Building Cooperation Among Groups in Conflict: An Experiment on Intersectarian Cooperation in Lebanon

Han Il Chang
New York University-Abu Dhabi
hc665@nyu.edu

Leonid Peisakhin
New York University-Abu Dhabi
leonid.peisakhin@nyu.edu

Abstract: Societies divided along ethnic or religious lines suffer from persistent conflict and underprovision of public goods. Scholarly understanding of how to strengthen intergroup cooperation remains limited. In this study, we set out to test the effectiveness of two interventions on intergroup cooperation: cross-group expert appeal and participation in a cross-group discussion. The laboratory-in-the-field experiment is set in Lebanon's capital Beirut and involves interactions between 180 Shia and 180 Sunni Muslim participants. We find that the expert appeal increases intersectarian cooperation in settings that do not entail reciprocal exchange. On average, cross-sectarian discussions do not improve cooperation, but those discussions where participants delve deeply into the conflict's causes and possible remedies are associated with greater cooperation. Neither intervention diminishes the effectiveness of sectarian clientelistic appeals. The policy implication of our study is that intergroup cooperation can be strengthened even in regions as bitterly divided as the Middle East.

Word Count: 9,965 words

July 2018

We thank for helpful comments Thad Dunning, Susan Hyde, Peter van der Windt, Jeff Jensen, Rebecca Morton, Cole Tanigawa-Lau, three anonymous reviewers, and seminar participants at UC-Berkeley, NYUAD, and the 2016 WESSI Alumni Workshop. We are grateful to Jad Mahmoud Halabi and Azhar Yerzhanova, who provided excellent research assistance, and to Alicia Jammal at *Information International*, whose help with the logistics of the study was invaluable. The project was funded by an NYUAD Research Enhancement Grant. It was approved by NYUAD's Institutional Review Board under protocol #080-2016.

1. Introduction

Societies that are home to multiple ethnic or religious groups are known to have lower levels of public goods provision (Alesina, Baqir, and Easterly 1999) and higher levels of conflict (Horowitz 1985). That is because cooperation is generally lower across group boundaries than within groups, and prevalence of ingroup favoritism impedes outcomes that are beneficial to society as a whole (Akerlof and Kranton 2000). Difficulty of cooperation across group boundaries is a common problem globally and has been particularly acute in the Middle East in recent decades. There, sectarian tensions have contributed to political instability and economic underperformance. In this project, we test the various means of encouraging cooperation across sectarian lines between Shia and Sunni Muslims.

There is a consensus in the literature on intergroup cooperation that differences in group identity undermine cooperation by inhibiting altruistic behavior toward 'the other' (Allport 1954; Tajfel et al. 1971; Chen and Li 2009) and by making sanctioning more difficult because of limited information about the other group (Habyarimana et al. 2007). We advance the literature by exploring how intergroup cooperation can be fostered. Much work in the social sciences has explored the determinants of cooperation *within* groups; ours is among the few studies in political science to examine how to build cooperation *across* group boundaries (e.g. Fearon and Laitin 1996; Miguel 2004; Cederman et al. 2011; Larson 2017).

Our starting point is the insight from the common ingroup identity theory (Gaertner et al. 1993) that intergroup cooperation can be increased through activation of a common overarching identity (for an empirical illustration see Miguel 2004). There are different ways to activate a common identity: through a top-down appeal or a bottom-up engagement. One way to build consensus and construct overarching identity is by facilitating a discussion aimed at establishing cross-group commonality between members of competing groups. Studies on the causes of intergroup cooperation (e.g. Kerr and Kaufman-Gilliland 1994) and also the related literature on deliberative democracy (e.g. Fishkin 1997) have considered the effectiveness of group discussions. A major limitation of experimental work in this vein is that the theoretical expectations have been tested mostly in settings where group identities are arbitrarily assigned (for a review see Balliet 2010). In this project, we explore the effectiveness of a discussion intervention in a setting with highly salient real-world religious identities. We also test the effectiveness of a completely novel intervention—an appeal by experts—on intergroup cooperation. Experts are commonly believed to influence public opinion either by providing new information or legitimating certain points of view (Zaller 1992; Henrich and Gil-White 2001).

We distinguish between two different types of cooperation: *unconditional* and *conditional*. Unconditional cooperation entails friendly behavior toward a member of a different group with no expectation of reciprocity. Conditional cooperation takes place only if one expects the opposite group to respond cooperatively. Unconditional cooperation is a form of altruistic behavior and signals a high regard toward the outgroup

and a lower level of ingroup favoritism. In the world of politics, voting for a candidate from the opposite ethnic or religious group out of support for their policy proposals can be an instance of unconditional cooperation insofar as one is voting in support of policy that might have no direct bearing on one's self-interest.

In a laboratory in the field experiment involving 180 Shia and 180 Sunni Muslims in Beirut, Lebanon, we test what effect watching a short pre-recorded *expert appeal* to cooperate and, separately, participation in a *cross-sectarian discussion* have on the level of intersectarian cooperation.[1] The first intervention—the expert appeal—is modeled after a televised political talk show. In an unscripted conversation, four prominent journalists—two Shias and two Sunnis—discuss Lebanon's problems caused by sectarianism and conclude that the country would benefit substantially from greater cooperation across sectarian lines. In the second intervention, after watching the expert appeal, participants in mixed six-person Sunni-Shia groups engage in face-to-face discussions about issues relating to intersectarian cooperation. The content of the expert appeal catalyzes subsequent group discussions in the same way that reading materials provide a prompt for discussions in most studies on deliberation and in the way that all

[1] Lebanon is home to a large Christian population. We exclude Christians from this study. Substantively, the problem of Shia-Sunni relations is of broad relevance across the region, including in countries with Christian populations. Logistically, inclusion of a third group would have made the study more difficult to execute and more costly.

real life discussions follow on from some stimulus. We back out the effect of group discussion by comparing that intervention against the baseline of the expert appeal.

The effect of the two interventions on the levels of intergroup cooperation is tested in a series of games. We measure unconditional cooperation by observing participants' vote for sectarian candidates in simulated elections and, separately, by recording how much of a certain fixed financial resource participants allocate between anonymous members of their own and the opposite sect in a standard other-other allocation game. Conditional cooperation is measured by observing the level of contributions in a standard public goods game.

The study also examines the effects of a clientelistic exchange by randomly exposing some participants to a financial incentive to vote for their co-sectarian in the election game. We introduce this manipulation out of a concern that clientelism in the form of one-shot payments in exchange for votes is common in Lebanon and across much of the Middle East (Lust 2009; Corstange 2016). The presence of this intervention allows us to test whether expert appeal and group discussion reduce the effectiveness of clientelism. Ours is the first study to consider how interventions designed to increase intergroup cooperation interact with clientelism.

We find that watching experts discuss the benefits of intersectarian cooperation increases unconditional cooperation, but there is no effect on conditional cooperation. Leveraging self-reported scores of cross-sectarian trust we find that the expert appeal intervention

fails to increase cross-sectarian trust, whereas increased trust might be a precondition for higher levels of conditional cooperation. We theorize that expert appeals might fail to improve unconditional cooperation because such appeals more easily activate intrinsic motivations for cooperation but are not very effective at building intergroup trust.

Participation in a cross-sectarian group discussion about cooperation appears to have no effect either on unconditional or conditional cooperation. However, we present suggestive observational evidence that the effectiveness of group discussion is highly heterogeneous and varies with substantive discussion depth. A deep and substantive exchange that touches on many topics relating to intergroup cooperation is correlated with higher levels of both types of cooperation. We also illustrate how clientelism impedes intergroup cooperation in divided societies. Neither intervention succeeds in diminishing the effectiveness of clientelism.

Our research suggests that intergroup cooperation can be strengthened under certain conditions even in places as bitterly divided as the Middle East. The primary policy implication of this study is that dissemination of consensual cross-sectarian expert calls to cooperate—a top-down intervention—may in fact bring about greater cooperation across group lines in unconditional exchanges. Cooperation conditional on one's expectation about the behavior of the other group is likely to improve only once intergroup trust levels increase. Our findings also indicate that intersectarian cooperation in divided societies is unlikely to take root as long as clientelistic vote buying remains a common practice.

2. Context

This project entailed interactions between residents of Beirut, the capital of Lebanon. Lebanon is a highly fractured society. Fissures run along the lines of religious sectarian identities and are deeply historically rooted. Sectarianism is institutionalized. The 1943 National Pact heralding Lebanon's independence from France and the 1989 National Reconciliation Accord (the Taif Agreement) that concluded the intersectarian civil war (1975-1990) organize the system of government around sectarian differences.

The very question of the relative size of the three largest sects is extremely sensitive, and the only national population census in Lebanon's fraught history took place in 1932. There is a general consensus that Shia and Sunni Muslims are now close to parity, each at about 30% of the population, and that Christians make up most of the remainder (Faour 2007). The three sects have quotas in parliament: 64 seats are assigned to Christians, 27 each to Shia and Sunni Muslims, and the remaining 10 seats to other Muslim minorities. Basic services are commonly provided along sectarian lines, and the population is highly segregated regionally within the country and by neighborhood in Beirut. Marital practices are regulated independently by separate confessional courts, and most high school students attend schools dominated by a single confession (Baytiyeh 2016). External actors exert pressure on behalf of the competing sects: Iran supports Lebanon's Shias, whereas Saudi Arabia supports the Sunnis.

This environment of institutionalized sectarianism gives rise to clientelism. Most commonly, clientelism takes the form of vote buying on the polling day (Corstange 2012, 2016), although longer-term clientelistic arrangements that entail selective service provision are also present (Cammett and Isaar 2010). About 55% of the Lebanese electorate can be inferred to have sold their votes in the 2009 parliamentary election according to Corstange (2012: 493).

3. Hypotheses

We hypothesize that both interventions—expert appeal and participant discussion—should increase unconditional and conditional intersectarian cooperation. Our starting point is the insight from common ingroup identity theory that models how activation of an overarching group identity might decrease bias among members of the smaller constituent groups that make up the larger whole (Gaertner et al. 1993; Gaertner and Dovidio 2000). In this case, the smaller constituent groups are Shia and Sunni residents of Lebanon. The overarching group identity is the Lebanese national identity.[2] We expect the common Lebanese national identity to come to the fore for those exposed to the interventions. According to common ingroup identity theory, increased salience of shared membership in the Lebanese national community should increase cooperation levels among Shia and Sunni by increasing altruistic regard and trust across the sectarian divide.

[2] Common Arab and Muslim identities are available as other overarching identities, but neither the experts nor the participants stressed these in discussions.

The literature suggests several mechanisms by which a shared overarching identity might be made more salient by the treatments. In the expert appeal treatment, the four experts—two Shias and two Sunnis—agree that inter-sectarian cooperation is necessary to overcome the ills plaguing Lebanese society and that all religious communities must band around the common national identity. We expect this treatment to increase the salience of the Lebanese national identity insofar as a consensus among experts from different sects is persuasive. That experts can be persuasive in affecting political attitudes and behavior is well-established empirically (Zaller 1992, Henrich and Gil-White 2001).

As to the effectiveness of group discussion in increasing the salience of the overarching group identity, studies on group decision making and deliberative discourse suggest that face-to-face communication—specifically the need to stress shared identity characteristics in order to persuade listeners—helps participants to recognize the existence of commonalities that bind all group members together thus increasing cooperation within the broader group (for a review see Delli Carpini, Cook, and Jacobs 2004). This leads us to hypothesize that group discussions that bring together participants from conflicting groups should result in increased mutual regard among the participants.

In the context of a laboratory-in-the-field experiment, we are able to test the observable implications of the theoretical arguments laid out above. Those implications are that

exposure to a discussion among experts that stresses the importance of a common Lebanese national identity and participation in a group discussion about the benefits of inter-sectarian cooperation should both lead to increased cooperation across sectarian lines. This gives rise to the following hypotheses:

> H1: Participants randomized into the *expert appeal intervention* will be more likely to cooperate with members of the opposite sect in exchanges entailing both unconditional and conditional cooperation.
>
> H2: Participants randomized into the *group discussion intervention* will be more likely to cooperate with members of the opposite sect in both types of exchanges.

The study also contains a vote buying intervention. Clientelism is effective because financial incentives offered by brokers might override individuals' social preferences. Given that the expert appeal and group discussion interventions strengthen the overarching group identity, we expect these interventions to also strengthen intrinsic motivations such as positive attitudes toward a member of the opposite religious or ethnic group. We hypothesize that appeal to intrinsic motivations will be more powerful than the pull of financial incentives, as those who accepted clientelistic payments might be made to feel ashamed of having taken money to vote for a co-sectarian. The test of this hypothesis entails a comparison of those who were randomized into the vote buying treatment and either the expert appeal or group discussion intervention (an interaction) against the baseline of those who were only subject to vote buying.

> H3: Participants assigned to expert appeal and group discussion interventions will be *less likely* to vote for a co-sectarian candidate in simulated elections after being offered a financial incentive to do so than those only subject to the vote buying treatment.

4. Research Design

This laboratory-in-the-field experiment is built around three interventions:

(1) Viewing of a pre-recorded expert appeal about the benefits of intersectarian cooperation,

(2) Participation in a small-group discussion about intersectarianism after seeing the expert appeal, and

(3) Being offered money to vote for a co-sectarian in the subsequent election game.

All tasks, from watching the expert discussion to playing the games, are undertaken within six-person groups in a series of face-to-face interactions. Each group is made up of three Shia and three Sunni participants; membership in these six-person groups remains fixed over the course of the study.

4.1 Study population

Participants are drawn from across Lebanon's capital Beirut with a view to assembling a representative sample of Shia and Sunni residents of the capital. Participants' characteristics are described in Table 1. Two-thirds are from neighborhoods dominated by a single sect, and the remainder are from mixed neighborhoods. This is a fair approximation of how Muslims are distributed across Beirut neighborhoods (Salamey and Tabar 2008). The average age of participants is 37.[3] On average, participants have 12 years of schooling. They come from households with monthly incomes of around $2,100. Men and women, and Shia and Sunni Muslims, are all equally represented. In

[3] Age was capped at 64 to maximize the chance that participants were literate.

Appendix A, we compare the demographic characteristics of our participants to Arab Barometer data on Muslim residents of Beirut and Lebanon. The comparison suggests that participants in this study are generally representative of the reference populations with regards to age, education level, and income. The Shia and Sunni communities are similar in size in Lebanon; this is also the case by design in our study.

Table 1. Description of Participants: Summary Statistics.

	Mean	*Standard Deviation*	*Min*	*Max*
Shia/Sunni, %	0.50	0.50	0	1
Female/Male, %	0.50	0.50	0	1
Age, years	37	13	18	64
Education level[+]	3.56	1.50	0	6
Monthly household income level[++]	1.06	1.24	0	6
Predominately Shia neighborhood, %	0.30	0.46	0	1
Predominately Sunni neighborhood, %	0.35	0.48	0	1
Level of understanding of tasks[+++]	2.64	0.60	0	3

[+] Education levels: 3-secondary; 4-technical.
[++] Income: 1-1,000,001 to 3,000,000LBP; 2- 3,000,001 to 5,000,000LBP.
[+++] Level of understanding: 2-understood instructions quite well; 3-understood instructions fully.
Note: Survey question wording in Appendix B.

4.2 Setting

On arrival, individuals were randomly assigned to six-member groups. Each group was balanced with regard to sect (three Shia and three Sunni), gender (three men and three women), and age (three participants aged 18-40 and three aged 41-64). None of the group participants had met beforehand, and we asked participants not to speak with one another until instructed to do so. Participants were asked to keep their names secret throughout the study because names can be a marker of sectarian affiliation.

Our assistants acted as group moderators: they informed participants that there are three Shia and three Sunni Muslims in the group without revealing which specific participants

belonged to which sect or introducing them, read out instructions for every task in Arabic, and moderated the group discussion. Assistants followed written scripts and were trained not to deviate from pre-agreed instructions. They were randomly assigned to groups, but in the analyses that follow we include a control for group moderator effects.

The experiment took place in November 2016 over the course of twelve separate sessions that were completed in eight days. Five six-person groups were recruited for every session. Participants were randomly assigned to sessions subject to the constraints of sect, gender, and age requirements. All those attending a specific session were subject to the same experimental manipulation. There was no communication across groups.

4.3 Interventions

Table 2. Experimental conditions.

	Baseline	Expert Appeal	Expert Appeal + Participant Discussion
No Clientelism	*Experimental condition 1*: 60 participants	*Experimental condition 2*: 60 participants	*Experimental condition 3*: 60 participants
Clientelism	*Experimental condition 4*: 60 participants	*Experimental condition 5*: 60 participants	*Experimental condition 6*: 60 participants

The experimental conditions are summarized in Table 2. There are three of them: baseline, expert appeal, and expert appeal + participant discussion. This structure is replicated twice—once in the absence of a clientelistic intervention (conditions 1-3) and once in the presence of clientelism (conditions 4-6). The total number of participants in the study is 360 for the total of 60 six-member teams. Ten groups, i.e. 60 individuals, were randomly assigned to each intervention. This is a relatively small number of

participants per intervention, which raises a concern about the study's statistical power. In designing the experiment, we faced a tradeoff between statistical power (increasing the number of participants per intervention) and the study's realism (including a clientelism intervention). We opted for including a clientelism intervention over increasing the number of subjects.

The expert appeal intervention is set up to resemble a political talk show. In the video, four experts—two Shia and two Sunni male journalists in their 50s—discuss pressing problems in contemporary Lebanese politics, the promise of intersectarian cooperation, and of the coming together around a common national identity in solving the problems. The discussion is moderated by a female Christian journalist and is unscripted. The four experts and the moderator are all prominent in Lebanese media, contributing opinion pieces to major newspapers and appearing on national television. In their public appearances, all the experts consistently advocate for intersectarian cooperation and stress the importance of shared Lebanese national identity. Forty-three percent of respondents remembered encountering at least one of the experts' names either on television or in print.[4] Respondents found the experts to be persuasive: 71% of those who watched the expert appeal considered all four experts to be either very or quite persuasive (for individual expert scores disaggregated by participants' sect see Appendix C).

[4] There is no difference in the rate of expert recognition among participants assigned to the expert appeal and group discussion conditions.

The expert appeal is 15 minutes in duration. The discussion opens with all the experts agreeing that sectarianism is the main problem in Lebanese politics. The nature of the problem is then illustrated: sectarianism is linked to corruption and nepotism, lack of a unifying Lebanese national identity, infrequent social mixing across sects, and institutional sclerosis and poor quality of basic public services. The discussion closes with participants appealing for intersectarian cooperation with the aim of creating a common Lebanese national identity so that problems of corruption, inefficiency in the provision of public services, and fostering of preconditions for violent conflict can be overcome. The transcript of the video (in English) and the video itself (in Arabic) are available in Appendices D and E.

The participant discussion intervention is designed to follow on from experts' appeal in order to ensure that participants talk specifically about intersectarian cooperation. In real-world interactions, political appeals are often followed by conversations in a group setting. Our treatment takes on the form of precisely such a group discussion where participants have a chance to react to the message that they had just been exposed to. Experiments on deliberation commonly provide a stimulus—usually a written text—prior to discussion. In this instance, a written prompt is replaced by a pre-recorded message from experts.

The group discussions lasted 28 minutes on average; the shortest was 18 minutes and the longest 39 minutes. Discussions were relatively freewheeling; group moderators were asked to pose a series of questions only in order to direct the conversation (for moderator

instructions see Appendix F). Group discussion transcripts (translated to English) are available in Appendix G.[5]

In designing the study we tried to approximate Lebanon's actual political environment. One obvious challenge to studying the effectiveness of expert appeals or deliberative discussions is that political behavior in Lebanon, just like in many developing democracies, is not simply a product of individual preferences but is subject to pressure from vote-buying intermediaries. The question then is not just whether our interventions can sway members of different sects to cooperate across the sectarian divide but, more realistically, whether these interventions can be effective in the presence of clientelism. To test this proposition we assigned 180 participants in their six-person groups to a clientelism intervention within which the three experimental conditions were replicated.

The clientelistic appeal was delivered by our confederates who acted as "election brokers." This happened in the context of checking participants' contact details. Participants were called to a private room to meet with our confederate one-on-one. We used two confederates: one Shia and one Sunni, both male, and in their mid-20s. Participants always met with the confederate who was their co-sectarian. Our confederates verified participants' contact information and then briefly described the first game, a simulated election with multiple rounds where participants would be asked to

[5] One participant objected to being recorded, and therefore transcripts are available for 19 of 20 group discussions.

vote for a Shia or Sunni candidate proposing to divide a fixed amount of money in different ways across different sects and that Shia or Sunni candidates were actors. Then those not randomized into the clientelism intervention were asked to return to the group, and those in the clientelism treatment were made the clientelistic offer. The confederate explained that unbeknownst to researchers a good friend and co-sectarian of the confederate (and therefore also a co-sectarian of the subject) was the actor delivering messages in the simulated elections. The broker then offered the subject $10 "in secret" to vote for the candidate of their sect across all the elections irrespective of the candidate's policy platform.[6] The script of what was said by election brokers is available in Appendix H. Subjects were free to accept or reject the money. Seventy-six percent of the participants accepted the offer; those who rejected it mostly did so stating explicitly that they wanted to hear the candidates' proposals about how the financial rewards were to be divided.

4.4 Outcome measures

The first game, designed to measure unconditional cooperation, was a simulated election and took place over the course of four rounds. In each round, two candidates—one Shia, one Sunni; both besuited men in their late 50s of similar physical appearance—appeared in person before the six-person group to deliver competing policy proposals with regards

[6] US dollars are used commonly alongside the national currency (Lebanese pounds), and all the transactions in this study took place in USD. Ten dollars is equivalent to about two hours of work on an average wage in Beirut.

to how to divide a certain amount of money between Shia and Sunni Muslims in the group. Prior to this, all participants were endowed with 40 tokens each (a single token is equivalent to $0.5), and 20 tokens were collected from every participant by way of what was described to participants as a tax. The resultant 120 tokens ($60) were placed at the center of the table; this amount was to be divided according to the proposal made by the winning candidate.

Over the course of four elections the candidates delivered competing proposals for how to divide "the tax" between the two sects. Votes were cast in writing by secret ballot. Election winners were determined at the very end of the study.[7] Instructions read out to participants by group moderators for this and all other games are available in Appendix I. Templates of decision sheets that participants had to complete in each game are in Appendix J.

The proposals delivered by election candidates are summarized in Table 3. The order of elections was fixed to facilitate the implementation of the study, and the two candidates alternated in delivering speeches. In every election, one of the candidates proposed equal distribution of tokens—20 per person—to both Shia and Sunni participants. The other candidate proposed some form of unequal distribution favoring his own sect by suggesting that either all or most of the tokens should go to his co-sectarians. The script of statements delivered from memory by the two candidates is in Appendix K. We define

[7] In the event of a tie, the election winner was determined by a coin toss.

sectarian voters as those who voted for a candidate from their sect across all four elections *irrespective of the candidate's policy proposal*.

Table 3: Candidates' messages in the election game.

	Shia candidate	*Sunni candidate*
Election 1	Sectarian distribution: Shia only (40/0)	Egalitarian distribution: (20/20)
Election 2	Egalitarian distribution (20/20)	Sectarian distribution: Sunni only (0/40)
Election 3	Sectarian distribution: Shia and Sunni (30/10)	Egalitarian distribution: (20/20)
Election 4	Egalitarian distribution (20/20)	Sectarian distribution: Shia and Sunni (10/30)

Note: In parentheses is the number of tokens offered to Shia and Sunni participants respectively.

The second game is the other-other allocation game, which is commonly used as a measure of unconditional cooperation and ingroup favoritism in an intergroup context (Chen and Li 2009; Grosskopf and Pearce 2017). In this game, participants were given 10 tokens that they had to spend and were asked to distribute these between an anonymous co-sectarian and an anonymous person from the opposite sect. Participants did not know which specific individuals in the group were affected. The allocation decision had to be recorded secretly in writing. In the other-other allocation game participants cannot benefit personally from their allocation decisions unlike in the simulated election where in the context of iterated votes it might be in participants' self-interest to support a candidate from the opposite sect. In short, the other-other allocation game provides a clean measure of altruistic regard toward the outgroup and, consequently, a clearer test of the hypothesized psychological channel behind increased cooperation.

The third and final game was the standard public goods game, which was used to measure the strength of conditional cooperation within groups.[8] In this task, participants had to decide how many of 10 tokens to keep for themselves and how many to surrender into the common pool. At the end of every round, all tokens in the common pool were multiplied by two, and the resultant sum was shared equally across all group members. Higher contributions to the common pool are reflective of a higher willingness to cooperate with others and specifically across sectarian lines. Expectations about others' willingness to contribute are central to the public goods game. We know from existing work that cooperation in heterogeneous groups is more difficult to achieve than in homogeneous settings (Balliet, Wu, and De Dreu 2014). Thus, the setting of a mixed group—mimicking a mixed society—is an especially hard case in which to expect cooperation.

The public goods game was repeated for five rounds. Before playing the game, participants had a chance to practice three hypothetical scenarios to see how the free-riding incentive operates in practice. During the game, every group member recorded their contribution in writing secretly, and the group moderator wrote out the individual

[8] The sequence of games and of elections in the first game were fixed for simplicity of implementation. Note that games' results were not revealed until the very end of the study. Fixed order of games is commonly used in experimental research where tasks are complex and a fixed sequence facilitates implementation (e.g. Chen and Li 2009; Gilligan et al. 2014).

contributions and resultant distributions on a whiteboard at the end of every round without linking individuals to specific contributions.

4.5 Compensation and Ethics

With the three games finished, participants were asked to complete a brief survey, and then earnings were tallied. One election of four and a single round of the public goods game were picked at random for the purposes of calculating the winnings from games one and three. Allocations from game two were added to these. In addition, every participant received a show-up fee of $18. Participants earned $48 on average, a little more than a day's average wage in Beirut. Those who had been assigned to the clientelism treatment and accepted the $10 reward kept that too.

We purposefully kept remuneration high in order to motivate participants to think carefully through their decisions and to try and imitate the high stakes of regular political interactions. This study involved deception for those assigned to the clientelism intervention. All those in the clientelism intervention were debriefed on the nature of deception; there were no adverse reactions in the debriefing.

5. Results

5.1. Balance

By design, there is the same number of women and men (50%) and Shia and Sunni participants (50%) across the six treatment conditions. The groups are also identical with

regards to age, wealth, and participants' ability to understand group moderator's instructions. Full results of a randomization check are in Appendix L.

In the analyses that follow, we include the following variables as controls: gender, sect, age, education, wealth, level of understanding of tasks, and type of residential neighborhood. The control variables are summarized in Table 1. We also cluster standard errors at the level of six-person groups to control for the possible confounding effect of group dynamics and include controls for group moderator effects and clientelism where appropriate. Groups are nested within sessions. In order to capture this dynamic we replicate the main analyses using multilevel modeling. All of the results from multilevel modeling analyses are reported in Appendix M; these results are consistent with ones reported in the body of the paper.

5.2. Unconditional cooperation

Table 4: Effect of Expert Appeal (panel A) and participant discussion (panel B) on likelihood of voting for a co-sectarian.

	A. Expert Appeal (Logistic)				B. Participant Discussion (Logistic)			
Model	1	2	3	4	5	6	7	8
Expert Appeal	-0.82*	-0.86*	-0.90**	-1.04**				
	(0.36)	(0.33)	(0.31)	(0.34)				
Participant Discussion					0.41	0.43	0.45	0.38
					(0.39)	(0.37)	(0.35)	(0.37)
Constant	-0.17	-0.62*	-0.10	3.77*	-0.99**	-1.61**	-1.37**	2.84
	(0.22)	(0.29)	(0.47)	(1.73)	(0.28)	(0.33)	(0.52)	(1.58)
Clientelism	N	Y	Y	Y	N	Y	Y	Y
Group moderator indicator	N	N	Y	Y	N	N	Y	Y
Control variables	N	N	N	Y	N	N	N	Y
Observations	238	238	238	212	238	238	238	205
R^2	0.03	0.06	0.10	0.14	0.01	0.06	0.10	0.13

*Note: Models 1-4: treatment effects in conditions 2 & 5 vs. baseline of 1 & 4. Models 5-8: treatment effects in conditions 3 & 6 vs. baseline of 2 & 5. Standard errors, clustered at group level, in parentheses. * p<0.05, ** p<0.01.*

The effects of expert appeal (panel A) and participant discussion (panel B) in the simulated elections are explored in Table 4.[9] The dependent variable takes on the value of 1 if a participant votes for a co-sectarian across all four elections irrespective of how the candidate proposed to divide the funds between Shia and Sunni participants.[10] The models are logistic regressions; the coefficients are non-exponentiated. For a substantive interpretation of the coefficients we computed marginal effects of varying the treatment variable from 0 to 1 while holding all other variables at their means. We expect both interventions to *decrease* the likelihood of voting for a co-sectarian candidate.

The results indicate that exposure to the expert appeal decreases sectarian voting thus increasing unconditional cooperation. Computing the marginal effects for the expert appeal intervention we find that those randomly assigned to the expert appeal treatment are less likely to vote along sectarian lines by about 19-23% (models 1 and 4 respectively). These effects are statistically significant. In contrast, participant discussion does not appear to diminish sectarian voting. If anything, the likelihood of voting for a co-sectarian after participating in a group discussion *increases* by 7-9%

[9] The number of observations in these and subsequent models is short of 240 because some participants made ambiguous markings on decision sheets. There is also missingness in analyses that include controls from survey questions.

[10] Distribution of votes by treatment group across elections is reported in Appendix N.

(models 8 and 5 respectively), although this effect is not statistically significant. Given that there are relatively few participants in this study, there might be some concern that the results in Table 4 are subject to Type-I error and are not representative of true population dynamics. To address this possibility we disaggregated election results by sect and treatment—the results are in Appendix O. The disaggregated pattern of voting is consistent with expectations: after exposure to expert appeal fewer co-sectarians vote for a candidate from their own sect who proposes to give all tokens to members of his sect.

Treatment effects in the other-other allocation game are explored in Table 5. Here, the dependent variable is the average number of tokens (of 10) allocated to an anonymous co-sectarian. The model is ordinary least squares (OLS). We expect the number of tokens given to co-sectarians to *decrease* as a result of both interventions.

Table 5: Effect of expert appeal (panel A) and participant discussion (panel B) on allocation of tokens to an anonymous co-sectarian.

	A. Expert Appeal (OLS)				B. Participant Discussion (OLS)			
Model	1	2	3	4	5	6	7	8
Expert Appeal	-1.03*	-1.03*	-1.03**	-1.11**				
	(0.42)	(0.40)	(0.36)	(0.39)				
Participant Discussion					0.41	0.41	0.40	0.49
					(0.35)	(0.34)	(0.30)	(0.30)
Constant	7.01**	6.60**	6.73**	10.88**	5.98**	5.77**	5.83**	7.98**
	(0.35)	(0.36)	(0.63)	(1.68)	(0.23)	(0.24)	(0.38)	(1.46)
Clientelism	N	Y	Y	Y	N	Y	Y	Y
Group moderator indicator	N	N	Y	Y	N	N	Y	Y
Control variables	N	N	N	Y	N	N	N	Y
Observations	238	238	238	212	239	239	239	206
R^2	0.04	0.07	0.13	0.19	0.01	0.02	0.07	0.14

*Note: Models 1-4: treatment effects in conditions 2 & 5 vs. baseline of 1 & 4. Models 5-8: treatment effects in conditions 3 & 6 vs. baseline of 2 & 5. Standard errors, clustered at group level, in parentheses. * p<0.05, ** p<0.01.*

Consistent with the findings in the election game, we find that exposure to the expert appeal diminishes the number of tokens allocated to co-sectarians by approximately one token from 7 to 6. This effect is statistically significant and suggests that the expert appeal increases altruistic regard toward the outgroup as we had hypothesized. As in the previous task, participation in a group discussion appears to *increase* allocation of tokens to co-sectarians by about 0.4 of a token, although, once again, this effect is not statistically significant. All in all, we find that exposure to the expert appeal seems to increase intersectarian cooperation in unconditional exchanges, while participation in a group discussion has no effect on this type of cooperation.

5.3. Conditional cooperation

To explore whether exposure to an expert appeal and participation in a group discussion might affect conditional cooperation in a setting involving calculations about the behavior of the opposite group we had the participants play a public goods game within their mixed-sect groups. Average contributions to the common pool remained stable across rounds at about 6 tokens (of 10); the pattern of average contributions is described in Appendix P. We hypothesized that contributions to the common pool and therefore the level of conditional cooperation will *increase* in both treatment conditions.

Table 6: Effect of expert appeal (panel A) and participant discussion (panel B) on number of tokens contributed to the common pool.

	A. Expert Appeal (OLS)				B. Participant Discussion (OLS)			
Model	1	2	3	4	5	6	7	8
Expert Appeal	0.32	0.32	0.36	0.39				
	(0.45)	(0.45)	(0.41)	(0.42)				
Participant Discussion					0.07	0.07	0.10	0.19
					(0.33)	(0.32)	(0.32)	(0.35)
Lagged group contribution	0.07**	0.07**	0.05**	0.05*	0.09**	0.09**	0.08**	0.08**
	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Constant	3.35**	3.51**	4.42**	-0.80	2.94**	3.05**	3.50**	-1.09
	(0.77)	(0.87)	(0.88)	(2.12)	(0.66)	(0.72)	(0.65)	(1.51)
Clientelism	N	Y	Y	Y	N	Y	Y	Y
Group moderator indicator	N	N	Y	Y	N	N	Y	Y
Control variables	N	N	N	Y	N	N	N	Y
Observations	960	960	960	856	960	960	960	828
R^2	0.30	0.30	0.26	0.28	0.29	0.29	0.27	0.31

*Note: Models 1-4: treatment effects in conditions 2 & 5 vs. baseline of 1 & 4. Models 5-8: treatment effects in conditions 3 & 6 vs. baseline of 2 & 5. Standard errors, clustered at group level, in parentheses. * p<0.05, ** p<0.01.*

The impact of the two interventions on conditional cooperation is explored in Table 6. The models are OLS; *Lagged group contribution* is a control for average group contribution in the preceding round of the game. Both interventions increase the amount that participants contribute, but these effects are not statistically significant across any specification. Contrary to expectations, our hypothesis about the treatments strengthening conditional cooperation is not supported by the data.

5.4. Clientelism

We hypothesized that the interventions would diminish the effectiveness of clientelism. We now test this hypothesis in the context of the simulated elections. The clientelism intervention incentivizes participants to vote for co-sectarians

irrespective of the content of the candidates' policy proposals. However, 24% of those offered the financial incentive declined to accept it. Therefore, in testing the effectiveness of vote buying we use two different measures of the treatment: whether the participant was randomly assigned to the treatment (*Clientelism*) and whether they accepted the financial incentive (*Accepted Clientelistic Offer*). It might be useful to think of this difference as that between the average treatment effect and the complier average causal effect.

Table 7: Effects of clientelism on likelihood of sectarian voting.

	A. Expert Appeal (Logistic)				B. Participant Discussion (Logistic)			
Model	1	2	3	4	5	6	7	8
Clientelism	0.94*	0.90			1.07**	1.06*		
	(0.38)	(0.53)			(0.36)	(0.52)		
Accepted Clientelistic Offer			1.53**	1.77**			1.80**	1.27*
			(0.42)	(0.62)			(0.40)	(0.56)
Expert Appeal	-1.04**	-1.10*	-1.15**	-0.92				
	(0.34)	(0.53)	(0.33)	(0.50)				
Clientelism × Expert Appeal		0.11						
		(0.74)						
Accepted Clientelistic Offer × Expert Appeal				-0.51				
				(0.87)				
Participant Discussion					0.38	0.38	0.58	0.09
					(0.37)	(0.56)	(0.36)	(0.52)
Clientelism × Participant Discussion						0.00		
						(0.74)		
Accepted Clientelistic Offer × Participant Discussion								1.11
								(0.86)
Constant	3.77*	3.81*	2.96	2.69	2.84	2.85	2.85	3.31
	(1.73)	(1.79)	(1.88)	(1.96)	(1.58)	(1.60)	(1.74)	(1.73)
Group moderator indicator	Y	Y	Y	Y	Y	Y	Y	Y
Control variables	Y	Y	Y	Y	Y	Y	Y	Y
Observations	212	212	212	212	205	205	205	205

R^2	0.14	0.14	0.19	0.19	0.13	0.13	0.19	0.20

Note: Accepted Clientelistic Offer: *coded as 0 for those not subject to clientelism and those who received the treatment but did not accept the offer. Models 1-4: data from conditions 1-2 & 4-5. Models 5-8: data from conditions 2-3 and 5-6. Standard errors, clustered at group level, in parentheses. * p<0.05, ** p<0.01.*

The effectiveness of the clientelism manipulation in encouraging sectarian voting and the impact of expert appeal and group discussion on clientelism are explored in Table 7. The dependent variable is a dummy indicating sectarian voting. The models are logistic, and the coefficients are non-exponentiated.

Consistent with expectations, we find that clientelism is highly effective at getting participants to vote along sectarian lines. Those subject to the clientelistic appeal are 21% more likely to vote for a co-sectarian (marginal effects from models 1 and 5). Among those who accept the clientelistic offer, the predicted probability of voting for a co-sectarian increases by 34% (models 3 and 7).

Whether the interventions mitigate the positive effect of clientelism on sectarian voting is tested using interaction terms between *Clientelism/Accepted Clientelistic Offer* and each of the treatments in models 2 and 4 for the expert appeal and models 6 and 8 for participant discussion. Neither of the interventions appears to diminish the effectiveness of clientelism. The interaction coefficients are consistently not statistically significant. Contrary to what we hypothesized, the results suggest that the treatments are ineffective in negating the effects of clientelism.

6. Discussion

The results give rise to two important questions. First, why does the expert appeal intervention fail to increase cooperation in exchanges that involve calculations about the likelihood of reciprocal response by the opposite group? Second, given the extensive literature on the importance of discussion for consensus building, why did participant discussion fail to increase cooperation? In addressing these questions we use observational data from a post-experimental survey and participant discussion transcripts.

6.1. Why no effect of interventions on conditional cooperation?[11]

Conditional cooperation entails calculations about the likelihood of a member of the opposite group reciprocating in response to a benevolent action. We hypothesize that inter-group trust might be a prerequisite for a positive expectation that a member of the competing group will reciprocate. We have measures of cross-sectarian trust from the survey that was administered after the games had been completed. Using survey data we constructed measures of absolute and relative trust in the opposite sect. The absolute cross-sectarian trust measure reflects a participant's trust in the opposite sect on a four-point scale, where 1 is "no trust at all" and 4 is "complete trust" (for question wording see Appendix B). The relative cross-sectarian trust measure captures how much a participant trusts the other group *relative* to her group.

[11] We re-ran the analyses in Table 6 using the Tobit model to test for the presence of ceiling effects that might be responsible for absence of treatment effects. The results are reported in Appendix Q and rule out ceiling effects.

The relative trust measure is on a seven-point scale and is calculated by subtracting a participant's trust in her own sect from her trust in the opposite sect.

Table 8. Effect of expert appeal and participant discussion on cross-sectarian trust (models 1-4), and effect of trust on contributions in a public goods game (models 5-6).

	A. Absolute cross-sectarian trust (Ordinal logistic)		B. Relative cross-sectarian trust (Ordinal logistic)		C. Contribution to public goods (OLS)	
	1	2	3	4	5	6
Expert Appeal	0.19		0.45			
	(0.31)		(0.28)			
Participant Discussion		0.42		0.15		
		(0.24)		(0.24)		
Absolute cross-sectarian trust					0.69**	
					(0.23)	
Relative cross-sectarian trust						0.46**
						(0.18)
Lagged group contribution					0.06**	0.06**
					(0.02)	(0.02)
Clientelism	Y	Y	Y	Y	Y	Y
Group moderator indicator	Y	Y	Y	Y	Y	Y
Control variables	Y	Y	Y	Y	Y	Y
Constant					-2.52	-1.12
					(1.48)	(1.58)
/cut1	-1.23	0.24	-1.50	-2.08		
	(1.60)	(1.36)	(1.50)	(1.45)		
/cut2	1.18	2.82	-0.18	-0.65		
	(1.57)	(1.31)	(1.50)	(1.41)		
/cut3	2.55	4.33	0.92	0.44		
	(1.58)	(1.31)	(1.52)	(1.40)		
/cut4			4.12	4.04		
			(1.58)	(1.48)		
/cut5			5.26	5.05		
			(1.76)	(1.63)		
/cut6			5.67	5.46		
			(1.83)	(1.69)		
Observations	176	179	173	178	1060	1044
R^2	0.05	0.04	0.03	0.02	0.32	0.31

*Note: Models 1 and 3: data from conditions 1-2 & 4-5. Models 2 and 4: data from conditions 2-3 & 5-6. Models 5 and 6: data from all experimental conditions. Standard errors, clustered at group level, in parentheses. * p<0.05, ** p<0.01.*

First, in panels A and B of Table 8 we establish whether the two interventions increased cross-sectarian trust. We analyze the effect of treatments on trust using the ordinal logistic regression. For ease of comparison, the regression in panel B is also ordinal logistic.[12] The coefficients are non-exponentiated. Exposure to the expert appeal and participation in a group discussion both appear to somewhat increase cross-sectarian trust. However, this effect is not statistically significant. These results suggest that the interventions failed to increase cross-sectarian trust.

We test for the existence of a relationship between cross-sectarian trust and conditional cooperation in Panel C of Table 8 where we regress cross-sectarian trust (absolute and relative in models 5 and 6 respectively) on the number of tokens that participants contributed to the common pool in the public goods game. The results are from an OLS regression. Consistent with expectations, we find that a one-unit increase in absolute trust (on a four-point scale) is associated with an additional contribution of 0.7 of a token (of 10). The magnitude of the correlation for relative cross-sectarian trust is similar. Both estimates are statistically significant. In short, higher levels of cross-sectarian trust do indeed tend to correlate with increased conditional cooperation across sectarian lines. Because the treatments seem to fail to increase cross-sectarian trust, they also fail to increase conditional cooperation.

[12] Results from an OLS regression are reported in Appendix R.

6.2. Heterogeneous effects of discussion quality

Results reported thus far suggest that participation in a group discussion has no effect on unconditional or conditional cooperation. This is contrary to initial expectations. One possible reason why might be that group discussions varied substantially as to their quality.

To test for possible heterogeneity in group discussion effects we constructed an index that measures their substantive depth. Substantive depth might matter insofar as discussions that lack it might fail to alter participants' preferences. A discussion depth index for the group is an aggregate of six individual discussion depth scores. A participant's position on each of 25 issues commonly brought up over the course of discussions forms the basis of her score. The 25 common issues fall into three broad categories: negative effects of sectarianism, causes of sectarianism, and ways to encourage cooperation. A participant receives 1 point if she said that the issue is relevant to cooperation; -1 if she stated that it was irrelevant; or 0 if she did not express an opinion on the issue. There is no double-counting of repeated opinions on the same issue. Coding rules for discussion depth scores are in Appendix S. Individual scores range from the low of 0 to the high of 11. Group scores are compiled by adding the individual scores of the six participants. The discussion depth scores for groups range from the lowest of 15 to the highest of 41.

Table 9: Effect of discussion depth on sectarian voting (task 1), average contribution in the other-other allocation game (task 2), and average contribution in the public goods game (task 3).

	A. Sectarian voting (Logistic)			B. Allocation to an anonymous co-sectarian in the other-other allocation game (OLS)			C. Contribution in the public goods game (OLS)		
Model	1	2	3	4	5	6	7	8	9
Group discussion score (15-41)	-0.09**	-0.13**		-0.13**	-0.14**		0.09**	0.09*	
	(0.03)	(0.04)		(0.03)	(0.03)		(0.02)	(0.04)	
Individual discussion score (0-11)			-0.29			-0.24			0.05
			(0.19)			(0.14)			(0.12)
Others' aggregate discussion score (11-36)			-0.11*			-0.12*			0.10*
			(0.05)			(0.05)			(0.05)
Female participation	-1.24**	-1.06**	-1.16**	-0.39*	-0.36	-0.40	0.50	0.86*	0.85
	(0.33)	(0.28)	(0.33)	(0.16)	(0.27)	(0.27)	(0.37)	(0.43)	(0.45)
Youth participation	0.40	-0.62	-0.38	0.48*	0.58*	0.58*	-1.00*	-1.23**	-1.23**
	(0.44)	(0.68)	(0.70)	(0.20)	(0.24)	(0.23)	(0.41)	(0.44)	(0.44)
Lagged group contribution							0.08**	0.07**	0.07**
							(0.01)	(0.02)	(0.02)
Constant	2.81**	11.80**	10.84*	9.91**	10.23**	9.71**	2.23*	-2.11	-2.33
	(1.06)	(4.52)	(4.77)	(0.61)	(1.93)	(2.20)	(1.12)	(1.78)	(2.12)
Clientelism	N	Y	Y	N	Y	Y	N	Y	Y
Group moderator indicator	Y	Y	Y	Y	Y	Y	Y	Y	Y
Control variables	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	114	95	95	114	95	95	456	380	380
R^2	0.25	0.38	0.39	0.16	0.23	0.24	0.27	0.33	0.33

*Note: All models use data from conditions 3 and 6. Standard errors, clustered at group level, in parentheses. * p<0.05, ** p<0.01.*

The results of the analyses examining the effect of discussion depth on the outcomes of the three games are presented in Table 9. The outcomes are a dummy for voting for a co-sectarian across four elections (Panel A), the number of tokens allocated to a

co-sectarian in the other-other allocation game (Panel B), and the size of contribution in the public goods game (Panel C). Panel A reports non-exponentiated coefficients from a logistic regression; OLS is used in panels B and C. We add two additional control variables that measure how active women and those under 40 were in the discussion, as existing studies suggest that women and young people might participate at lower rates (Karpowitz and Mendelberg 2014).[13]

In models 1-2, 4-5, and 7-8 we examine the effect of the overall depth of group discussion in each of the tasks. Higher quality of discussion appears to be strongly associated with more cooperation across sectarian lines. A move from the shallowest to the deepest discussion is associated with a decrease in the probability of sectarian voting of 54% in the election game (model 2), allocation of 3.6 fewer tokens to a co-sectarian in the other-other allocation game (model 5), and an increase in the contribution in the public goods game of 2.3 tokens (model 8). This result suggests that the overall ineffectiveness of the group discussion intervention masks a great deal of heterogeneity.

One possible concern about these associations is that our score of discussion depth stands in for something else, for instance, discussion tone or duration, or length of interventions by group moderators. To rule out these concerns we replicated the

[13] These measures are constructed based on moderators' evaluations of the dynamics of each group discussion.

analyses in Table 9 and included measures of discussion tone[14] (Appendix T), as well as discussion duration and the length of moderator interventions (Appendix U). Inclusion of these additional variables does not change the magnitude or significance of the reported coefficients.

Another potential concern is that, given that discussion depth was not assigned randomly, both active contribution to discussion and higher propensity to cooperate across sectarian lines might be a product of some unmeasured set of factors. For instance, in Appendix V, we demonstrate that older participants had higher individual discussion scores. The data allow us to disaggregate the effect of the depth of a participant's own contribution to group discussion from the effect of being in an especially engaging discussion. We have discussion scores for individuals and also aggregate discussion scores for the five remaining group members. If the quality of one's own contributions to group discussion is the only thing that matters in the association between discussion depth and intersectarian cooperation then the score for others' contributions to discussion should not be significant.

This is what we test for in models 3, 6, and 9 of Table 9. The results suggest that it is the quality of others' and not one's own comments that shape one's willingness to cooperate across sectarian lines. More substantive comments by other group

[14] Discussion tone is calculated by using automated sentiment analysis; negative sentiment is subtracted from positive to create the score.

members seem to be associated with more cooperative behavior by individual members of the group, and this association is statistically significant. The coefficient for an individual's discussion score is consistently in the right direction, but it never reaches statistical significance.[15]

We have to be cautious not to overinterpret these associations. Discussion depth was not varied randomly in this study, and our results are based on the analyses of only 19 group discussions. Therefore, much more work would be needed to confirm the validity of these suggestive associations.

7. Conclusion

In a laboratory-in-the-field experiment among the residents of Beirut, Lebanon, we set out to examine how to increase cooperation between Shia and Sunni Muslims. Our results indicate that exposure to a cross-sectarian consensual message calling for greater cooperation and delivered by experts can effectively increase unconditional intersectarian cooperation, i.e. exchanges in settings where no calculations about reciprocal action are

[15] Another way to tackle this concern is to compare the level of individual contributions when a participant finds herself in a group where discussion is flagging to one where there is engaging discussion. In Appendix W, we report the findings from these pairwise differences of means comparisons. We find that inactive participants consistently cooperate more across sectarian lines when they find themselves in a group where others participate actively.

necessary. Pro-cooperation appeal by experts has no effect on conditional cooperation where expectation of reciprocity is a precondition for cooperative behavior because the appeal seems to fail to increase cross-sectarian trust. Participation in a group discussion with members of the opposite sect does not appear to increase either type of cooperation in an experimental setting. We did find suggestive evidence that the effect of discussion might be highly heterogeneous. A substantive discussion touching on many aspects of intersectarian relations appears to be associated with higher levels of unconditional and conditional cooperation. Finally, we find that neither intervention succeeds in diminishing the effectiveness of clientelism.

These findings give rise to several important questions for follow-up research. First, the tentative finding on the heterogeneous effect of group discussion suggests that substantive discussions might be important, and additional research is needed to establish how discussion partners might be incentivized to engage in a more substantive exchange (e.g. Humphreys, Masters, and Sambu 2006). Second, in the aggregate, only the expert appeal has any substantive effect on increasing cooperation (only of the unconditional variety) among the Shia and Sunni Muslims in this study. This hints at a possibility that top-down attempts to engender inter-group cooperation—ones that are propagated by authoritative figures perhaps in televised debates, on the radio/podcasts, or in the press—might be more effective than bottom-up efforts to achieve the same goal (like group discussions, for instance) given that several conditions might have to be met for the

bottom-up efforts to be effective.[16] Finally, with regards to conditional cooperation, our findings indicate that cross-sectarian trust might be an important prerequisite for conditional cooperation across sectarian lines. More research is needed to establish how to augment cross-sectarian trust.

As with any experiment, a few words of caution are in order. First, with only sixty participants per intervention, this study has relatively low statistical power, and some of our conclusions must therefore be tentative. Given that for most of the non-statistically significant treatment coefficients, associated standard errors are large, we acknowledge the problem of low power but are not extremely concerned about it because the point estimates are close to zero. Furthermore, in subsidiary analyses, we also demonstrated that effects reported in regressions are also present when the data are disaggregated by sect and treatment. In settling on a study with relatively few participants per intervention we prioritized realism over statistical power by including clientelistic appeals instead of having fewer interventions with more subjects per treatment.

On a different note, there are important issues of external validity to consider. In Lebanon, Shia and Sunni Muslims are close to being numerically balanced. In contrast,

[16] In this study, the group discussion intervention combines a top-down appeal with bottom-up engagement, as group discussion follows on from watching the expert appeal. Therefore, we are not able to cleanly disentangle the effect of top-down and bottom-up interventions.

in many Middle Eastern countries, one group dominates the other numerically (e.g. Sunnis in Pakistan or Afghanistan) or politically (e.g. Sunnis in Bahrain). In future research, it would be important to explore whether the treatments have the same effect in societies where the two groups are highly unequal either in terms of relative numbers or political influence.

The primary policy implication of this project is that dissemination of a cross-sectarian appeal by experts to cooperate across group boundaries may increase intersectarian cooperation in exchanges that do not entail calculations about the likelihood of reciprocal response. These results from the laboratory suggest the need for a field intervention to test this relationship in a less controlled and more realistic setting. Our findings also suggest that a reduction in the incidence of clientelism is likely to result in greater intersectarian cooperation. Given what we know about the pervasiveness of clientelism in many young democracies (Hicken 2011), the practice of vote buying is likely to continue to pose a major impediment to voting in support of candidates from ethnic or religious groups other than one's own.

References:

Akerlof, George and Rachel Kranton. 2000. "Economics and Identity." *Quarterly Journal of Economics*, 115(3): 715-753.

Alesina, Alberto, Reza Baqir and William Easterly. 1999. "Public Goods and Ethnic Divisions." *Quarterly Journal of Economics* 114(November): 1243-84.

Allport, Gordon. 1954. *The Nature of Prejudice*. Reading, Massachusetts: Addison-Wesley Publishing Company.

Balliet, Daniel. 2010. "Communication and cooperation in social dilemmas: A meta-analytic review." *Journal of Conflict Resolution* 54(1): 39-57.

Balliet, Daniel, Junhui Wu, and Carsten KW De Dreu. 2014. "Ingroup Favoritism in Cooperation: A meta-analysis." *Psychological Bulletin* 140(6): 1556-1581.

Baytiyeh, Hoda. 2016. "Has the Educational System in Lebanon Contributed to the Growing Sectarian Divisions?" *Education and Urban Society* (April 2016): 1-14.

Cammett, Melanie and Sukriti Isaar. 2010. "Bricks and mortar clientelism: Sectarianism and the Logics of Welfare Allocation in Lebanon." World Politics 62(3): 381-421.

Cederman, Lars-Erik, Nils Weidmann, and Kristian Gleditsch. 2011. "Horizontal Inequalities and Ethnonationalist Civil War: A Global Comparison." *American Political Science Review* 105(3): 478-495.

Chen, Yan and Sherry Xin Li. 2009. "Group Identity and Social Preferences". *American Economic Review* 99: 431–57.

Corstange, Daniel. 2012. "Vote Trafficking in Lebanon." *International Journal of Middle East Studies* 44(3): 483-505.

Corstange, Daniel. 2016. *The Price of a Vote in the Middle East: Clientelism and Communal Politics in Lebanon and Yemen*. New York, NY: Cambridge University Press.

Delli Carpini, Michael, Fay Lomax Cook, and Lawrence Jacobs. 2004. "Public Deliberation, Discursive Participation, and Citizen Engagement: A Review of Empirical Literature." *Annual Review of Political Science* 7: 315-344.

Faour, Muhammad. 2007. "Religion, Demography, and Politics in Lebanon." *Middle Eastern Studies* 43(6): 909-921.

Fearon, James and David Laitin. 1996. "Explaining Interethnic Cooperation." *American Political Science Review* 90(4): 715-735.

Fishkin, James. 1997. *The Voice of the People: Public Opinion and Democracy*. New Haven, CT: Yale University Press.

Gaertner, Samuel, John Dovidio, Phyllis Anastasio, Betty Bachman, and Mary Rust. 1993. "The Common Group Identity Model: Recategorization and the Reduction of Intergroup Bias." *European Review of Social Psychology* 4(1): 1-26.

Gaertner, Samuel and John Dovidio. 2000. *Reducing Intergroup Bias: The Common Intergroup Identity Model*. New York, NY: Taylor and Francis.

Gilligan, Michael, Benjamin Pasquale, and Cyrus Samii. 2014. "Civil War and Social Cohesion: Lab-in-the-Field Evidence from Nepal." *American Journal of Political Science* 58(3): 604-619.

Grosskopf, Brit and Graeme Pearce. 2017. "Discrimination in a Deprived Neighborhood: An artefactual field experiment." *Journal of Economic Behavior and Organization* 141: 29-42.

Habyarimana, James, Macartan Humphreys, Daniel Posner, and Jeremy Weinstein. 2007. "Why Does Ethnic Diversity Undermine Public Goods Provision?" *American Political Science Review* 101(4): 709-725.
Henrich, Joseph and Francisco Gil-White. 2001. "The Evolution of Prestige: Freely Conferred Deference as a Mechanism for Enhancing the Benefits of Cultural Transmission." *Evolution and Human Behavior* 22: 165-96.
Hicken, Allen. 2011. "Clientelism." *Annual Review of Political Science* 14: 289-310.
Horowitz, Donald. 1985. *Ethnic Groups in Conflict*. Berkeley, CA: University of California Press.
Humphreys, Macartan, William Masters, and Martin Sandbu. 2006. "The Role of Leaders in Democratic Deliberations: Results from a Field Experiment in São Tomé and Príncipe." *World Politics* 58(04): 583-622.
Karpowitz, Christopher and Tali Mendelberg. 2014. *The Silent Sex: Deliberation and Institutions*. Princeton: Princeton University Press.
Kerr, Norbert and Cynthia Kaufman-Gilliland. 1994. "Communication, Commitment, and Cooperation in Social Dilemmas." Journal of Personality and Social Psychology 66(3): 513-529.
Larson, Jennifer. 2017. "Networks and Interethnic Cooperation." *Journal of Politics* 79(2): 546-559.
Lust, Ellen. 2009. "Democratization by Elections? Competitive Clientelism in the Middle East." *Journal of Democracy* 20(3): 122-135.
Miguel, Edward. 2004. "Tribe or Nation? Nation Building and Public Goods in Kenya versus Tanzania." *World Politics* 56 (April): 327-62.
Salamey, Imad and Paul Tabar. 2008. "Consociational Democracy and Urban Sustainability: Transforming the Confessional Divides in Beirut." *Ethnopolitics* 7(2-3): 239-263.
Tajfel, Henri, M. G. Billig, R. P. Bundy, and Claude Flament. 1971. "Social Categorization and Intergroup Behaviour." *European Journal of Social Psychology* 1(2): 149-178.
Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. New York, NY: Cambridge University Press.