Journal of Public Economics 122 (2015) 80–93

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Leadership and the voluntary provision of public goods: Field evidence from Bolivia

B. Kelsey Jack [a,*], María P. Recalde [b,1]

[a] Tufts University, Department of Economics, 314 Braker Hall, Medford, MA 02155, United States
[b] International Food Policy Research Institute, 2033 K St NW, Washington, DC 20006, United States

ARTICLE INFO

Article history:
Received 12 August 2013
Received in revised form 3 October 2014
Accepted 14 October 2014
Available online 5 November 2014

JEL codes:
H41
D82
O13
C93

Keywords:
Leadership by example
Sequential giving
Public goods
Voluntary contributions
Field experiment
Bolivia

ABSTRACT

We conduct a controlled field experiment in 52 communities in rural Bolivia to investigate the effect that local authorities have on voluntary public good provision. In our study, community members pool resources to provide environmental education material for local schools. We find that voluntary contributions increase when democratically elected local authorities lead by example. The results are driven by two factors: (1) authorities, like other individuals, give more when they are called upon to lead than when they give in private, and (2) high leader contributions increase the likelihood that others follow. Both effects are stronger when authorities, as compared to randomly selected community members, lead by example. We explore two underlying sources of leadership influence. First, we provide evidence that the effect of a leader's contribution is not limited to signaling the value of the public good. Second, we examine how leader characteristics affect the likelihood that others follow. Specifically, our study shows that authority influence is driven by a combination of formal leadership status, observable characteristics, and the amount that authorities contribute when they give publicly before others.

1. Introduction

Leaders play a central role in the resolution of collective action problems. Existing evidence demonstrates that leaders affect growth at the aggregate level (Jones and Olken, 2005) and influence the choice of public goods provided at the local level (Chattopadhyay and Duflo, 2004). Most studies of leadership and public good provision focus on public goods that are provided by the government.[2] In spite of the importance of voluntary contributions for the resolution of local-level collective action problems, less is known about the effect leaders have on the voluntary provision of public goods. Recent work has shown that leaders can affect voluntary contributions through informal taxation (Olken and Singhal, 2011), sanction enforcement (Grossman and Baldassarri, 2012), and reciprocity (Beekman et al., 2014). This paper examines another mechanism by which leaders may affect voluntary contributions to local public goods: leadership by example.

In a voluntary contribution setting, leadership by example arises when individuals make sequential decisions, and the choice made by the first mover (the leader) influences the contributions of others. Substantial theoretical and experimental literature has shown that first movers can affect voluntary contributions in sequential decision settings through free-riding (Varian, 1994), information signaling (Hermalin, 1998, 2007; Vesterlund, 2003; Potters et al., 2005, 2007; Andreoni, 2006), reciprocity (Andreoni et al., 2002; Meidinger and Villeval, 2003; Gächter et al., 2010, 2012) and social status (Kumru and Vesterlund, 2010; Eckel et al., 2010). Due perhaps to the challenge of empirical identification of leadership influence in field settings, no study has examined how the example set by individuals who occupy actual leadership positions affects the voluntary contributions of the groups they lead. Our paper begins to fill this gap in the literature by conducting a randomized field experiment in rural Bolivia that investigates two questions: (1) Can local leaders (authorities) affect voluntary public good provision through their example?, and (2) If so, why? Our experiment examines the effect of leadership on the contributions of both leaders and followers, and provides suggestive evidence on causal mechanisms underlying leadership influence.

* Corresponding author. Tel.: +1 617 627 3138.
E-mail addresses: kelsey.jack@tufts.edu (B.K. Jack), m.recalde@cgiar.org (M.P. Recalde).
[1] Tel.: +1 202 862 6479.

[2] For example, Chattopadhyay and Duflo (2004) study the effect of female leadership on policy decisions in India; Reinikka and Svensson (2004) investigate the political capture of public education funds in Uganda; Humphreys et al. (2006) study leadership influence on public deliberations about future public resource use in São Tomé and Príncipe; and Besley et al. (2012) analyze political influence in public resource allocation decisions in India.

http://dx.doi.org/10.1016/j.jpubeco.2014.10.003

We implement a controlled field experiment in 52 socially and politically independent communities, each of which has its own elected local authority.[3] In our experiment, a representative sample of community members pool resources to provide environmental education books for the local school.[4] We employ a between-subject design that solicits voluntary contributions in a natural decision setting and compare total contributions when an authority makes an initial public voluntary contribution – and other group members make private voluntary contributions after observing the authority's choice – to two types of controls: one in which a randomly selected community member makes an initial public contribution and one in which all contributions are private. Two of the three treatments are implemented simultaneously in each community, facilitating the use of fixed effects to address unobservable community-level confounds.

Our results show that local authorities increase average contributions when they lead by example. The effect is unique to authorities; randomly selected individuals have little effect on overall giving when they lead. We decompose treatment effects into leader and follower responses to leadership. Our results show that authorities not only increase their own contributions when they lead by example, they also influence follower contributions. Followers of authorities are more likely to make a low contribution after observing a low leader contribution and more likely to make a high contribution after observing a high leader contribution, relative to contributions in the simultaneous setting. Follower contributions are less strongly influenced by the contributions of randomly selected leaders.

We offer two pieces of suggestive evidence on why leaders affect public good provision in our setting. First, our study is designed to identify information signaling as a mechanism through which leaders influence followers. We exogenously vary whether or not participants receive information about the quality of the public good. Uninformed participants are more responsive than informed participants to the example set by community authorities, but even informed followers adjust their contributions in the direction of the leader's contribution when the authority sets the example. This result suggests that other mechanisms such as social status or reciprocity contribute to the observed influence of authority leaders. Second, we examine the relative importance of the authority's formal leadership position in the community and his or her observable characteristics. In our study, community members randomly selected to lead by example who more closely resemble local authorities have a greater influence over the contribution decisions of followers, i.e. they have the same effect on provision as do authorities in a leadership role. This finding provides suggestive evidence that authorities are influential because of their characteristics, not just the formal position they hold.

Our study is the first to examine how local authorities affect voluntary public good provision without the use of sanctions or coercion and thus makes several contributions to the literature. First, we empirically identify leadership by example as a mechanism through which local authorities can affect real voluntary contributions in a development setting. Second, we show that the leadership influence of local authorities on aggregate provision is explained both by their own contribution and the effect that they have on the contribution decisions of others. Third, we offer novel evidence on one of the most studied channels underlying leadership by example – information signaling – and show that its empirical relevance depends on who is in the leadership role. Finally, we provide suggestive evidence that authorities are influential because of their formal position in the community, their elevated contributions when placed in a leadership role, and their observable characteristics; traits such as education and wealth, which are correlated with several potential mechanisms of authority influence, matter.

Our study relates to a small but growing number of controlled field studies that examine the relationship between leaders and public good contributions in developing countries. Using public good games in the field, Grossman and Baldassarri (2012) find that individuals elected within the experiment – who are not local authorities – are more effective at sanctioning low voluntary contributions, while Beekman et al. (2014) show that voluntary contributions are lower in communities that have corrupt officials. Kosfeld and Rustagi (forthcoming) study how the sanctioning style of leaders, elicited through a third party punishment game, is correlated with local forest conservation outcomes. More similar to our study, d'Adda (2012) conducts an artefactual field experiment in 8 villages in rural Colombia that investigates how social information interacts with social status, defined endogenously along leadership dimensions, in a repeated voluntary contribution setting. Her results show that high status individuals (leaders elected within the experiment) are more likely to make high contributions and are less influenced by the contribution decisions of others. Our study is unique in this literature in that we study actual authorities, vary leadership exogenously, use a one-shot decision setting in which both leaders and followers can react to leadership, and study voluntary contributions to an actual good that benefits the community.

In trading off the control of the laboratory for the richness of the field, our study makes some compromises. First, in order to investigate both leader and follower responses to leadership, we allow leader contributions to arise endogenously in our experiment. This design feature reveals whether or not authorities take advantage of leadership opportunities, but prevents us from cleanly separating the effect of leader contributions from leader characteristics and leader type when analyzing follower responses. Second, a small number of communities could not comply with treatment randomization for idiosyncratic reasons. Our findings are robust to correcting for any resulting selection bias. Third, like most field studies, our results are specific to a particular setting at a particular point in time. While Bolivia's unique decentralized local governance arrangement allows us to examine the effect of local authorities on voluntary public good provision, other types of leaders may be more influential in other contexts. The fact that we find suggestive evidence that leadership by example operates through multiple channels increases the likelihood that our results are relevant to other settings where different types of opinion leaders exist and where one or more channels of leadership influence are relevant.

The paper proceeds as follows. Section 2 offers a conceptual framework for leadership in voluntary contributions. Section 3 describes the experimental context and design. Section 4 describes the main results, treatment heterogeneities, and robustness checks. Section 5 concludes.

2. Conceptual framework

Early theoretical literature on sequential giving showed that leadership by example is weakly detrimental for voluntary public good provision when information is perfect and individuals are solely motivated by altruism (Varian, 1994). This result emerges because the positive externalities generated by voluntary contributions introduce a free-riding incentive that induces first movers, leaders, to make low initial contributions that force followers to provide the public good. Subsequent theoretical and experimental studies have, in contrast, shown that sequential giving can be beneficial for public good provision. Three primary classes of mechanisms underlie these positive results: (1) information signaling, (2) reciprocity, and (3) social status.

First, models of information signaling have shown that sequential giving can have beneficial effects for voluntary public good provision when the common value of the public good is uncertain. If the leader

[3] We refer to the elected local leader as the "authority" to differentiate the leadership role assigned in the experiment from the formal authority position elected local leaders occupy at the community level.

[4] Environmental education books provided through the experiment are accessible to all community members (non-excludable), but exhibit rivalry. Because contributions exert a positive externality that is non-rival and non-excludable on anyone who cares about the provision of education material in the local school, we consider them a pure public good from the donor's perspective.

has an informational advantage over others, he or she can signal if the public good is of high (low) value by making a high (low) contribution that induces others to follow (Hermalin, 1998, 2007; Vesterlund, 2003; Potters et al., 2005; Andreoni, 2006).[5] Although level predictions are conditional on the information that the leader possesses, information signaling is always welfare enhancing in this setting. In our field setting, authorities may have superior information about the value of the public good. They may have been elected precisely because of this informational advantage, or may have acquired information through their formal leadership role.[6] If individuals have independent private valuations for the public good that are not common knowledge, then strategic uncertainty limits free-riding and leadership by example can also increase public good provision (Bag and Roy, 2011).

Second, sequential giving can positively affect public good provision when individuals have reciprocity, equity and/or fairness concerns (Meidinger and Villeval, 2003; Huck and Rey-Biel, 2006; Potters et al., 2005; Potters et al., 2007; Gächter et al., 2012). Leaders who make high contributions crowd-in the contributions of others. They cannot use their first mover advantage to free-ride (as in Varian, 1994) because reciprocal followers punish free-riding at a cost.[7] In our field setting, authorities may generate more reciprocity among community members due to their authority position. They may also cause a reduction in strategic uncertainty when coordination incentives are present, generated by reciprocity and/or fairness concerns (Bicchieri, 2005).

Third, social status can affect public good provision when individuals with high social status lead and followers like to be associated with high status others (Kumru and Vesterlund, 2010) or want to acquire status themselves (Bracha et al., 2009).[8] Status can also be modeled as the location individuals occupy in a social network (center vs. periphery), and explain leadership influence on voluntary contribution decisions through the number of agents that observe the leader's choice (Eckel et al., 2010). In a field setting, authorities may have higher social status than the average community member because they are wealthier, more educated, and even possess higher social status as a direct result of the formal leadership position they occupy in the community. Observable characteristics associated with status may therefore provide a clue as to the importance of status in the field.

Any or all of the mechanisms described above may be active when any individual leads by example. Their importance, however, may be amplified when a local authority leads. To empirically test this assertion we conduct an experiment that varies whether or not a local authority or a random individual leads by example. Using the private, simultaneous contribution setting as the benchmark for comparison, we test whether Varian's (1994) free-riding predictions hold. We do not develop a theoretical model capable of providing specific predictions by leader type and mechanism because the variety of channels described in this section may coexist simultaneously in our study setting.[9] We rely on the existing literature to provide intuition as to how each mechanism works and focus instead on testing for the first time how leadership by example affects voluntary public good provision in a field setting when actual leaders take on a leadership role.

2.1. Hypotheses

Our study is designed to test four groups of benchmark hypotheses derived after the pure altruism model of voluntary public good provision (Varian, 1994). Rejection of the null hypotheses is consistent with all channels underlying positive leadership by example effects (information signaling, reciprocity, and social status). Hypothesis 4 is the exception, which refers specifically to information signaling. We generate the variation necessary to test Hypothesis 4 by randomizing access to information about the public good among followers in our experimental set up, which is described in greater detail in Section 3.

Hypothesis 1. Leadership by example (sequential giving) weakly decreases total contributions relative to the simultaneous contribution setting.

Null Hypothesis 1 establishes the main result of Varian (1994): sequential giving is weakly detrimental for public good provision because of free-riding. The alternative hypothesis is that sequential giving increases total contributions relative to the simultaneous contribution setting, particularly when authorities give first. This is consistent with any or all of the underlying mechanisms discussed in this section. Free-riding incentives need not be absent, only outweighed by the other mechanisms.

Hypothesis 2. Individuals give less when they lead by example (give publicly before others), than when they give in a private simultaneous setting.

Null Hypothesis 2 predicts that contribution leaders (first movers) free-ride off of followers. The alternative hypothesis is that individuals make higher contributions when they give publicly before others than when they give in a private simultaneous setting. Specifically, we expect authority contribution leaders to give more than randomly selected contribution leaders both because they recognize the positive influence that they can exert on others and because more of the mechanisms that make contribution leaders influential may be relevant when a local authority leads.

Hypothesis 3. Follower contributions are decreasing in the contributions made by contribution leaders. They decrease (increase) with leadership by example, relative to the simultaneous contribution setting, when leader contributions increase (decrease).

Null Hypothesis 3 establishes free-riding. The alternative to Hypothesis 3 predicts that the best response function of followers is flat or upward sloping. We expect followers to positively respond to the contributions made by first movers, and more so to the contributions made by authorities, because authorities likely amplify the mechanisms that underlie positive leadership by example effects.

Hypothesis 4. Information does not affect follower response to first mover contributions.

Null Hypothesis 4 establishes the straw man hypothesis that leader contributions do not convey information. The alternative is that leader

[5] Field studies have investigated information signaling within the context of charitable contributions. For example, Karlan and List (2012) examine information signaling within the context of matching grants in charitable giving, while Smith et al. (forthcoming) investigate peer effects in charity fundraising conducted by individuals on-line. The form of leadership we study differs from these field studies in that we require the leaders to set an example by making a one-time costly and unrecoverable contribution before others; the leader does not observe the decisions of others and cannot make contributions at a later time.

[6] Miller and Mobarak (forthcoming) study this informational channel of leadership influence within the context of opinion leader influence on technology adoption decisions in Bangladesh. They show that when opinion leaders (including local authorities) unanimously decide to accept (reject) the adoption of a new technology, the likelihood of adoption by other community members increases (decreases).

[7] Andreoni et al. (2002) and Gächter et al. (2010) compare simultaneous and sequential giving in linear public good games in the laboratory when information is perfect and show that although leaders try to free-ride off of followers, followers punish free-riding by giving less than their best response function predicts.

[8] Related to this literature is also the empirical work on prestige and visibility motives for giving. See, for example, Harbaugh (1998), Ariely et al. (2009) and Karlan and McConnell (2014).

[9] Deriving specific predictions from a model that incorporates all mechanisms of leadership influence and distinguishes between different types of leaders requires strong assumptions about preferences and about leader characteristics, and is made complex by the multiplicity of different channels at work in our setting. Our experimental design does not generate the variation necessary to distinguish between mechanisms, so including such a model would do little to aid in the interpretation of our results.

contributions do convey information and that followers are less responsive to the example set by contribution leaders the more information they possess about the public good.

3. Experimental design

We employ a between-subject design with three treatments that (a) identifies the effect authorities have on voluntary public good provision when they lead by example, (b) distinguishes the influence of authorities from that of non-authorities in the community, and (c) incorporates additional variation to explore the importance of one of the most studied mechanisms behind leadership by example: information signaling. Before turning to the design of the experiment and its implementation, we describe the study setting, which informs our design.

3.1. Study setting

The experiment was conducted in 52 communities located in the Rio Grande–Valles Cruceños region of Bolivia, in collaboration with a non-governmental organization, Fundación Natura. The setting is useful for the study of leadership by example in public good provision for three reasons. First, decentralization in Bolivia extends all the way to local level administrative units called Organizaciones Territoriales de Base (OTBs). OTBs are independent social and political units; in our study setting they are small in size (the average OTB contains 27 households), meet regularly as a group, and are poorly integrated with outside markets. Each OTB has an elected representative (OTB president) who serves as the formal authority in the community. OTB presidents are elected in public meetings through majority vote. There is no clear term limit but most remain in office for a couple of years. The position is unpaid. OTB presidents are in charge of requesting funds from the municipal government, of developing local projects, of interacting with outsiders, and of organizing collective work. The fact that these authorities exist in all communities allows us to use a single, pre-existing definition of local leader, rather than relying on subjective approaches that might vary from community to community. Although other leadership positions exist in the OTB, including vice president, treasurer, and secretary, none has the same level of importance as the OTB president. We use the term OTB interchangeably with community in the remainder of the paper because there is perfect overlap between OTBs and communities in our sample.

Second, political parties and organizations have little presence in the study area.[10] Anecdotal evidence indicates that OTB presidents do not actively seek office and have no intention of pursuing a political career. They accept the authority position when selected by their peers, but find the responsibility costly in terms of the effort and time. This helps mitigate political factors that might confound our experimental design, but may also increase the likelihood of observing leadership by example effects in our setting. Given the paucity of field evidence on the topic, we consider this to be an advantage.

Third, a detailed census of 130 communities was conducted in the area by Fundación Natura in 2010. The census includes information that facilitated the randomization of communities and households into treatments, and provides us with most of the controls used in our analysis of experimental results. Study communities are spread over five municipalities across a very rural 7400 km^2 area. Though the study area is culturally and socioeconomically fairly homogenous, some communities are considerably more remote than others.

The experimental design uses a naturally occurring decision setting – a community meeting – to solicit contributions to environmental education books for the local school. Community meetings occur regularly in our study setting, and are organized through local authorities to address community business or at the request of outside individuals or organizations. We followed the standard local practices that community outsiders employ to organize meetings, to make the decision setting as natural as possible.

Environmental education books were chosen as the public good in our experiment for several reasons. First, all communities in our study have a local primary school, to which books could be donated. Second, 40% of households in our study site identified environmental protection as one of the top values that should be taught to children. All communities are located inside a watershed that was declared protected in 2007. Environmental issues such as trash, water pollution, and soil erosion are thus very salient in the area. Third, although environmental education books are not a pure public good, they exhibit several relevant characteristics: books are non-excludable and generate social spillovers.[11] We consider them a pure public good from the donor's perspective. Specifically, donations generate a positive externality that is non-rival and non-excludable on anyone who cares about the provision of education material in the local school.[12]

From a practical perspective, books made it possible for us to examine voluntary contributions to a local public good in a setting in which even small contributions could ensure positive levels of provision. They also minimized trust confounds by allowing us to deliver the good acquired by the contributions at the end of the experiment. The books used in the experiment were purchased from a non-government organization that specializes in producing environmental education material in Bolivia. Seven different books were available, and were sold at a zero-profit price of 10 Bs. per book.

3.2. Treatments

Our experiment employs a between-subject design with three treatments. In each treatment, subjects complete a survey in exchange for money and are subsequently given the opportunity to make a voluntary contribution to environmental education books for the local school.

The treatments vary the way in which community members make voluntary contribution decisions. In a No Leader Treatment (NL), individuals make private simultaneous contributions. In a Random Leader Treatment (RL), a randomly selected individual is asked to make his or her voluntary contribution publicly before others. In an Authority Leader Treatment (AL), the community authority (OTB president) is asked to make his or her voluntary contribution publicly first. We refer to these first movers as contribution leaders. In both the RL and AL treatments, other participants make private voluntary contribution decisions after observing the contribution leader's public choice. The sequential decision setting we employ follows the laboratory literature, which typically uses a 2-person public good game with induced preferences to study leadership by example (see, for example, Potters et al., 2005, 2007; Kumru and Vesterlund, 2010; Gächter et al., 2010). Our experiment uses the same setup, in which an exogenously chosen first mover makes a public contribution decision, but with a group of size n. While it does not make a difference whether follower contributions are private or public in the 2-person game employed in the laboratory, follower contributions are private in our experiment to provide a test of leadership by example that is clean of sequential considerations introduced by public follower contributions.

[10] Sixty percent of households in our study (located in 51 communities) indicated in the 2010 census that political syndicates do not exist at the OTB level. Six out of 41 local authorities indicated that they attend political syndicate meetings, but none indicated that they have occupied authority positions at the syndicate level in the past.

[11] All community members can access the education material available in the local school. Teachers are present in the school during the week and can grant school access. Community authorities and members of the parent–teacher association also have keys to the school.

[12] Focus groups were conducted at the pilot stage in communities close to but not included our study sample to select the public good used in the experiment and to assess its perceived quality. Qualitative evidence gathered at the pilot stage showed that most participants perceived the books to be of high quality, address an important topic, and appeal to children.

The NL treatment establishes a benchmark scenario that we use as a control in our experiment. Comparison of NL and AL determines if local authorities affect voluntary public good provision through their example. Comparison of RL and AL determines whether AL treatment effects are specific to authority leaders or are a generic response to leadership. We conduct two simultaneous treatments per community and use fixed effects to control for community characteristics that may affect both leader and follower contributions. Whenever AL is conducted in a community, the authority is assigned to the AL treatment. Whenever the AL treatment is not conducted in a community, the authority is assigned to the NL treatment. No authority is thus ever present in the RL treatment.

The design introduces an *information manipulation* in all treatments that gives half of all participants, and always the contribution leader, the opportunity to inspect the public good before making a voluntary contribution decision.[13] If the leader's contribution conveys information about the value or quality of the public good then uninformed follower contributions should move in the direction of the leader's contribution (Vesterlund, 2003; Potters et al., 2005, 2007; Andreoni, 2006). While visual inspection may not resolve all uncertainty about the quality of the public good, the effect of the leader's contribution should be muted or reversed for informed followers, who rely less on the leader's contribution to update their beliefs about the quality of the public good. Comparison of uninformed and informed follower decisions within treatment isolates the effect of information signaling. Comparison across treatments informs us about its differential importance by contribution leader type.

3.3. Implementation

Table 1 summarizes the stages of implementation and refers the reader to the relevant Appendix section where the translated (English language) script is provided. Data collection and study implementation occurred over several stages, beginning with a census of the region implemented by the study partner in 2010. OTBs included in the census but missing OTB-level information or without a local primary school were excluded from our study, as were communities smaller than 15 or greater than 80 households in size. The final eligible sample consisted of 52 OTBs.

OTBs were randomly assigned to one of three possible pairwise combinations of the NL, RL and AL treatments and 12 households from each community were randomly sampled for participation, using the min max T procedure (Bruhn and McKenzie, 2009) and testing balance on both OTB and household variables from the 2010 census.[14] The randomization thus delivered both balanced characteristics across treatment and a representative sample of households for participation in the study.

The implementation of the experiment began with a visit by the study team to each community 4 to 7 days prior to the intervention. The team met with community leaders, scheduled the community meeting, and delivered invitations worded to minimize experimenter demand effects that might arise if the meeting was organized by the partner NGO (see Appendices B.1 through B.3). Specifically, written invitations were delivered in person to the heads of the 12 households selected through the randomization process, which always included the OTB president. At the time of invitation, individuals were told that they could earn up to 45 Bs. for attending the meeting and that only one person per household could attend.[15] On the day of the experiment, invited households were reminded of the time and location of the meeting (see Appendix B.4). At all stages of the invitation process, unavailable households were replaced using a list of alternates generated through the original randomization.[16]

Two types of attrition affect the final study sample. The first is selection into the study, which occurred before the experimental session was conducted and does not affect the internal validity of our results. Appendix Table A.1 provides a description of how the sample of households selected through randomization differs from the final sample of participants. The second is selection into treatment, which occurred in 6 sessions assigned to the AL treatment, where the authority was not present on the day of the experiment. These sessions were run with either the NL or RL treatment (whichever was not assigned to the other concurrent session in the same community). In four cases, to preserve balance in the number of sessions across treatments, the next available community scheduled to receive a combination of RL and NL treatments received an unplanned AL treatment. These incidents are documented in Appendix Table A.2.

Selection into treatment is potentially problematic for our empirical strategy. Authority absences, however, appear to have been idiosyncratic. Authorities were not present the day of the experiment because they had to attend classes in the municipal capital, had to take care of medical emergencies, or were away harvesting crops. No systematic differences between sessions selected into and out of treatment are detected in our data (see Appendix Table A.2). We show balance on the characteristics of participants for the final experimental sample in Appendix Table A.3. The AL treatment has fewer females, and participants in the NL treatment have slightly fewer assets, measured as a raw sum of durables owned by the household. These results persist when community fixed effects are included and are partly driven by the fact that authorities differ from other community members on a number of observable characteristics (described in Appendix Table A.4), and are more often present in the AL treatment. We control for these and other observable characteristics throughout our analysis of experimental results.

Each experimental session consisted of three parts and took place at the local school or in another centrally located community building. Throughout the implementation, efforts were made to keep the process similar to a typical community meeting. Appendix B.5 provides the script used to conduct the session.

In Part 1 of the meeting, individuals arrived to the designated meeting place, registered, received an envelope, an ID number, and consent forms. IDs ranging from 0 to 11 were distributed at random to participants with the exception of ID 0, which was always given to the OTB authority. The subject that was randomly assigned ID 6 acted as contribution leader in the RL treatment. Subjects with ID 0 and 6 were not aware that they would have a special role in the experiment. All subjects were informed that they would earn 35 Bs. by completing a questionnaire and 10 Bs. by attending the full meeting. At the time of soliciting consent, subjects knew that they would be asked survey questions but were not aware that they would be asked to make a voluntary contribution. Part 1 of the experiment took approximately 20 min.

In Part 2, subjects were split into two groups based on their ID number, which allowed more seating space for each of the participants while also facilitating the implementation of two simultaneous treatments. The experimenter and assistant were rotated to ensure balance across

[13] The information manipulation is implemented in such a way that informed agents know who else is able to inspect the public good, but uninformed agents have no knowledge of the informational advantage possessed by others.

[14] Randomization was balanced at the household level on: the number of rooms in a house, the education of the household head, the number of children under 16 per household, a stated preference for instilling environmental values in children, perceptions of community cooperation and decision-making, attitudes toward outsiders and participation in past community meetings. The distance to market, the number of households in the community, and the municipality were all balanced at the community level. Most balancing variables are included as controls. In the analysis, a few variables are replaced with superior measures of the underlying characteristic of interest, such as the use of household assets measured by the raw sum of durables owned by the household instead of number of rooms as a proxy for wealth. Analysis that includes households that did not participate in the experiment relies exclusively on the balance variables.

[15] 45 Bs. is approximately 6.50 US dollars and is equal to the daily wage for agricultural work in the study setting.

[16] In select cases, no alternates from the list were available and substitutions were based on convenience. Convenience replacements were made in 19 cases.

Table 1
Study implementation timeline.

Step	Timing relative to session	Stage	Details	Supporting documentation
1	1 year before	Census	Census of 120 communities (>3000 households) conducted by NGO partner as part of a long run project conducted in the area.	
2	1 to 3 months before	Randomization	Random sample of households selected using census data and pairwise combination of treatments assigned to communities.	
3	1 week before	Invitation	Community meeting organized in each OTB with the approval of OTB representatives. Sampled households delivered invitations to attend a community meeting. Appointment made with school teacher to conduct a survey the day of the intervention.	Appendices B.l, B.2, B.3
4	Day of	Reminder	Invited households reminded about the time and location of the community meeting. When an invited household was unable to attend an alternative household selected through randomization procedures was invited to participate in the study.	Appendix B.4
5	Day of	Session	Community meeting, including substages (single script)	Appendix B.5
		1. Registration	Participants arrive at the community meeting, receive an identification number, and give consent to participate in the study.	
		2. Session split	Meeting randomly split into two sessions using even and odd identification numbers.	
		3. Questionnaire	Participants in each session complete a 17 question survey in exchange for 35 Bs.	Appendix C.l
		4. Questionnaire payment	Payment of 35 Bs. given to participants who complete the survey, in 7 units of 5 Bs. coins.	
		5. Information manipulation	Participants with even ID numbers are asked to step outside of the room and are given the opportunity to inspect environmental education books.	
		6. Contribution decisions	Participants are given the opportunity to make voluntary contributions to acquire environmental education books for the local school.	
		7. Post-decision questionnaire	6 question survey administered to participants after all donations are made.	Appendix C.2
		8. Delivery of show up fee	10 Bs. show up fee given to all participants in an envelope.	
		9. Reconvene	Participants are taken back to a single room, where the community meeting is reconvened.	
		10. Delivery of books	The total amount contributed by participants in both sessions is announced and the corresponding number of books is delivered publicly.	
		11. Conclusion	Participants leave the room.	
6	Day of	Teacher survey	Survey administered to a teacher who preferably did not participate in the meeting or had a household member participate.	Appendix C.3

treatments. In each group, subjects completed a survey containing questions unrelated to the study (see Appendix C.1), in exchange for their experimental earnings. Questions were read out loud to participants, who answered using paper and pencil.

Regardless of the answers provided, all subjects were given 35 Bs. in 5 Bs. coins upon completion of the survey. Participants with even ID numbers were then asked to step out of the room; contribution leaders always had even ID numbers. Even numbered subjects were shown the environmental education books and given the opportunity to inspect them, but were not told how the books would be used in the session. Subjects with odd ID numbers were not told the purpose of this interruption, and were asked to answer one additional survey question to pass the time.[17] Participants with even IDs returned to the room after 5 min.

Following the information manipulation, the contribution decision was presented to subjects. Subjects were told that the money earned by completing the survey was theirs to keep and that they could contribute as much or as little as they wanted to the environmental education books for the local school. Books were displayed in front of the room and subjects were given general information about their cost and content. They were informed that for every 10 Bs. contributed by all community members (in both sessions of the experiment) the school would receive one book.[18] Participants knew that 7 different volumes of the books were available and that they would be delivered on-site at the end of the experiment.

To make their voluntary contributions, subjects were asked to place the money they wished to contribute in an envelope that had their ID number marked on the inside. Contribution decisions were done in private behind a cardboard partition. If the session was assigned the RL or the AL treatment, the contribution leader – referred to by his or her ID number – was asked to publicly announce the amount of his or her contribution as it was placed in the envelope. All other participants were called one by one to make their private voluntary contribution in the back of the room behind the cardboard partition. Complete confidentiality was stressed to ensure that subjects would not be concerned about sanctioning or other social rewards or punishment. The order by which subjects were called upon to make their contributions depended on the seating arrangement. Participants were not allowed to talk while contribution decisions were being made.

After all participants made their contributions, subjects were asked to complete a survey with 6 questions on household socio-demographics and perceptions of teaching quality in the local school (see Appendix C.2). The purpose of these questions was to collect individual-level information that was not available through the census or was outdated.[19] Once the final survey was completed, subjects received a 10 Bs. show up fee. This marked the conclusion of Part 2, which took approximately 60 min.

Part 3 of the experiment started once both experimental sessions were over. All participants returned to the same room and the total amount contributed by subjects was announced. The environmental

[17] Subjects with odd ID numbers observed those with even numbered IDs leave the room, but none asked about the purpose of the interruption and no explicit explanation was provided. Given that the subjects had already been split into two groups based on ID number, further amendments to the group structure based on ID number may not have been surprising at this stage of the implementation.

[18] Participants were additionally informed that contributions would be rounded up if the total amount contributed by all participants was not a multiple of 10. This ensured that we never kept any of the contributions made by subjects.

[19] Note that the census occurred one year prior to the experiment. We chose not to re-collect substantial socio-demographic information to avoid imposing an additional time burden on participants, and chose to gather data on key time-varying questions that were relevant to our experiment instead.

education books were counted in public and given to the community authority or school representative in front of all subjects. The final part of the experiment took approximately 10 min. The entire session lasted between 90 and 120 min.

4. Results

We observe the decisions of 580 subjects in 104 sessions of the experiment, which were conducted between May and July 2011. Each session included between 4 and 6 subjects; a total of 9 to 12 individuals participated in the experiment in each of the 52 communities included in our sample. Table 2 summarizes these and other session characteristics, our primary outcome variables, and the set of variables that we use as controls throughout the analysis. We control for a number of factors that may affect contributions both of leaders and of followers, including past leadership roles in the OTB, whether the individual is a teacher (12% of OTB presidents are also teachers), the number of children in the individual's household that attend the local school, and whether the individual was exposed to the information manipulation.

We begin by describing the raw contribution data before outlining our empirical strategy. Fig. 1 shows histograms of contribution decisions by participant type and experimental treatment.[20] Panel (a) combines all participants and shows that 5 and 10 Bs. are the most common contribution levels across treatments. A contribution of 5 Bs. is the median contribution in the study population and the minimum non-zero

Table 2
Summary statistics.

	Mean	Std. dev.	Min.	Max.	N
	(1)	(2)	(3)	(4)	(5)
Outcome variables					
Contribution	8.236	5.830	0	40	580
Contribution > 0 Bs.	0.472	0.500	0	1	580
Contribution ≥ 10 Bs.	0.941	0.235	0	1	580
Total contributions	45.930	16.500	10	120	104
Contribution: NL treatment	7.767	6.117	0	35	188
Contribution: RL treatment	7.960	5.106	0	35	201
Contribution: AL treatment	8.987	6.202	0	40	191
Contribution: leaders, RL	10.556	4.595	5	20	36
Contribution: leaders, AL	14.279	8.344	5	40	34
Contribution: followers, RL	7.394	5.048	0	35	165
Contribution: followers, AL	7.841	4.969	0	35	157
Control variables					
Individual received information	0.529	0.500	0	1	580
Female	0.291	0.455	0	1	580
Years of education	4.733	3.489	0	20	580
Household assets (sum of durables)	1.972	1.551	0	9	580
Num. children attending local school	0.640	1.074	0	10	580
Evaluated teacher as good or excellent	0.607	0.489	0	1	580
Caring for environment is top value	0.395	0.489	0	1	580
Participated in all OTB meetings	0.366	0.482	0	1	580
Participates in OTB projects	0.659	0.475	0	1	580
Always agrees with community decisions	0.672	0.470	0	1	580
Always trusts NGOs	0.407	0.492	0	1	580
Held past OTB leadership position	0.103	0.305	0	1	580
Teacher	0.038	0.191	0	1	580
Experimenter indicator	0.500	0.502	0	1	104
Session size	5.577	0.516	4	6	104
Share of session female	0.287	0.192	0	0.833	104
Other variables					
Community size	26.803	13.042	15	75	52
Travel time to nearest market (min)	171.442	131.590	15	720	52
Pupils provide their own books	0.635	0.486	0	1	52

Notes: Variables are measured at the individual level unless otherwise noted. Contributions, experimenter indicator, session size, and share of session female are measured at the session level. Community size, travel time to nearest market and pupils provide their own books are measured at the community level. All variables listed under the control variables heading are the standard set of controls used in the analysis.

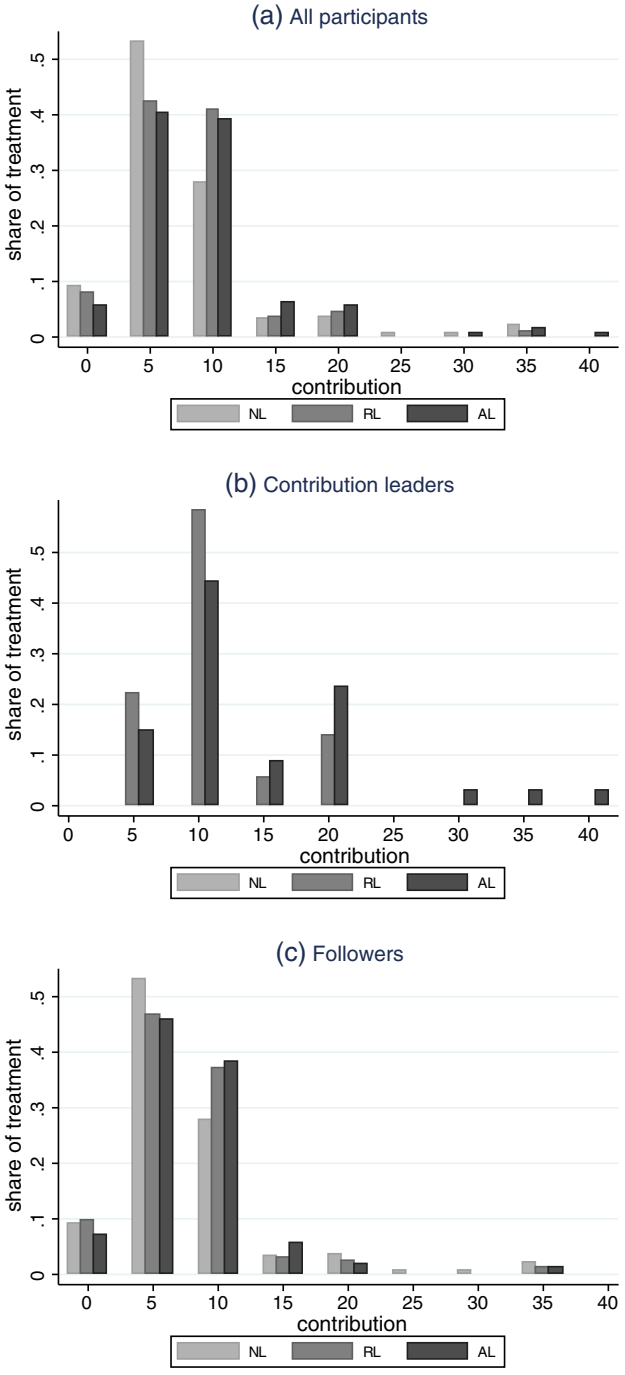

Fig. 1. Histogram of contributions by participant type and treatment. Notes: Figures represent histograms of contributions in bins of 5 Bs. Each of the figures describes contributions for different samples of participants: all (top), contribution leaders only (middle) and followers only (including the NL treatment, bottom). The shading describes each of the three experimental treatments, and the histogram plots the share of each treatment in the different contribution bins.

[20] Contributions are classified in 5 Bs. bins that reflect the monetary unit used to pay subjects in the experiment. A small number of subjects made contributions using their own coins. These are rounded to the closest 5 Bs. interval in Fig. 1 but not in the remainder of the analysis.

contribution level. Ten Bs., on the other hand, is the smallest contribution that has a direct impact on the provision of environmental education books.[21] The AL and RL treatments show a pattern of first order stochastic dominance over the NL treatment. Differences between the AL and RL treatments are statistically insignificant, but suggest that authorities increase provision more than random individuals when they lead by example.[22] Panels (b) and (c) break contributions down by participant type. This presentation of raw data suggests that authority leaders give more than individuals randomly selected to lead by example (Panel b), and that contribution leaders make higher contributions than followers (Panel b vs. c).[23]

Although the raw data suggests that contributions increase with leadership by example, these results ignore observable and unobservable factors that drive both leader and follower contributions. To quantify treatment effects parametrically we regress contributions on the experimental treatments and include community fixed effects. Community fixed effects remove any spurious correlation between leader and follower giving, driven by unobservable community-level factors. Even with fixed effects, we include the individual- and session-level controls defined in Table 2, to address the slight imbalance of observable characteristics discussed in Section 3.

We estimate treatment effects using three different specifications, which we describe in general here and in detail immediately preceding each set of results. First, to examine effects on average contributions, we use a continuous measure of contributions and estimate treatment effects using ordinary least squares. Second, we take into account that the experiment was implemented using 5 Bs. coins and use an ordered logit model in which each 5 Bs. interval constitutes a separate categorical giving bin.[24] Third, given that the greatest mass of contributions occurs at 5 and 10 Bs. (see Fig. 1) treatment effects may be concentrated around the median level of giving. We therefore estimate treatment effects on the probability of giving above the median ($\geq$10 Bs). We revert to OLS for the median regressions, though our results are qualitatively similar if we use a conditional logit model with fixed effects. All tables show OLS and ordered logit estimates, include fixed effects, and show results with and without controls. In our discussion of the results, we describe point estimates from the specifications with controls.

In our main analyses, we assume that the selection documented in the implementation section is idiosyncratic. Robustness checks that address selection into treatment, including an instrumental variable correction and a sample restriction to the compliant sub-sample of communities, are presented in Section 3. They are consistent with our main results.

[21] Note that there are two possible focal points for contributions in our study: 10 Bs., the minimum individual contribution that changes the supply of the public good directly, and 70 Bs., the aggregate contribution level that supplies a full set of books. Ten Bs. is also the value of the show-up fee. The modal contribution of 5 Bs. in the NL treatment suggests that neither focal point is particularly salient for participants, since it under-provides relative to both.

[22] Two-sided Wilcoxon–Mann–Whitney rank sum tests reject the null hypothesis that the contributions of all participants in AL and NL or RL and NL were drawn from the same underlying distribution ($p < 0.01$ and $p < 0.10$ respectively). Differences between AL and RL are not statistically significant; though the two-sided Wilcoxon–Mann–Whitney rank sum test has a p-value < 0.15.

[23] Two-sided Wilcoxon–Mann–Whitney rank sum tests provide p-values < 0.05 and <0.01 respectively.

[24] To accommodate community-level fixed effects in the ordered logit specification, we use Baetschmann et al.'s (2011) "blow up and cluster" (BUC) approach, which generates dichotomous outcomes at each of k thresholds and estimates each binary outcome using conditional maximum likelihood. The method relies on the restriction that the log odds associated with each threshold is the same, but is shown to be robust to outcome categories with few observations. In our data, the dependent variable acquires values $k \in \{0, 5, 10, 15, 20\}$ when contributions fall respectively in (0, 5), (5, 10), (10, 15), (15, 20) and (20,40). Contributions of 20 Bs. or more are grouped together because few observations exceed the value of 20, and no random leader makes contributions in excess of this amount (see Fig. 1).

4.1. Main results

We begin by analyzing the effect of leadership by example on total and individual contributions, then split the analysis to focus on the behavior of leaders and on the response of followers.

4.1.1. Total contributions

To test Hypothesis 1, the weak negative effect of leadership by example predicted by free-riding, we estimate:

$$y_{ic} = \alpha + \beta_1 AL_{ic} + \beta_2 RL_{ic} + \Gamma X_{ic} + \eta_c + \epsilon_{ic} \quad (1)$$

where y_{ic} represents the contribution made by individual i in community c, AL denotes the Authority Leader treatment, RL the Random Leader treatment, X_{ic} is the vector of individual- and session-level controls, η_c is a community fixed effect, and ϵ_{ic} is an error term clustered at the community level.

Results are shown in Table 3. Columns 1 and 2 show estimates aggregated at the session level, where i indexes the experimental session and X_{ic} is a vector of average session-level characteristics. Having an authority lead by example increases the total contributions in an experimental session by 9.36 Bs. (s.e. 5.21), or approximately one environmental education book (column 2). Columns 4 and 6 show that this translates to an average individual-level increase of 1.03 Bs. (s.e. 0.71) or a 0.48 (s.e. 0.28) increase in the log odds of contributing an additional 5 Bs. coin. The median regressions in Columns 7 and 8 indicate that the likelihood that a contribution exceeds the median increases by 15% (s.e. 0.07) when a local authority gives first (column 8). We can thus comfortably reject the null hypothesis that leadership by example weakly decreases total contributions when an authority gives first.

What happens when a random community member leads by example? As Table 3 shows community members randomly selected to lead by example do not affect total or average contributions (columns 1 through 6). We are therefore unable to reject the null hypothesis that leadership by example weakly decreases public good provision when a random individual gives first. Random leaders do, however, increase the probability that contributions exceed the median by approximately 10% (s.e. 0.05, column 8). As shown by the p-value for a test of equal coefficients, the RL and AL treatments are significantly different from each other in most specifications. Having established that authorities increase total contributions when they lead by example, we turn next to the analysis of why authorities increase public good provision — whether the effect is driven by their own contributions, the contributions of followers or both.

4.1.2. Leader contributions

We next examine how authority and randomly selected contribution leaders adjust their behavior when leading by example. To test the null hypothesis that contribution leaders contribute less when they lead by example than when they give in a private simultaneous setting and use their first mover advantage to free-ride (Hypothesis 2), we restrict our analysis to contribution leaders in the RL and AL treatments and to individuals in the NL treatment, which includes elected authorities. The NL participants thus form a counterfactual for leader behavior, when contribution decisions are private. We regress contribution decisions on leadership treatment and authority status as follows:

$$y_{ic} = \alpha + \delta Authority_{ic} + \beta_1 AL_{ic} + \beta_2 RL_{ic} + \Gamma X_{ic} + \eta_c + \epsilon_{ic} \quad (2)$$

where $Authority_{ic}$ indicates that individual i in community c is an elected authority in the NL treatment, and all other variables use the same notation described in Eq. (1). The coefficient on $Authority_{ic}$ therefore captures any difference in the private contribution behavior of authorities and other community members. The coefficients β_1 and β_2 reflect the change in contribution behavior displayed by random individuals and

Table 3
Total contributions (all participants).

	Total				Individual			
	Continuous		Continuous		Ordered logit		≥10 Bs.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
RL	−0.941	−0.643	−0.183	−0.268	0.136	0.122	0.100*	0.096*
	(3.865)	(4.600)	(0.652)	(0.663)	(0.235)	(0.252)	(0.055)	(0.054)
AL	6.952	9.361*	1.257*	1.029	0.607**	0.481*	0.174**	0.149**
	(4.280)	(5.215)	(0.739)	(0.710)	(0.270)	(0.278)	(0.066)	(0.066)
N	104	104	580	580	580	580	580	580
Test: RL = AL (p-value)	0.102	0.040	0.072	0.071	0.045	0.103	0.273	0.402
Controls		Yes		Yes		Yes		Yes
Dep. variable mean, NL	42.947		7.767		7.649		0.378	

Notes: N = 104 in columns 1 and 2, N = 580 in columns 3–6. Columns 1 and 2 show OLS estimates of treatment effects on total session contributions. Columns 3–4 and 7–8 show OLS estimates of treatment effects on individual contributions. Columns 5–6 show log odds ratios from an ordered logit regression (see text for details). All specifications include community fixed effects and standard errors clustered at the community level. Controls refer to the full set of individual- and session-level controls shown in the balance table.
* $p < 0.10$.
** $p < 0.05$.

Table 4
Leader contributions.

	Continuous		Ordered logit		≥10 Bs.	
	(1)	(2)	(3)	(4)	(5)	(6)
Private giving by an authority	−0.921	−0.290	−0.005	0.416	0.006	0.014
	(2.257)	(1.987)	(0.815)	(0.776)	(0.171)	(0.155)
Public giving by a random leader (RL)	2.400*	1.690	1.076***	1.219***	0.347***	0.327***
	(1.266)	(1.379)	(0.391)	(0.460)	(0.080)	(0.095)
Public giving by an authority leader (AL)	7.192***	5.851**	2.144**	1.481*	0.424**	0.322*
	(2.637)	(2.482)	(0.917)	(0.890)	(0.196)	(0.187)
Controls		Yes		Yes		Yes
Dep. variable mean, NL	7.751		7.701		0.372	
Tests (*p*-values)						
Public giving RL = AL	0.098	0.141	0.244	0.774	0.706	0.978
Authority private = public	0.084	0.149	0.190	0.484	0.242	0.343

Notes: N = 258. The sample consists of individuals who led by example in the RL and AL treatments and all subject in the NL treatment. Authority refers OTB presidents in NL. Columns 1–2 and 5–6 present OLS estimates of treatment effects on leader contributions. Columns 3 and 4 show log odds ratios from an ordered logit regression (see text for details). All regressions include community fixed effects and standard errors clustered at the OTB level. Controls refer to the full set of individual- and session-level controls shown in the balance table.
* $p < 0.10$.
** $p < 0.05$.
*** $p < 0.01$.

authorities when they give publicly before others relative to non-authorities who give in private in the NL treatment. The regression is analogous to a difference in difference set up that includes leadership position and authority status, where β_1 represents the total effect for an authority in a contribution leadership role. The regression does not describe differences in the contribution behavior of authorities within community because community fixed effects are used and authorities can only be present in one of the two simultaneous sessions conducted in each community (AL or NL, specifically). With this caveat in mind, the test for Authority public versus private reported in the last row of Table 4 tests whether average authority contributions are significantly different when they are made in public in the AL treatment from when they are made in private in the NL treatment. In our design, like in most field settings, the contribution decisions that leaders make may be shaped both by their anticipated effect on follower decisions and by reputational considerations.

Table 4 shows that we can reject the null hypothesis of contribution leader free-riding. Both authorities and non-authorities increase their contributions in response to leadership by example. Authorities give 5.85 Bs. (s.e. 2.48) more when they lead by example than the average individual in the NL treatment (Columns 1 and 2). Randomly selected contribution leaders, on the other hand, give an additional 1.69 Bs. (s.e. 1.38). In all specifications, the coefficient on AL (β_1) is equal to or larger in magnitude than the coefficient on RL (β_2). Though standard errors are large, the difference between authority and non-authority leader contributions is marginally significant in some specifications. There is thus only weak evidence supporting the notion that authorities who lead by example give more than non-authorities who do so. Nevertheless, in contrast with the null hypothesis of first mover free-riding, these results suggest that part of the total increase in public good provision in the AL treatment is explained by the direct effect of the contribution of authorities who give publicly before others.

Interestingly, the differences in authority and non-authority contributions arise solely in response to leadership. The coefficient on the authority status indicator (δ) is small and imprecisely estimated, indicating that authority contributions are not different from the contributions of other community members when they contribute privately in the NL treatment. The final row of Table 4 shows that the difference in authority contributions in public and in private is marginally significant in only some specifications. It is important to note that we only observe 8 authorities giving in the NL treatment, so the authority status indicator variable is identified off a very small number of observations.[25]

[25] The authority does not always make the first (private) decision in the NL treatment. Controlling for order effects or the presence of an authority does not impact the outcome of either authority or non-authority contributions in the NL treatment.

Table 5
Follower contributions.

	Continuous				Ordered Logit		≥ 10 Bs. (OLS)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Random leader (RL)	-0.388	-0.367	-0.896	-0.802	-0.441	-0.381	-0.086	-0.093
	(2.411)	(2.437)	(0.834)	(0.940)	(0.470)	(0.476)	(0.100)	(0.106)
Authority leader (AL)	-0.010	0.315	-1.512***	-1.098	-1.113***	-1.064***	-0.208***	-0.181**
	(1.023)	(1.185)	(0.361)	(0.895)	(0.201)	(0.393)	(0.049)	(0.083)
RL x leader contribution	-0.027	-0.033						
	(0.238)	(0.238)						
AL x leader contribution	0.005	-0.020						
	(0.078)	(0.084)						
RL x leader contribution ≥ 10 Bs.			0.448	0.265	0.492	0.420	0.212*	0.218*
			(1.117)	(1.164)	(0.507)	(0.506)	(0.112)	(0.115)
AL x leader contribution ≥ 10 Bs.			1.856*	1.334	1.420***	1.312**	0.372***	0.329***
			(0.968)	(1.334)	(0.374)	(0.530)	(0.085)	(0.109)
Controls		Yes		Yes		Yes		Yes
Dep. Variable mean, NL treatment	7.767		7.767		7.728		0.378	
Total effects								
RL + RL X leader contribution ≥ 10 Bs.			-0.448	-0.537	0.051	0.039	0.126**	0.124**
			(0.791)	(0.773)	(0.269)	(0.271)	(0.062)	(0.060)
AL + AL X leader contribution ≥ 10 Bs.			0.344	0.236	0.307	0.248	0.164**	0.148**
			(0.898)	(0.906)	(0.315)	(0.332)	(0.070)	(0.070)
Tests (p-values)								
RL=AL	0.891	0.806	0.501	0.827	0.188	0.277	0.279	0.539
RL x leader contrib. = AL x leader contrib.	0.912	0.964	0.420	0.606	0.177	0.263	0.284	0.525
RL total effect = AL total effect			0.489	0.475	0.450	0.531	0.619	0.753

Notes: N = 510. Columns 1-4 and 7-8 show OLS estimates of treatment effects on follower contributions. Columns 5 and 6 show log odds ratios from an ordered logit regression with fixed effects (see text for a details). All specifications include community fixed effects and standard errors clustered at the OTB level. Controls refer to the full set of individual- and session-level controls shown in the balance table.
* $p < 0.10$.
** $p < 0.05$.
*** $p < 0.01$.

4.1.3. Follower contributions

Now we turn to the behavior of followers to test whether they respond negatively to the contribution decisions of leaders (Hypothesis 3), and whether the response differs by leader type. We exclude contribution leaders from the analysis and compare the behavior of followers in the RL and AL treatments to individuals who contribute privately in NL. The estimating equation is:

$$y_{ic} = \alpha + \beta_1 AL_{ic} + \beta_2 RL_{ic} + \beta_3 AL_{ic} \times y_c^{AL} + \beta_4 RL_{ic} \times y_c^{RL} + \Gamma X_{ic} + \eta_c + \epsilon_{ic} \quad (3)$$

where y_{ic} represents the contribution made by follower i in community c, and y_c^T for $T \in \{AL, RL\}$ represents the contribution made by the contribution leader in treatment T. The effect of the different leader types cannot be completely separated from the fact that leader type (RL or AL) is correlated with leader contribution decisions and leader characteristics (see Appendix Table A.4). Thus, treatment effects on followers should be interpreted as the combined effect of leader type, leader characteristics, and endogenous leader contributions.[26]

Columns 1 and 2 of Table 5 show OLS estimates of Eq. (3) with a continuous measure of y_c^T on the right hand side. The linear effect of continuous leader contributions on follower giving is statistically insignificant for both authority leaders (AL × leader contribution) and randomly selected leaders (RL × leader contribution), as are the level effects of the leadership treatments (AL and RL). The specification is restrictive in that it estimates the average effect of an increase in leader contributions on follower responses. Potential theories underlying a leadership by example effect require neither monotonicity nor linearity in the response function of followers. Before describing the rest of Table 5, we turn to less parametric analyses presented in Fig. 2, which shows the follower best response function (marginal effects) to leader contributions, including a quadratic term. The relationship is imprecisely estimated in Panel (a), where follower response is continuous, but is similar and more precise when follower response is binary in Panel (b). Fig. 2 suggests that authority leader contributions have a positive and approximately linear effect on the probability that followers give above the median (≥10 Bs.), while random leader contributions have a concave effect that is increasing up to 10 Bs. To accommodate these nonlinearities, the remainder of Table 5 tests for asymmetries in the response to high and low leader contribution, with the split at 10 Bs.

We examine the effect of high and low leader contributions on continuous follower contributions in columns 3 and 4 and categorical contributions in columns 5 through 8 of Table 5. Authority leaders who give less than 10 Bs. (coefficient on AL) insignificantly decrease follower giving by approximately 1.10 Bs. (s.e. 0.90, column 4) relative to the NL treatment. They significantly decrease the log odds that followers give an additional coin by 1.06 (s.e. 0.39, column 6), and significantly reduce the probability that follower contributions exceed the median by 18% (s.e. 0.08, column 8). Relative to a low authority leader contribution, a high AL contribution has a positive effect on follower contributions. Specifically, a follower of an authority who gives above the median gives 1.33 more Bs. (insignificant, column 4), has a significant 1.31 greater log odds of giving an additional 5 Bs. coin (column 6) and is 33% more likely to give above the median than is a follower who observes an authority leader contribute below the median (column 8).

In general, the coefficients on the RL treatment variables in Table 5 are of the same sign, smaller magnitude and less precisely estimated than the corresponding AL treatment effects. Column 8 shows that a random leader who gives at least 10 Bs. increases the probability that followers give at least 10 Bs. by 12.4% (s.e. 0.06) relative to the No Leader treatment and by 22% (s.e. 0.12) relative to random leaders who give

[26] We did not vary the amount authorities and non-authorities give when they lead by example because doing so would require letting subjects know that the leader is not freely choosing the amount they wish to contribute (in order to avoid using deception). This may generate a different response to leadership by example and would not be able to capture the leader response to leadership opportunities that we analyze in the previous subsection.

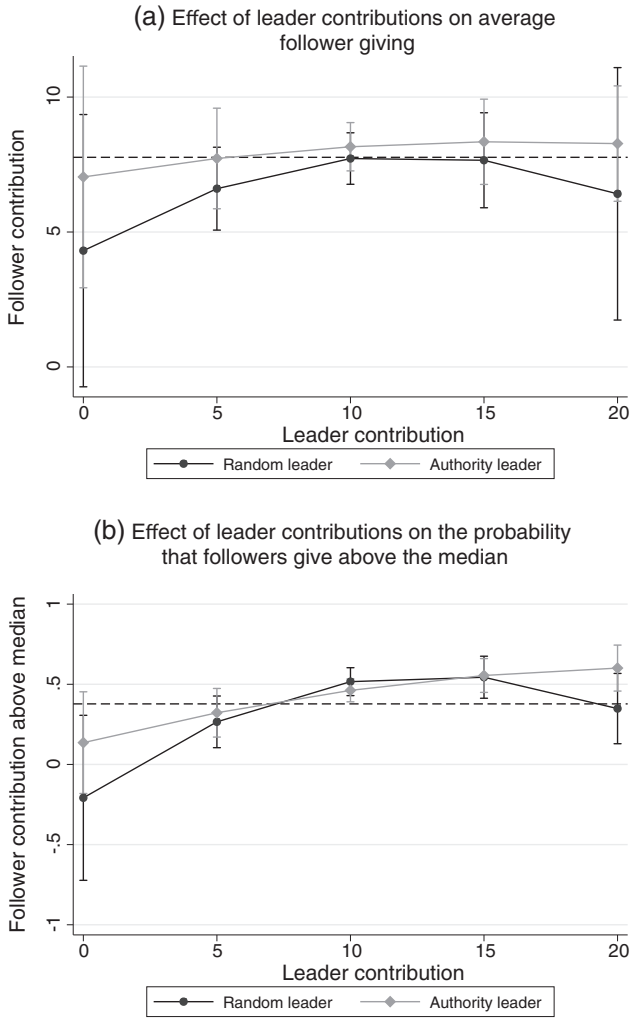

Fig. 2. Leader influence over followers. Notes: Figures represent the marginal effects of regression coefficients for random and authority leader contributions. See text for a description of the regressions. We present estimates for leader contributions up to 20 Bs. because this is the region of common support (see Fig. 1, panel b).

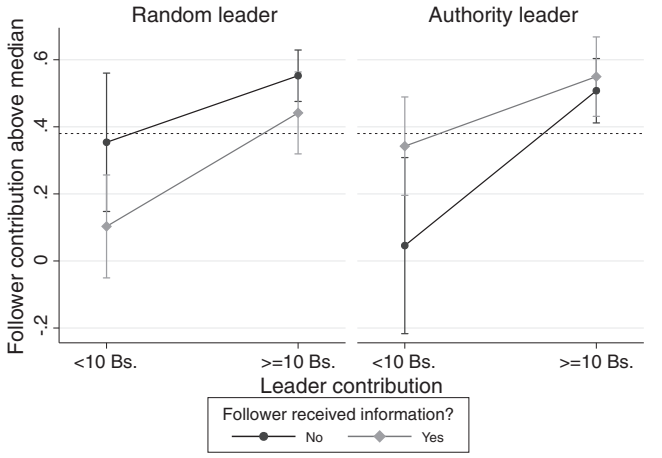

Fig. 3. Information signaling and leader influence. Notes: Figures represent the marginal effects of regression coefficients for random and authority leader contributions by information condition. See text for a description of the regression. The horizontal dashed line provides a reference for average contributions above the median in the NL treatment pooled across information conditions.

less than 10 Bs. This is the only specification in which random leaders are seen to have a significant influence over follower contributions. The differences in follower behavior across leader types are statistically imprecise; p-values from the relevant t-tests are reported in the final rows of Table 5. In rejection of the null hypothesis of free-riding, we find some evidence that follower contributions positively respond to the example set by authority and random leaders, with relatively stronger evidence for authority leaders.

4.2. Heterogeneous treatment effects

We turn next to the question of why leadership affects voluntary contributions in our setting, by exploring treatment heterogeneities in both leader and follower contribution decisions. First, we test Hypothesis 4, which establishes that followers' response to leader contributions does not differ based on their exposure to information. Second, we examine heterogeneities in leader influence based on the observable characteristics of leaders.

4.2.1. Followers: information signaling

Recall that the cross cutting information manipulation generated exogenous variation in the information available to study participants about the quality of the public good. Null Hypothesis 4 establishes that there is no information signaling and thus no difference in the response of informed and uninformed followers to leader contributions. The alternative hypothesis is that leadership by example signals information, so the contribution decisions of leaders have greater influence on the decisions of uninformed followers, because uninformed followers are more dependent on the information conveyed by the leader's decision than are informed followers.

Fig. 3 shows the marginal effects from a regression of the probability that the follower contributes above the median on the three-way interaction of leader type, leader contribution and whether the follower received information about the quality of the public good. As in previous analyses, we include a full set of controls, OTB fixed effects, and cluster standard errors at the OTB level. The follower's response function to the leader's contribution is plotted, with 90% confidence intervals. The average contribution in the No Leader treatment, which does not significantly differ across information conditions, is shown for reference by the dashed horizontal line.[27] While we lack precision in many of the relevant statistical tests, we observe several suggestive patterns in follower response to leader contributions across information conditions which allow us to reject the null hypothesis that information does not affect contributions.

First, when authority leaders contribute below the median, uninformed followers are significantly less likely to make a contribution above the median than are informed followers and participants in the NL treatment. When an authority makes a contribution above the median, both informed and uninformed followers are more likely to make a contribution above the median than are individuals in the NL treatment. Consequently, the difference in the response to the authority leader's contribution is greater for uninformed followers, as seen by the slope of the best response function of uninformed followers of authority leaders in the right hand panel. Second, the differences across information conditions in the response to a random leader's contribution are considerably weaker and not consistent with the predictions of information signaling. Information has a negative and significant level effect on follower contributions. Uninformed followers of random leaders are more likely to make contributions above the median when the leader contributes above the median, both relative to the simultaneous setting

[27] Specifically, in the described regression, the coefficient on information in the NL treatment is 0.014 (s.e. 0.829).

and to informed random leader contributions below the median. Statistical tests included in Appendix Table A.5 show that, with the exception of information, the responses by leader type are statistically indistinguishable. We discuss reasons why our information manipulation might be seen as a lower bound on the information signal provided by the authority leader's contribution in the conclusion.

4.2.2. Leaders: individual characteristics

As discussed in Section 3, authorities differ from the average community member on a number of dimensions, including gender, education, assets, and community participation (see Appendix Table A.4). As a result, the influence of authority leaders may be driven not by the position that they hold in the community but by their observable characteristics. Some relevant traits, such as education and wealth, may allow leaders to generate better information signals, trigger more reciprocity, and have stronger social influence regardless of their status as elected authorities. Holding an authority position may, on the other hand, convey an additional influence that extends beyond observable characteristics.

Though our study is not designed to explicitly investigate how the observable characteristics of leaders explain leadership influence (i.e. none of our four hypotheses pertain to leader characteristics), we take advantage of the fact that randomly selected contribution leaders vary in the degree to which they resemble the average elected authority. We construct an "authority propensity score" using a probit regression of authority status on the six characteristics where authorities significantly differ from the rest of the community: gender, education, wealth, participation in community meetings, trust in NGOs and the likelihood that they are teachers.[28] Each contribution leader is assigned an authority propensity score between 0 and 1, which describes the resemblance of each contribution leader to the average authority in the study. The top panel of Fig. 4 shows the resulting distribution of contribution leader types across the propensity score.

To directly examine the relationship between leader characteristics and leader influence, we construct a new outcome variable: the absolute difference between leader and follower contributions. The bottom panel of Fig. 4 shows the marginal effects from a regression of this measure of leader influence on an interaction of authority propensity score and leadership treatment, controlling for the leader contribution amount, individual- and session-level controls, and community fixed effects, with standard errors clustered at the community level. The figure shows that random leaders are more influential the more they look like the typical authority in the study. At low authority propensity scores, authorities are significantly more influential than are random leaders. However, random leaders who resemble the typical authority in the study area exert an influence that is indistinguishable from that of an authority leader with similar characteristics. It is worth noting that leader contributions are also weakly correlated with the characteristics on which authorities differ from other individuals. Though the regression results shown in Fig. 4 control for leader contribution, some of the effect of contribution leader characteristics may be due to the size of their initial contribution.

The results in Fig. 4 indicate that random leaders with authority-like characteristics are just as influential as local authorities. This suggests that at least some of the influence that authorities have when they lead by example is driven by their observable characteristics, and may indicate that communities choose their leaders based, in part, on observable characteristics that are correlated with influential leadership.

[28] Each of the covariates used in the probit regression is balanced after imposing common supports. We implement the propensity score matching using the algorithm developed by Becker and Ichino (2002). Note that authorities differ from the rest of the population on several participation-related characteristics, including participation in OTB meetings and projects and agreement with OTB decisions. We focus on one of these variables: participation in OTB meetings.

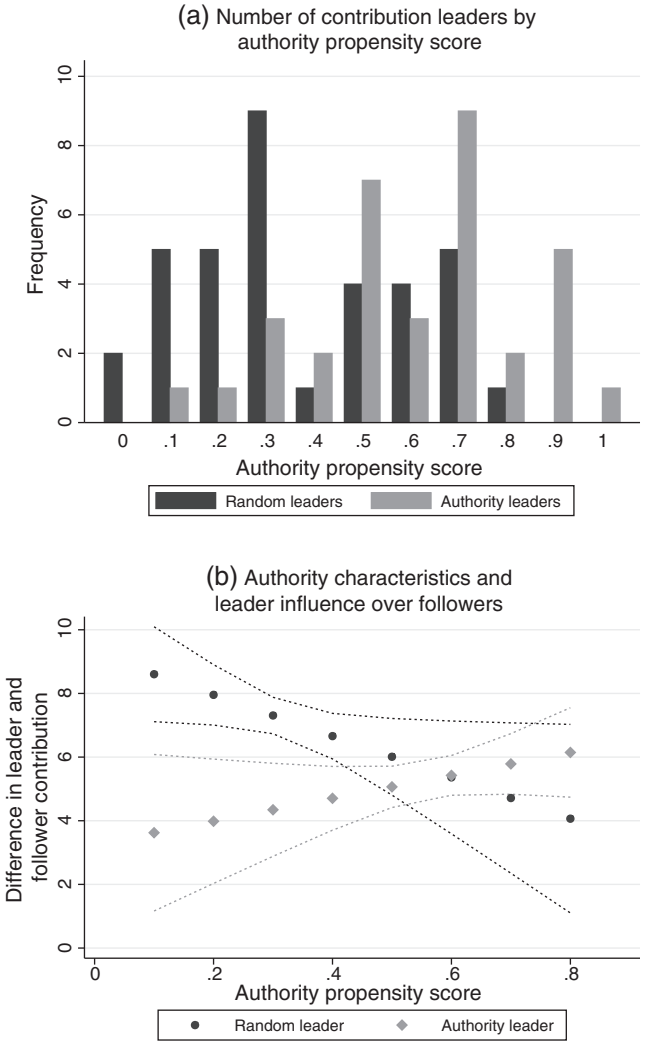

Fig. 4. Authority characteristics and leader influence. Notes: The top panel shows the distribution of contribution leader types across the constructed authority propensity score, which measures how closely the contribution leader resembles an elected authority. The bottom panel shows the marginal effects of the interaction of contribution leader type and authority propensity score on a measure of leader influence. See text for a description of the regression.

4.3. Robustness checks

We use two types of robustness checks to address possible selection bias resulting from non-compliance with the assigned treatment in some communities. First, we calculate two stage least square estimates of treatment effects that use treatment assignment to the AL treatment as an instrument for administered AL treatment, as follows:

$$y_{ic} = \delta_1 RL_i + \delta_2 \widehat{AL}_{ic} + \Gamma X_{ic} + \eta_c + \epsilon_{ic}. \quad (4)$$

The fitted values $\widehat{AL}_{ic}$ are obtained from the first stage regression

$$AL_{ic} = \theta D_{ic} + \Gamma X_{ic} + \varepsilon_{ic} \quad (5)$$

where D_{ic} is an indicator for assignment to the AL treatment for individual i in community c. Second, we restrict our analysis to the sample of communities that complied with treatment assignment and estimate treatment effects directly as in our main specifications.

Appendix Table A.6 presents revised estimates of treatment effects on total contributions. Overall, the results look similar to the main specifications, and are stronger under the instrumental variable specification

in most cases. This strengthening of the results under the IV specification is due to the relatively low contributions among the four replacement authority leaders. The limited sample analysis sometimes lacks statistical power because of the loss of sample size. The same robustness specifications are carried out for the leader and follower results. These are presented in Appendix Tables A.7 and A.8 respectively, and are also consistent with our main results.

5. Conclusion

Local authorities in developing countries often wield substantial power, and some evidence shows large authority fixed effects in community development outcomes, including the provision of public goods (Chattopadhyay and Duflo, 2004; Miguel and Gugerty, 2005). What role do local authorities play? Can they help communities overcome collective action problems and sustain higher levels of voluntary public good provision? If so, how? A number of channels present themselves: sanctioning or rule enforcement, moral suasion, liaison with outside resources, reciprocity, and leadership by example. Our study offers novel evidence on the latter mechanism.

We implement small group experiments in 52 communities in rural Bolivia to examine the role that locally elected authorities play in the voluntary provision of public goods when they lead by example. In our setting, authorities exert a significant influence over voluntary contributions even without the ability to monitor, sanction, or coerce. On average, total contributions increase by approximately 20% when the group is led by an elected local authority who makes an initial public contribution. Authorities significantly increase their contribution decisions when they lead by example relative to when they contribute in a private, simultaneous decision setting where their contributions do not differ from those of the average community member. Leader decisions may be driven both by their anticipated effect on follower contributions and by reputational considerations. Authorities also positively influence the contribution decisions of followers, to a marginally but insignificant greater extent than do random individuals who lead by example. The effect on followers is driven by a combination of authorities' leadership status, their contribution amounts, and their observable characteristics.

Our design explores one of the best-studied mechanisms underlying a positive effect of leadership by example: information signaling. We find that the predictions of information signaling are consistent with the influence that authorities have on their followers, but cannot explain all of their influence. Our information manipulation may be seen as generating a lower bound on the true relevance of information signaling for several reasons. First, we allow randomly selected followers to inspect the environmental education books rather than providing contribution leaders an additional informational advantage as is typically done in laboratory studies. However, if followers believe that the leader is better able to extract information from inspection of the books or has additional information about the books and their value, then the information asymmetry between the leader and informed followers may persist. Second, uninformed followers were not aware that contribution leaders had access to additional information, which left the existing information asymmetries between the contribution leader and followers intact. While telling the uninformed participants that contribution leaders were informed might have increased treatment effects, it might have also generated experimenter demand effects that otherwise biased our estimates of the natural importance of information signaling in our study setting.

Methodologically, our study offers an innovative approach to studying endogenously arising behavior within groups in field settings. The inclusion of community fixed effects allows us to address many of the concerns associated with unobservable factors driving leader and follower contributions within communities. While the use of community fixed effects eliminates endogeneity concerns at the community level, they may still exist at the session level. We test for session level correlates of leader contributions and find only one significant explanatory variable out of 13 tested.[29] We also borrow best practices from both the lab and the field in a number of other design features, including precise measurement of selection into the study, using earned money in the voluntary contribution decision, and studying contributions that acquire an actual good that benefits the community. Combining the rigor and insights of laboratory studies with the complexities of social interactions in the field offers a promising direction for future research. Particularly where leadership is concerned, stepping outside of the laboratory can generate insights about how actual leaders influence their followers and how the characteristics of individuals and groups interact.

Our findings are broadly consistent with existing studies of leadership in field-lab settings, which have shown that leaders can improve local cooperative outcomes. In rural Colombia, d'Adda (2012), for example, shows that leaders increase total contributions in a repeated simultaneous donation setting when contributions are public. Kosfeld and Rustagi (forthcoming) show that the sanctioning style of local leaders in the lab is correlated with the cooperative success of communities in actual conservation practices in Ethiopia. Beekman et al. (2014) show that communities with corrupt officials display lower levels of cooperation elicited through contributions in public good games in Liberia. These studies, like ours, raise the question of whether "good" leaders make communities more successful at resolving local level collective action problems or if cooperative communities elect "good" leaders to begin with.

Our results hold constant other means of influence that local authorities have at their disposal, such as sanctioning or coercion, which may be relatively more or less important than leadership by example in sustaining voluntary public good provision in other settings. In places with less decentralization, greater heterogeneity, or more corrupt leaders, other actors within the community may be relatively more influential than locally elected authorities. The comparative statics we observe in response to leadership by example may therefore generalize to other types of opinion leaders and different forms of leadership influence. Our methodology offers the ability to identify individual influence over voluntary contributions at the local level, and could be applied to settings where local leadership is less objectively defined. By exploring the causal mechanisms underlying leadership by example, we offer suggestive evidence that leadership by example may be particularly important where informational asymmetries are large or where opinion leaders stand out from followers on observable characteristics. Further research is needed to identify other mechanisms by which leadership by example affects voluntary contributions in field settings, and to understand its importance relative to other means of authority influence.

Acknowledgments

We thank Fundación Natura, the Sustainability Science Program at Harvard University, and the Agricultural Technology Adoption Initiative for support. Nigel Asquith, María Teresa Vargas and the rest of the staff at Natura provided invaluable support and guidance. Nava Ashraf, Lee Benham, Matthieu Chemin, Giovanna d'Adda, Sarah Jacobson, Margaret McConnell, Mushfiq Mobarak, Mary Shirley, Lise Vesterlund, Randy Walsh and audience members at the Inter-American Development Bank, IFPRI, Carnegie Mellon, Harvard, Pomona College, Trinity College, UC Santa Barbara, University of Wisconsin — Madison, University of Pittsburgh, Yale, the AERE summer meeting, the SITE experimental workshop, ESA, NTA, NEUDC, and the CAGE Conference on Generosity and Wellbeing provided useful comments. We thank Blaine Pellicore and Filipe Miranda for research assistance.

[29] Specifically, we see that out of the full set of individual- and session-level controls, only average session-level assets measured as the raw sum of durables owned by the household is associated with leader contributions, which confirms that leaders are not systematically adjusting their behavior based on the random group of followers to which they were assigned. The full table of results is available on request.

Appendix. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.jpubeco.2014.10.003.

References

Andreoni, J., 2006. Leadership giving in charitable fund-raising. J. Public Econ. Theory 8 (1), 1–22.

Andreoni, J., Brown, P.M., Vesterlund, L., 2002. What makes an allocation fair? Some experimental evidence. Game Econ. Behav. 40 (1), 1–24.

Ariely, D., Bracha, A., Meier, S., 2009. Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. Am. Econ. Rev. 99 (1), 544–555.

Baetschmann, G., Staub, K.E., Winkelmann, R., 2011. Consistent estimation of the fixed effects ordered logit model. IZA Discussion Paper, p. 5443.

Bag, P.K., Roy, S., 2011. On sequential and simultaneous contributions under incomplete information. Int. J. Game Theory 40 (1), 119–145.

Becker, S.O., Ichino, A., 2002. Estimation of average treatment effects based on propensity scores. Stata J. 2 (4), 358–377.

Beekman, G., Bulte, E., Nillesen, E., 2014. Corruption, investments and contributions to public goods: experimental evidence from rural Liberia. J. Public Econ. 115, 37–47.

Besley, T., Pande, R., Rao, V., 2012. Just rewards? Local politics and public resource allocation in south India. World Bank Econ. Rev. 26 (2), 191–216.

Bicchieri, C., 2005. The Grammar of Society: The Nature and Dynamics of Social Norms. Cambridge University Press.

Bracha, A., Heffetz, O., & Vesterlund, L. 2009. Charitable giving: the effect of exogenous and endogenous status. Unpublished manuscript.

Bruhn, M., McKenzie, D., 2009. In pursuit of balance: randomization in practice in development field experiments. Am. Econ. J. Appl. Econ. 200–232.

Chattopadhyay, R., Duflo, E., 2004. Women as policy makers: evidence from a randomized policy experiment in India. Econometrica 72 (5), 1409–1443.

d'Adda, G., 2012. Leadership and Influence: Evidence from an Artefactual Field Experiment on Local Public Good Provision. University of Zurich Department of Economics Working Paper, p. 59.

Eckel, C.C., Fatas, E., Wilson, R., 2010. Cooperation and status in organizations. J. Public Econ. Theory 12 (4), 737–762.

Gächter, S., Nosenzo, D., Renner, E., Sefton, M., 2010. Sequential vs. simultaneous contributions to public goods: experimental evidence. J. Public Econ. 94 (7), 515–522.

Gächter, S., Nosenzo, D., Renner, E., Sefton, M., 2012. Who makes a good leader? Cooperativeness, optimism and leading-by-example. Econ. Inq. 50, 953–967.

Grossman, G., Baldassarri, D., 2012. The impact of elections on cooperation: evidence from a lab-in-the-field experiment in Uganda. Am. J. Polit. Sci. 56 (4), 964–985.

Harbaugh, W.T., 1998. The prestige motive for making charitable transfers. Am. Econ. Rev. 88 (2), 277–282.

Hermalin, B.E., 1998. Toward an economic theory of leadership: leading by example. Am. Econ. Rev. 88 (5), 1188–1206.

Hermalin, B.E., 2007. Leading for the long term. J. Econ. Behav. Organ. 62, 1–19.

Huck, S., Rey-Biel, P., 2006. Endogenous leadership in teams. J. Inst. Theor. Econ. 253–261.

Humphreys, Macartan, Masters, William A., Sandbu, Martin E., 2006. The role of leaders in democratic deliberations: results from a field experiment in São Tomé and Príncipe. World Polit. 58 (04), 583–622.

Jones, B., Olken, B., 2005. Do leaders matter? National leadership and growth since World War II. Q. J. Econ. 120, 835–864.

Karlan, D., List, J., 2012. How can Bill and Melinda Gates increase other people's donations to fund public goods? NBER Working Paper, p. 17954.

Karlan, D., McConnell, M.A., 2014. Hey look at me: The effect of giving circles on giving. J. Econ. Behav. Organ. 106, 402–412.

Kosfeld, M., Rustagi, 2014. Leader punishment and cooperation in groups: experimental field evidence from commons management in Ethiopia. Am. Econ. Rev. (forthcoming).

Kumru, C.S., Vesterlund, L., 2010. The effect of status on charitable giving. J. Public Econ. Theory 12 (4), 709–735.

Meidinger, C., Villeval, M.C., 2003. Leadership in teams: signaling or reciprocating? MSE Working Paper.

Miguel, E., Gugerty, M.K., 2005. Ethnic diversity, social sanctions, and public goods in Kenya. J. Public Econ. 89 (11), 2325–2368.

Miller, G., Mobarak, A.M., 2014. Learning about new technologies through social networks: experimental evidence on non-traditional stoves in rural Bangladesh. Mark. Sci. http://dx.doi.org/10.1287/mksc.2014.0845 (forthcoming).

Olken, B.A., Singhal, M., 2011. Informal taxation. Am. Econ. J. Appl. Econ. 3, 1–28.

Potters, J., Sefton, M., Vesterlund, L., 2005. After you: endogenous sequencing in voluntary contribution games. J. Public Econ. 89 (8), 1399–1419.

Potters, J., Sefton, M., Vesterlund, L., 2007. Leading-by-example and signaling in voluntary contribution games: an experimental study. Economic Theory 33 (1), 169–182.

Reinikka, R., Svensson, J., 2004. Local capture: evidence from a central government transfer program in Uganda. Q. J. Econ. 119 (2), 679–705.

Smith, S., Windmeijer, F., Wright, E., 2014. Peer effects in charitable giving: evidence from the (running) field. Econ. J. http://dx.doi.org/10.1111/ecoj.12114 (forthcoming).

Varian, H.R., 1994. Sequential contributions to public goods. J. Public Econ. 53 (2), 165–186.

Vesterlund, L., 2003. The informational value of sequential fundraising. J. Public Econ. 87 (3–4), 627–657.