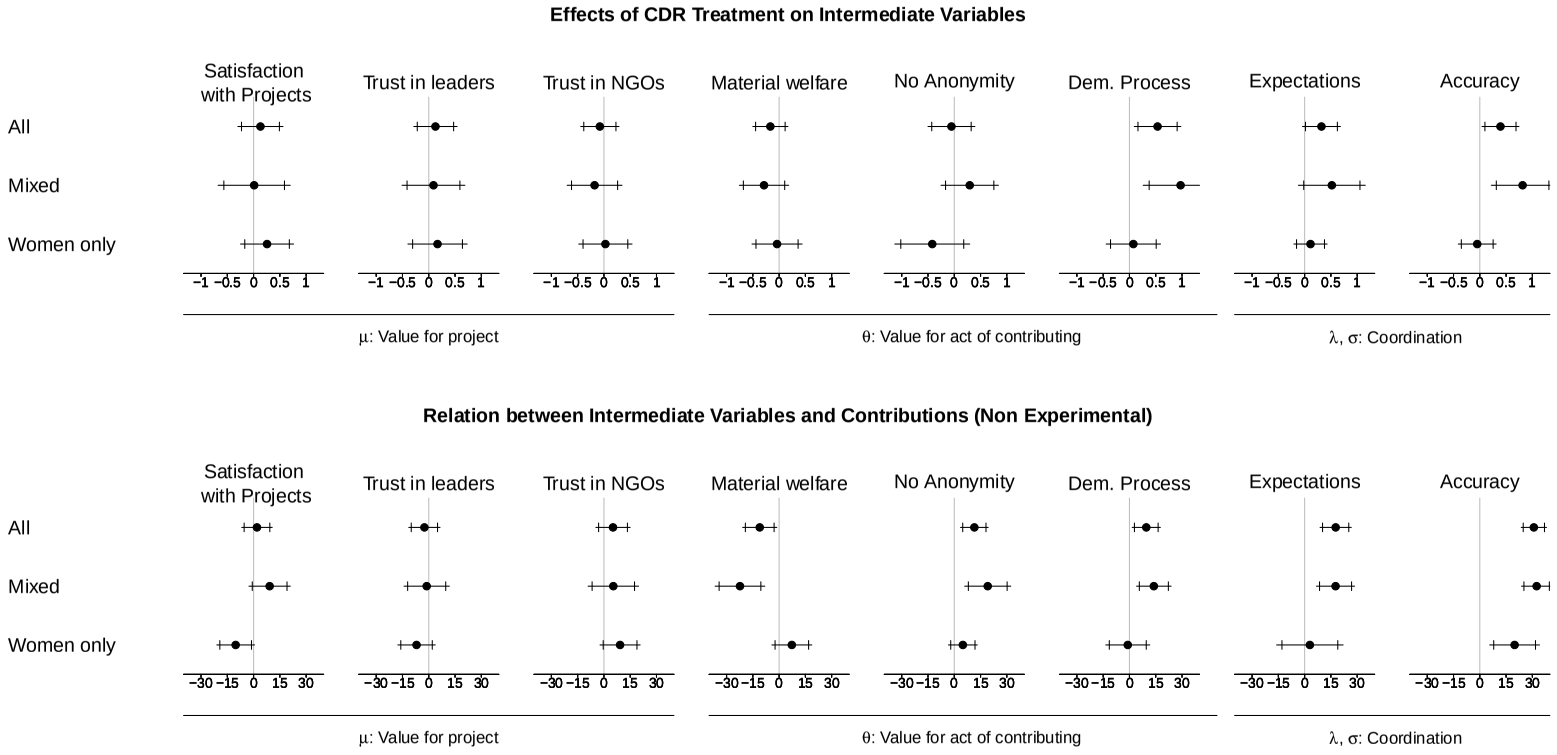

Figure 2: Evidence on Mechanisms. Upper panel shows the effects of CDR on various intermediate variables; lower panel shows the relation (regression coefficients) between these intermediate variables and outcomes. Note that while estimates in the upper panel are identified due to randomization, those in the lower panel are not. See Table 7 for variable definitions and summary statistics.

Figure 2 about here

There is more support for the idea that CDR *facilitated* coordination conditional on players having coordination preferences—that is that CDR effects may have operated through λ or σ. We constructed indicators of each player's beliefs about the contributions of other players as well as indicator of the accuracy of those beliefs. From the final two columns of Figure 2, we see evidence that in the mixed communities, CDR caused players to expect other players to contribute more. Expectations are also more closely associated with actual contributions at the village level in CDR areas for the mixed gender game, suggesting that players do have reasonable knowledge regarding the likely play of their peers on which they could condition their strategy if they wanted to. The former effect (on levels) is not robust to conditioning on quarter though the latter (accuracy) is.

This evidence is consistent with either CDR having a direct effect on community members' ability to coordinate, say, by increasing information sharing, or by increasing mobilization activity which got more people "on the same page." Given that we observe these effects only in the mixed communities and that under the simplest interpretation, one would expect a direct effect on information-sharing to affect both men and women, the it seems more likely that these total effects result from differential mobilization strategies. We turn to assessment of these mobilization strategies next.

Figure 3 about here

5.2 Direct versus indirect effects

The analysis of the last section focussed on total effects of CDR on preferences and beliefs. We now examine the case for CDR's impact via an indirect effect through mobilization rather than a direct effect on preferences and beliefs.

If CDR had an impact primarily via a large, direct effect on preferences and beliefs of both men and women in treated communities (P_1), then it is difficult to explain why we found, in

section 3, that CDR caused a large increase in contributions in the villages where both men and women played the game but had zero (or a slightly negative estimated) impact in villages where only women played.

On the other hand, if the CDR program had a large direct impact on the preferences and beliefs of *just men*, then it is hard – though perhaps not impossible – to explain why CDR caused a large increase in contributions by women in the mixed groups, but not in the communities where only women played. Hypothetically, it is possible that CDR directly affected men's preferences and beliefs so as to favor contributing, that women knew or anticipated this, and that they had coordination preferences and so wanted to match what they expected that men (and women) would do in the mixed communities. In the all-women communities, by contrast, perhaps contributions were expected to be higher regardless of CDR treatment because women were expected to be more community-minded than men, or because it was anticipated that they could more easily solve the collective action problem for some reason.[28] Player expectations about whether men or women would contribute more are broadly consistent with this latter hypothesis. We asked game players if they thought men or women would contribute more in the game, or about the same amount. 68% of women in the all-women communities thought women would give more than men (had men had a chance to play), compared to 52% of women in the mixed communities, a highly statistically significant difference. In addition, we find evidence consistent with players having coordination preferences.

[28]For example, in these communities it was known that there were half as many potential game players to inform and mobilize. Further, anecdotally, when it was announced that only women could be chosen to play the game at a community meeting, it often seemed as if the women in the audience took this as a challenge to demonstrate the community-mindedness of their gender.

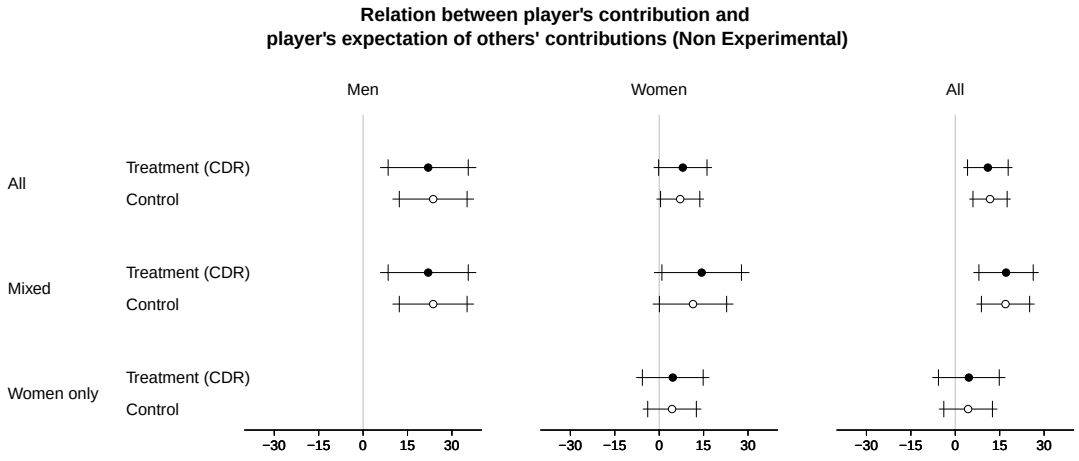

Figure 3: Regression estimates for player's beliefs about other's average contribution (scaled to 1 sd in control communities). Dependent variable is player's contribution in Liberian Dollars. Sample broken into 8 disjoint strata according to gender of respondent, CDR treatment status, and gender treatment status. All estimates account for village-level fixed effects. For variable definitions see Table 7.

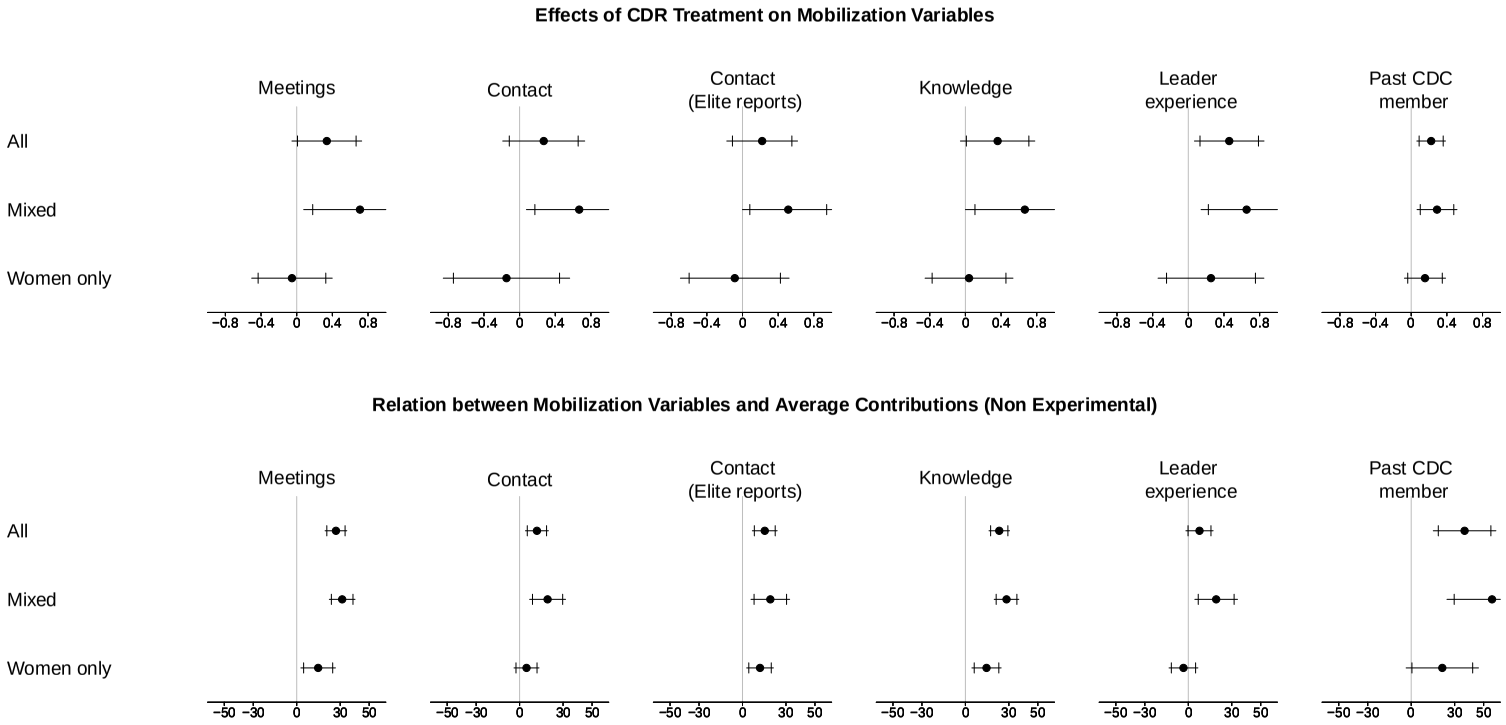

Figure 4: The upper panel shows the CDR effect on mobilization variables; the lower panel shows the relation (regression coefficients) between these intermediary variables and outcomes. Note that while estimates in the upper panel are identified by the randomization, those in the lower one are not. For variable definitions see Table 7.

Figure 4 about here

An alternative hypothesis, however, is that CDR mattered primarily via an impact on mobilization capabilities, and that it affected the use of these capabilities in the mixed communities but not in the all-women communities. Consistent with this idea, we find strong evidence that the CDR program caused an increase in mobilization activity in the mixed communities but not in the all-women communities. Figure 4 summarizes these results. In the mixed communities, CDR treatment caused significantly greater reports of community meetings, game-player knowledge, game-player reports of contact by elites, and elites' reports of contact with potential players.[29] By contrast, in communities where only women participated in the contribution game, CDR is not associated with more community meetings, game player knowledge, or contact efforts.

This mechanism – mobilization activity – mirrors the pattern observed with contributions: a positive impact of CDR in mixed gender groups, and zero effect in the communities where only women could play the game. Moreover, we show in the lower panel of Figure 4 the community-level correlations between measures of mobilization activity and average contributions in the game (the path P_3-P_5, which we note again is not identified by randomization). More mobilized communities gave significantly more in the contribution game, in both the mixed and all-women groups. Thus mobilization seems to matter everywhere, even if the CDR program only affected its extent in the mixed groups.[30]

These same patterns hold at the individual level, when we compare women in the mixed groups to women in the all women groups on their responses to the questions about community

[29]Among the three sets of respondents, CDR's impact on the community representative reports of contact effort are positive but not statistically significant. For the game player reports, these effects in the mixed groups weaken slightly in terms of statistical significance when we block on quarter status, though the estimated ATEs remain positive and non-trivial in substantive terms. Elite reports on contact activity remain strong and significantly positive even when controlling for quarter.

[30]The non-parametric method of "mediation analysis" proposed by Imai, Keele, Tingley, and Yamamoto (2011) estimates that in the mixed communities the "average causal mediation effect" of our mobilization measures on average contributions is 19LD out of a total effect of 43LD, or 44% (the estimate drops to 7.4 of 29LD, or 25%, if we control for the pre-treatment variable quarter). This is using the factor score for the three composite measures of meetings, contact and game-player knowledge, and of course making the "sequential ignorability" assumption under which the mediation effect is identified.

meetings, knowledge of the game and projects, and contact experience. We find generally strong positive effects of CDR treatment on women in the mixed groups, but no effects at all for women in the all women groups (including controlling for quarter).

Mobilization effects can also be observed in the choice of leadership. Recall that communities were asked to select a set of representatives to receive and manage the funds for the public goods project. The final columns in Figure 4 show the effect of CDR on survey responses to questions about the prior leadership activity of these representatives. According to game players, the CDR treatment is associated with significantly more experienced representatives where the game was played by a mixed gender group, but no significant difference in experience where only women were selected to play.[31] The differences may stem in part from the use communities made of structures put in place by the CDR program. We asked the community representatives directly "were you a member of a donor-sponsored community development council?" Since the CDR program created CDCs, we might expect to find a positive effect and indeed we do: about 68% say "yes" in the CDR communities, versus 46% in the control communities, a strongly significant difference. [32] However, when we distinguish between mixed and all-women groups, we find that the effect is larger in mixed, while smaller and statistically insignificant (though positive) in the all-women groups (final column, Figure 4).

These results suggest that CDR's impact occurs, at least in part, through changes in the ability and willingness of community leaders to mobilize and share information. They also provide a partial answer to the puzzle of why the CDR program increased community collective action where both genders could participate in the contribution game, but not where only women played: CDR increased mobilization activity in the mixed communities, but not in the all-women communities, and mobilization activity is strongly associated with bigger

[31]Similar patterns hold if we examine the representatives' own accounts of leadership experience, but not if we use reports by the chiefs. The pattern is not explained by the gender composition of community representatives. The all-women groups did have a somewhat higher frequency of women as community representatives (59% versus 48% in the mixed groups), but the pattern persists when we condition on gender of the representatives.

[32]The fact that control communities sometimes made use of donor-sponsored community councils reflects the fact that a variety of donors, outside of the organizations involved in this CDR project, create new community institutions to support their activities.

contributions. This might be because the CDR program affected men more than women, and men mobilized women (in addition to men) when they knew that men could play the game, or women sought to match how others were expected to play. Or it might be that leaders in the mixed communities drew on skills, experience, or networks created during the CDR program, whereas leaders in the all-women communities handed off the problem of organizing to play the contribution game to (possibly more traditional) women leaders and networks that were not affected by the CDR program. We return to these possibilities below.

We close this analysis by considering a decomposition of CDR's direct effects and effects via mobilization activity on these several measures of game player beliefs and preferences, using the approach of Imai, Keele, Tingley, and Yamamoto (2011). For consistent estimates, this method requires the satisfaction of an assumption that Imai et al. call "sequential ignorability." This is a demanding assumption and we highlight that inferences from this analysis are valid only to the extent that the assumption holds.[33] The results, provided in Table 5, reinforce the view that the effects of CDR worked through a coordination mechanism and through the employment of democratic processes, and moreover that a large share of these effects operate though greater community level mobilization and not simply through direct effects on the attitudes and beliefs of citizens. [34]

Table 5 about here

[33]The key requirement of "sequential ignorability" in this setting is that, conditional on actual treatment, it is as if the value of intermediary variables are randomly assigned relative to potential outcomes. One can see logics by which this assumption could be violated in this setting. Consider for example the intermediary outcome "Trust in leaders"; to justify the use of the Imai et al approach here we can allow that particiaption in CDR may affect the level of trust, but, conditional on this effect, there should be no third variables that affect both the level of trust and behavior in the game. In this example, the assumption might be violated if the size of village, for example, independently reduced the level of trust and the value of the public good, and hence the contributions of players.

[34]Note that the decomposition does not include the intermediate variable capturing material effects, since it isn't plausible that the mobilization activities for the matching funds experiment operated through improvements in material well-being.

Table 5: CDR effects, Direct and Indirect (via mobilization),
Mixed Communities Only

	Intermediate variable	Mob. Effect	Direct Effect	Total Effect
μ	Project satisfaction	0.2	-0.2	0
	Leader trust	0.017	0.051	0.068
	NGO trust	0.058	-0.165	-0.107
θ	Not anonymous	0.238	0.103	0.341
	Dem process	0.443*	0.545*	0.988**
λ	Expectations	0.363*	0.099	0.462
σ	Accuracy	0.329*	0.488*	0.816**

Note: Effect of CDR on intermediate variables divided between a direct pathway and a pathway through mobilization measures. Intermediate variables have standard deviation of 1.
** : $p < .01$; * : $p < .10$

6 Conclusions

A field experiment in which villages in northern Liberia were randomly assigned to receive a community-driven reconstruction program provides evidence that the introduction of new institutions and practices can alter patterns of social cooperation in a way that persists after the program's conclusion. Villages exposed to the development program exhibit higher subsequent levels of social cooperation than those in the control group, as measured through an innovative matching funds experiment that enables us to observe individual and community level contributions to a public good. The estimated effect is of a similar magnitude to the effect of a quadrupling of interest rates. These results suggest that changes in community capacity for collective action can take place over a short period of time; can be the product of outside intervention; and can develop without fundamental changes either to the structure of economic relations or to more macro-level political processes. However, we found strong heterogeneity in the measured effect of the CDR program. When the matching funds experiment was carried out with mixed gender groups, CDR had a very large impact; when we invited only women to participate, the estimated effects were zero or negative. These mixed results underscore the importance of specifying and assessing the mechanisms by which CDR changes

political and social outcomes.

Commonly proposed mechanisms include the idea that elections to CDCs can increase accountability and trust in local leadership; that participatory processes can lead to selection of more widely valued projects; and that the experience of working together on a CDD/CDR project can increase post-conflict reconciliation or involve marginalized community members in ways that lead to greater community cohesion. Our study finds little or no evidence that the program increased trust in community leaders, led to the selection of more highly valued projects, or caused people to put more weight on community welfare. It also appears highly unlikely that the CDR program changed individual preferences for public goods relative to income by increasing community income levels, since we find no evidence that the program had a significant impact on measures of average material well-being.[35]

Nonetheless, we do find strong evidence that exposure to CDR caused greater mobilization efforts by leaders in the communities where both men and women could be selected to provide matching funds. In these "mixed-gender" settings, CDR treatment communities selected leaders with greater experience who ran more meetings, communicated more with village members, and imparted a greater understanding of the process. We also found some evidence that CDR led communities to use more democratic methods in the selection of community representatives and projects, and that CDR is associated with more common expectations among game players about how others in the community would act. Again, however, these CDR effects were observed only in communities where both men and women could participate in the contribution game, and not in communities where it had been announced that only women would be selected.

When we introduced the gender variation to the development project experiment, we conceived of it as a cross-cutting manipulation distinct from the CDR treatment. However, it can also be thought of as providing two ways of measuring the impact of the CDR program.

[35]This is perhaps not surprising given that this particular CDR program was conceived to target governance and community collective action capacity rather than incomes and livelihoods; for example, direct income-generating projects were not allowed. Other studies have found evidence of welfare improvements where aid projects facilitated economic activity and focused not only on local public goods (Barron, Humphreys, Paler, and Weinstein, 2009; Casey, Glennerster, and Miguel, 2012)

In effect, we have one measurement of the effect of the CDR program on collective action problems facing mixed gender groups, and one for problems facing women-only groups. We found evidence that CDR improved collective action capacity for the former, but not the latter. What accounts for the difference?

One possible explanation focuses on the interaction of heterogeneous effects and strategic rationales. Under this account, the interventions involved and affected men more than women, women in the mixed groups sought to match what they expected others would contribute, and this was expected to be higher in the CDR-treated communities because of the CDR effect on men.[36]

A second possibility is that CDR increased the number or experience of a cadre of leaders who could be deployed to inform and mobilize the community to contribute in the matching funds challenge. Because the CDR program had mandated mixed-gender structures (the CDCs, for example), leaders in the CDR-treated, mixed gender communities may have found it natural to employ people and networks who had prior experience mobilizing collective action across gender lines. By contrast, in the all-women communities, the male leadership may have simply handed the problem off to women leaders, who have well-developed, traditional networks and social institutions that could be used to mobilize the women of the community.[37] Consistent with this hypothesis, it appears that more of the community representatives in the mixed-gender communities had prior experience in CDCs than in the all-women communities.

If this hypothesis is correct and if the mechanism generalizes, it suggests that the value of CDD/CDR programs might be less in their having a broad impact on the beliefs and preferences of community members about democracy, governance, or the value of inclusion, but rather in their effect on what might be called *leadership capital* – local leaders' skills and experience in coordinating and mobilizing collective action for collective action problems

[36]They expected higher contributions of women in all-women communities because women were expected to be more community-minded in general, and/or because there were fewer women to be informed and mobilized about the contribution game.

[37]Patriarchal norms have long dominated traditional and customary institutions in Liberia. At the same time, women were able to carve out a distinct space for power and leadership. While elder, male chiefs governed villages, mutually exclusive secret societies for men and women operated in parallel. The Sande, or female secret society, gave women the opportunity to voice their views in mixed gender settings and exercise authority and control over women in the community (Fuest (2008)).

that must be addressed by broad cross sections of a community. If leaders confront social dilemmas of this kind, prior exposure to inclusive decision-making of the form promoted by external actors may have large effects. However, because existing structures of authority are often organized around single-gender or other more narrowly constituted groups, this new leadership capital may prove redundant.

This interpretation resonates with a finding from a study in Afghanistan by Beath, Christia, and Enikolopov (2013). In that study, researchers assessed the effects of exposure to a CDR intervention on behavior in a wheat distribution exercise. The study found that when communities were unconstrained in determining how wheat would be distributed, there were no improvements attributable to the CDR intervention. However distribution improved on their measures when local development committees were formed *and delegated* to manage the distribution task, compared to control areas where there was neither a CDR intervention or delegation. The authors interpret the result as indicative of gains from delegation. Our results suggest instead an interpretation that points to a weakness of the CDR model: that the model may create effective institutions but not provide the incentives for communities to make use of them endogenously. This interpretation also helps square our results with the negative results of Casey, Glennerster, and Miguel (2012) and Humphreys, Sanchez de la Sierra, and van der Windt (2013). Rather than employing experimental games, these studies examined behaviour in more naturalistic settings that gave communities significant discretion over *who* should resolve the collective action challenges they faced. Our results suggest that in these cases it is possible that CDR altered capacity but that any such gains counted for little when communities could employ alternative institutional structures.

Our focus on mechanisms challenges the underlying causal model and suggests a need to reconsider the promises made by advocates of interventions like these that seek to alter the decision making processes in developing areas. We saw that exposure to participatory and inclusive decision-making through CDR does not necessarily lead individuals and communities to embrace more democratic processes, include marginalized groups, and better hold their leaders to account. Decisions that communities make about how to organize themselves depend

very much on the kinds of challenges they confront. We found evidence that the leadership capital built up through CDR facilitates the mobilization of cross-community participation and engagement, but only in a context in which we decided, as outsiders, that everyone in the community had an equal role to play in addressing the challenge. In the real world, where communities themselves choose who to involve in addressing local concerns, traditional approaches to problem solving may suffice, and any benefits of CDR may only be in evidence when cross-cutting organization is necessary.

References

Acemoglu, D. (2010): "Why Development Economics Needs Theory," *The Journal of Economic Perspectives*, 24(3), 17–32.

Andrews, M. (2013): *The Limits of Institutional Reform in Development*. Cambridge University Press.

Aumann, R. (1976): "Agreeing to disagree," *The annals of statistics*, 4(6), 1236–1239.

Avdeenko, A., and M. J. Gilligan (2013): "International Interventions to Build Social Capital: Evidence from a Field Experiment in Sudan," Discussion paper, Working Paper, NYU.

Baldassarri, D., and G. Grossman (2011): "Centralized sanctioning and legitimate authority promote cooperation in humans," *Proceedings of the National Academy of Sciences*, 108(27), 11023–11027.

Bardhan, P. (2002): "Decentralization of governance and development," *The Journal of Economic Perspectives*, 16(4), 185–205.

Barron, P., M. Humphreys, L. Paler, and J. Weinstein (2009): "Community based Reintegration in Aceh," *Indonesian Social Development Papers*.

Beath, A., F. Christia, and R. Enikolopov (2013): "Do Elected Councils Improve Governance? Experimental Evidence on Local Institutions in Afghanistan," 6510, MIT Political Science Department Research Paper No. 2013-24.

Bowles, S., and H. Gintis (2004): "The evolution of strong reciprocity: cooperation in heterogeneous populations," *Theoretical Population Biology*, 65(1), 17–28.

Casey, K., R. Glennerster, and E. Miguel (2012): "Reshaping Institutions: Evidence on Aid Impacts Using a Pre-Analysis Plan," *Quarterly Journal of Economics*, 127(4), 1755–1812.

Coleman, E. A., and M.-C. Lopez (2012): "Reconstructing Cooperation from Civil Conflict: Experimental Evidence from Colombian Development Policy," Working Paper.

Dal Bó, P., A. Foster, and L. Putterman (2010): "Institutions and Behavior: Experimental Evidence on the Effects of Democracy," *American Economic Review*, 100(5), 2205–2229.

Fearon, J. D., M. Humphreys, and J. M. Weinstein (2009): "Development Assistance, Institution Building, and Social Cohesion after Civil War: Evidence from a Field Experiment in Liberia," Unpublished manuscript.

Freedman, D. A. (2008): "On Regression Adjustments to Experiments with Several Treatments," *Annals of Applied Statistics*, 2, 176–96.

Fuest, V. (2008): "This is the Time to Get in Front: Changing Roles and Opportunities for Women in Liberia," *African Affairs*, 107(427), 201–224.

Gerber, A. S., and D. P. Green (2012): *Field experiments: Design, analysis, and interpretation*. WW Norton.

Green, D. P., S. E. Ha, and J. G. Bullock (2010): "Enough Already about Black Box Experiments: Studying Mediation Is More Difficult than Most Scholars Suppose," *The ANNALS of the American Academy of Political and Social Science*, 628(1), 200–208.

Habyarimana, J., M. Humphreys, D. Posner, and J. Weinstein (2007): "Why Does Ethnic Diversity Undermine Public Goods Provision?," *American Political Science Review*, 101(4), 709–725.

Haley, K. J., and D. M. Fessler (2005): "Nobody's watching?: Subtle cues affect generosity in an anonymous economic game," *Evolution and Human Behavior*, 26(3), 245–256.

Hamman, J., R. A. Weber, and J. Woon (2011): "An Experimental Investigation of Electoral Delegation and the Provision of Public Goods," *American Journal of Political Science*, forthcoming.

Humphreys, M., R. Sanchez de la Sierra, and P. van der Windt (2013): "Exporting Institutions: Evidence from a Field Experiment in Congo," Unpublished ms., Columbia University.

Imai, K., L. Keele, and D. Tingley (2010): "A General Approach to Causal Mediation Analysis," *Psychological Methods*, 15, 309–334.

Imai, K., L. Keele, D. Tingley, and T. Yamamoto (2011): "Unpacking the Black Box of Causality: Learning about Causal Mechanisms from Experimental and Observational Studies," *American Political Science Review*, 105(4), 765–789.

Kling, J. R., J. B. Liebman, and L. F. Katz (2007): "Experimental Analysis of Neighborhood Effects," *Econometrica*, 75(1), 83–119.

Ledyard, J. O. (1995): "Public Goods: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel, and A. E. Roth, pp. 111–94. Princeton University Press, Princeton, NJ.

Levitt, S. D., and J. A. List (2007): "What do laboratory experiments tell us about the real world," *Journal of Economic Perspectives*, pp. 153–174.

Ludwig, J., J. R. Kling, and S. Mullainathan (2011): "Mechanism Experiments and Policy Evaluations," *Journal of Economic Perspectives*, 25(3), 17–38.

Mansuri, G., and V. Rao (2004): "Community-Based (and Driven) Development: A Critical Review," *World Bank Research Observer*, 19(1), 1–39.

——— (2012): *Localizing Development: Does Participation Work?* The World Bank, Washington, DC.

Mutz, D., and R. Pemantle (2011): "The Perils of Randomization Checks in the Analysis of Experiments," Draft, 2011; U Penn.

Nunn, N., and L. Wantchekon (2011): "The Slave Trade and the Origins of Mistrust in Africa," *American Economic Review*, 101(7), 3221–52.

Putnam, R. (1994): *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton University Press, Princeton, NJ:.

Richards, P. (2005): "Community Cohesion in Liberia," *Social Development Paper*, 21.

Voigtländer, N., and H.-J. Voth (2012): "Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany*," *The Quarterly Journal of Economics*, 127(3), 1339–1392.

Voors, M. J., E. E. M. Nillesen, P. Verwimp, E. H. Bulte, R. Lensink, and D. P. V. Soest (2012): "Violent Conflict and Behavior: A Field Experiment in Burundi," *American Economic Review*, 102(2), 941–64.

Whitt, S., and R. K. Wilson (2007): "The dictator game, fairness and ethnicity in postwar Bosnia," *American Journal of Political Science*, 51(3), 655–668.

Online Appendices: Balance Table, Summary Statistics, Variable Definitions and Detailed Results

A Timeline

Table 6: Timeline of study

Date	Event
April/May 2006	Baseline survey of 1702 households in Voinjama and Zorzor, asking questions about household characteristics and opinions on governance and community functioning.
Summer 2006	Identification of 83 communities with hub towns and satellite villages meeting population thresholds for the IRC project
October 2006	Public lotteries held in Voinjama and Zorzor allocate 42 communities to CDR treatment and 41 to control.
November 2006-March 2008	IRC implements CDR programs in the 42 treatment communities.
March/April 2008	Follow-up panel survey of households, material welfare indicators and attitudes on governance, community functioning.
Late June-September 2008	Matching funds experiment conducted in T and C hub villages, a behavioral measure of community collective action capacity. Surveys of contribution game players, village chiefs, and communities representatives.

B Balance Table

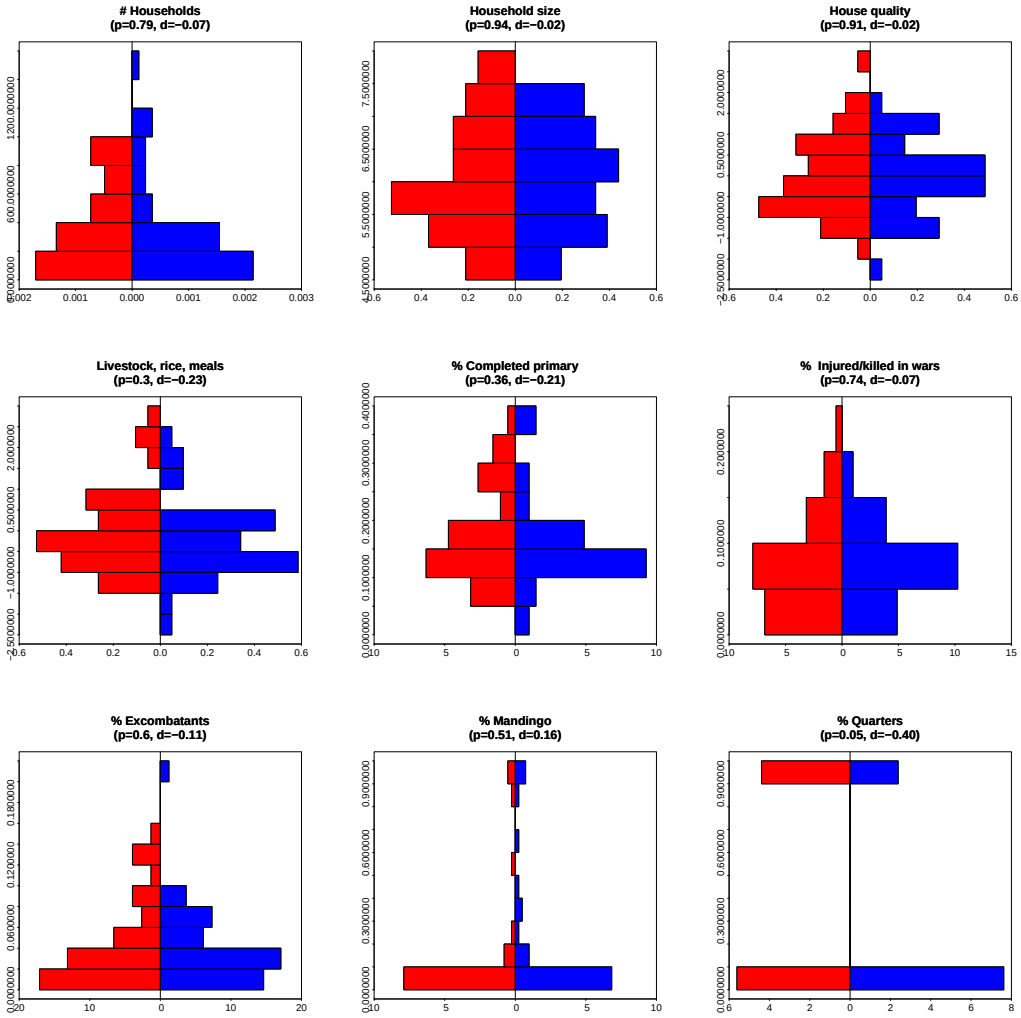

Figure 5: Notes: Figures show histograms of each variable for treatment (blue) and control group (red). Graph titles report p values from t-tests of the null hypothesis of no relation between each variable and treatment assignment and "d" values which report the difference in means between treatment and control groups in standard deviations of the control group values. For variable definition see Table 7.

C Construction of indices

Given m measures $\{X_1, X_2, \ldots, X_m\}$, mean effects indices were constructed as follows.:

Define:

$$\zeta_j(X_1, X_2, \ldots, X_m) = \frac{1}{m} \sum_{k=1}^{m} \left(\frac{\overline{x}_{kj} - \frac{1}{n_c} \sum_{j \in C} \overline{x}_{kj}}{\sqrt{\frac{1}{n_c} \sum_{h \in C} \left(\overline{x}_{kh} - \frac{1}{n_c} \sum_{j \in C} \overline{x}_{kj} \right)^2}} \right) \tag{7}$$

where each variable X_i is coded so that positive values have a substantively positive interpretation, C is the set of of control community indices and $\overline{x}_{kj}$ is the average outcome on measure k in community j. Then our index is given by:

$$\xi_j(X_1, X_2, \ldots, X_m) = \frac{\zeta_j - \frac{1}{n_c} \sum_{j \in C} \zeta_j}{\sqrt{\frac{1}{n_c} \sum_{h \in C} \left(\zeta_j - \frac{1}{n_c} \sum_{j \in C} \zeta_j \right)^2}} \tag{8}$$

In practice we define variables at the community level (the level at which treatment is assigned) and calculate the averages and standard errors in equations 7 and 8 using non-missing data only, thus in some cases for a given unit, ζ_j may use data for only a subset of variables from all units.

D Variable Definition and Summary Statistics

Table 7: Table of Variables used in text (Part I)

Variable	Definition	Source	n	Min	Max	sd
CDR	Treatment: Village Participation in IRC CDR project	Project files	83	0	1	0.5
Mixed	Treatment: Communities in which 50% men and 50% women took part in the public goods experiment	Project Files	83	0	1	0.5
Women only	Treatment: Communities in which only women took part in the public goods experiment (1-Mixed)	Project Files	83	0	1	0.5
Interest rate	Treatment: Multiplier on individual contributions ($0 = \times 2, 1 = \times 5$)	Project Files	83	0	1	0.5
Quarter	Indicator for whether a hubtown is a neighborhood of an urban area	Project Files	83	0	1	0.47
Items in Figure 5	(Baseline measures)					
# Households	# Households in village	United Nations	83	33	1500	315
Household size	Village Average	BL: 16	79	4.5	7.9	0.8
House quality	Mean effects index: # Rooms and doors, mud brick or better walls, zinc roof, piped water (6 items).	BL: 45-47.	79	-2.1	2.7	0.93
Food / livestock	Mean effects index: # Chickens, sheep, guinea fowl, # meals/day, tins of rice planted (5 items).	BL: 45, 49, 56.	79	-2.2	2.8	0.97
Primary education	Household members that completed primary education	BL: 28	79	0	38	8.2
Share injured or died	Percentage	BL: 32, 34	79	0	23	4.6
Share excombatant	Percentage	BL: 33	83	0	22	4.1
Share Mandingo	Percentage	BL: 27	83	0	100	28
Variable	**Definition**	**Source**	n	**Min**	**Max**	**sd**
Figure 2						
Satisfaction with projects	Mean effects index: Project was among the most important things for village; liking projects was a factor in decision making (2 items)	GS: 47 (or 50(2)), 67a,	83	-3.42	1.51	1.02
Trust in leaders	Mean effects index: Are leaders trustworthy (2 items)? Would they steal from community? Would they use money effectively? Would representatives benefit disproportionately? Or the chief? Was believing money would be used well important for decisions? (7 items)	GS: 18-20, 32 (or 41), 51, 52, 67c	83	-2.09	2.13	0.97
No anonymity	Did subjects think it likely that others villagers would find out how much they contributed?	GS: 67	83	-1.07	3.06	1.03

Table of Variables used in text (Part II)

Variable	Definition	Source	n	Min	Max	sd
Figure 2	Continued					
Social Desirability	Mean effects index: Index on whether individuals would appeal first to local or international community for assistance? likely effectiveness of the appeal? (4 items)	EL: 100	83	-2.54	2.12	0.87
Participation Index	Mean effects index: were there community meetings to discuss game? who organized them? were community representatives elected? were meetings public? was there competition for positions? did some say that they preferred different representatives? are Community Representatives relatives of the chief? was the process of selecting projects considered important for decisions? (8 items)	GS: 25, 27, 33-36, 37 (or 39), 67b	83	-2.17	2.28	0.97
Expectations	Mean effects index: Of the other players, how many contributed all? how many nothing? what did most do?	GS: 58–59	83	-3.5	1.32	0.88
Accuracy of expectations	Mean effects index: Squared deviation between expected and actual number contributing 300, squared deviation between expected and actual number contributing nothing.	GS: 58–59	83	-2.97	1.43	0.89
Figure 4						
Meetings	Mean effects index: Existence and participation in community meetings (5 items)	GS: 24–26, 28, 54	83	-1.98	1.99	0.89
Contact	Mean effects index: Were game players contacted about the game, or the project? were they asked to stay home? urged to contribute?	GS: 55, 56, 57, 66	82	-0.8	1.4	0.5
Contact (Elite Reports)	Mean effects index: Did elites contact population to discuss the game to discuss the project? Encourage people to spread the word? did they tell them to stay home? urge them to contribute? (3 items)	CR: 64(1), 65, 66, 67	83	-2.22	1.33	0.93
Knowledge	Mean effects index: Do players know what the project is? do they know who organized the meetings? how many attended? who the community representatives are? what the projects are? (5 items)	GS: 22, 27-9, 42	83	-1.72	2.46	0.98
Leader Experience	Mean effects index: Have community representatives played a leadership role in the community before?(2 items)	GS: 31, 40	83	1.61	5.52	0.89
Past CDC member	What share of community reps were former CDC leaders?	CR: 34	75	0	1	0.37

Key: Baseline survey (BL) (LINK), Endline survey (EL) (LINK), Postgame survey (GS) (LINK), Community representative (CR) survey (LINK), Chief Game survey (CG) (LINK)

E Table Versions of Figures in Text

Table 8: CDR Effect on Mobilization Measures, and Relation between Mobilization Measures and Contributions (Figure 4)

	CDR Effect			Impact on Contributions		
	All	Mixed	Women only	All	Mixed	Women only
Meetings	0.338	0.710	-0.053	27.1	31.4	14.8
	(0.093)	(0.028)	(0.830)	(0.000)	(0.000)	(0.020)
Contact	0.271	0.668	-0.147	12	19.3	4.8
	(0.260)	(0.030)	(0.699)	(0.004)	(0.004)	(0.287)
Contact (Elite reports)	0.22	0.513	-0.087	15.4	19.2	12.1
	(0.285)	(0.065)	(0.782)	(0.001)	(0.007)	(0.014)
Knowledge	0.361	0.667	0.041	23.4	28.4	14.6
	(0.101)	(0.056)	(0.87)	(0.000)	(0.000)	(0.008)
Leader experience	0.459	0.654	0.254	7.7	19.2	-3.3
	(0.021)	(0.011)	(0.403)	(0.111)	(0.014)	(0.517)
Past CDC member	0.224	0.29	0.155	36.9	55.8	21.5
	(0.009)	(0.014)	(0.211)	(0.001)	(0.001)	(0.100)

Notes: First three columns give CDR treatment effect on standardized measures of mobilization activity (p values in parentheses, from randomization inference). Second three columns give coefficients for average individual contribution for community regressed on the standardized mobilization measures (p values from regression).

Table 9: Mechanisms: CDR Effects on Behavior and Attitudes, and Behavior/Attitude Relationship with Contributions (Figure 2)

	CDR Effect			Impact on Contributions		
	All	Mixed	Women	All	Mixed	Women
Satisfaction with Projects	0.235	0.143	0.332	7.7	12.4	-1.0
	(0.316)	(0.698)	(0.245)	(0.069)	(0.033)	(0.859)
Trust in leaders	0.127	0.09	0.167	-2.5	-1.2	-7.0
	(0.569)	(0.79)	(0.578)	(0.587)	(0.859)	(0.209)
No Anonymity	-0.054	0.294	-0.419	11.5	19.1	5.0
	(0.819)	(0.308)	(0.252)	(0.006)	(0.007)	(0.252)
Social Desirability	-0.077	-0.178	0.028	5.2	5.4	9.2
	(0.683)	(0.513)	(0.91)	(0.303)	(0.471)	(0.124)
Participation	0.492	0.994	-0.034	14.2	21.7	-0.9
	(0.024)	(0.004)	(0.902)	(0.001)	(0.000)	(0.885)
Expectations	0.32	0.518	0.112	17.7	17.7	3.0
	(0.098)	(0.131)	(0.496)	(0.000)	(0.003)	(0.759)
Accuracy	0.393	0.815	-0.049	30.9	32.4	19.9
	(0.037)	(0.008)	(0.786)	(0.000)	(0.000)	(0.009)

Notes: First three columns give CDR treatment effect on standardized measures of community behavior and attitudes (p values in parentheses, from randomization inference). Second three columns give coefficients for average individual contribution for community regressed on standardized behavior and attitude measures (p values from regression).

Table 10: Relation between player's contribution and player's expectation of others' contributions (Figure 3)

	In CDR communities			In no-CDR communities		
	All comm's	Mixed	Women only	All comm's	Mixed	Women only
All	11.0	17.2	4.6	11.7	17.0	4.3
	(0.008)	(0.002)	(0.463)	(0.001)	(0.001)	(0.385)
Women	8.0	14.4	4.6	7.1	11.4	4.3
	(0.108)	(0.08)	(0.463)	(0.079)	(0.099)	(0.385)
Men	22.0	22.0		23.7	23.7	
	(0.008)	(0.008)		(0.001)	(0.001)	

Notes: Regression estimates. Dependent variable is player's contribution in Liberian Dollars; independent variable is player's belief about others' contributions (scaled to 1 sd in control communities). Sample broken into 8 disjoint strata according to gender of respondent, CDR treatment status, and gender treatment status. All estimates account for village-level fixed effects.

F Extra Tables

Table 11: Linear regression with controls

	Model 1	Model 2	Model 3	Model 4	Model 5
(Intercept)	227.79***	200.54***	219.38***	235.17***	238.37***
	(6.05)	(7.64)	(7.89)	(8.89)	(9.11)
CDR treatment	14.58†	43.05***	31.99**	27.04**	26.20**
	(8.56)	(10.56)	(9.67)	(9.25)	(9.21)
All Women		53.20***	45.78***	43.62***	40.50***
		(10.68)	(9.62)	(9.09)	(9.30)
CDR*All Women		−55.82***	−47.17***	−43.56**	−41.13**
		(15.12)	(13.55)	(12.83)	(12.86)
Quarter			−34.26***	−36.76***	−37.50***
			(7.33)	(6.95)	(6.93)
# households (100s)				−8.80**	−8.23**
				(2.71)	(2.72)
Share Mandingo					−17.11
					(12.12)
N	82	82	82	82	82
RMSE	38.75	34.17	30.35	28.62	28.44
R^2	0.04	0.27	0.43	0.50	0.51
adj. R^2	0.02	0.24	0.40	0.47	0.47

Standard errors in parentheses

† significant at $p < .10$; $^{*}p < .05$; $^{**}p < .01$; $^{***}p < .001$

Table 12: Mobilization in Communities where Women Only Played the Game versus Both Genders

	Mixed	Women Only	Difference	p value
Meetings	0.389	0.303	-0.086	0.721
Contact	0.42	0.154	-0.266	0.441
Contact (Elite reports)	0.413	0.001	-0.412	0.124
Knowledge	0.493	0.247	-0.246	0.426
Leader experience	4.676	4.664	-0.012	0.962
Past CDC member	0.736	0.636	-0.100	0.361

Notes: Comparison of intermediary outcomes between women-only and mixed gender CDR areas on mobilization measures. For variable definitions see Table 7.