

Please do not cite or circulate without author permission

Deliberative Decision-Making and the Public Good: A Behavioral Lab Experiment in Kenya

Tara Grillos
Department of Political Science
Purdue University

Prepared for delivery at the Workshop on the Ostrom Workshop (WOW6) conference
Indiana University Bloomington, June 19-21, 2019
© copyright 2019 by author

Acknowledgements: I wish to thank the staff at Busara Center for Behavioral Economics for their support and feedback in preparing the experiment for implementation, in particular Jennifer Adhiambo, Mary Chebet, Joseph Kimani, Kelvin Kihindas and Channing Jang. For helpful input into the research design and analytical methods, I would like to acknowledge Jay McCann, Mike Touchton, Brian Wampler, and Richard Zeckhauser. For comments on previous drafts of this paper, I thank participants in the 2018 Association for Public Policy Analysis and Management conference, 2019 NYU-CES Conference on Experimental Political Science, and the Political Science Research Workshop at Purdue University. This experiment was implemented thanks to financial support from the Exploratory Research in the Social Sciences Grant from Purdue University's College of Liberal Arts, and early work on this project was supported by the Harvard Law School Project on Negotiations' Graduate Student Grants Program.

ABSTRACT

Over the past twenty years, participation by citizens in government decision making has been promoted all over the world. While there are normative reasons to encourage more inclusive decision-making processes, costly and time-intensive collective decision-making processes are often justified on the grounds that they may also improve collective outcomes. In political theory, deliberative discussion, in particular, is believed to be more transformative than a mere aggregation of individual preferences, potentially changing both the outcomes of the decision process and the value that participants place on those outcomes. If outcomes in question relate to a public good, these effects have the potential to encourage individuals to contribute to that public good despite individual incentives to deviate. One of Ostrom's original design principles for successful commons management was that those affected by rules participate in decision-making related to them. But these referred to endogenously created institutions, whereas the use of participatory processes in public policy more often involves efforts to engage citizens or stakeholders in externally organized decision-making processes. Despite widespread interest, empirical evidence in support of such co-produced decision-making remains weak.

This study uses a behavioral lab experiment in Kenya to examine whether engagement in group decision-making regarding creation of a public good leads to greater contributions to the public good. I find that participation that involves deliberation (but not a simple majority rule) results in better collective outcomes. This effect is achieved primarily through better strategic decision making, which minimizes the costs associated with socially desirable behavior. However, I also find an indirect effect of deliberation, mediated by preference change, on individual effort exerted on behalf of a collective good. This provides rigorous empirical evidence in support of claims by political theorists that deliberation may transform citizens, rather than merely aggregate preferences or coerce agreement.

INTRODUCTION

Participation in decision-making by citizens and stakeholders has been widely lauded as a method for improving outcomes in democratic governance (Fung & Wright, 2001), environmental management (Koontz & Thomas, 2006; Reed, 2008), and international development (Mansuri & Rao, 2004, 2012). The World Bank alone has invested billions of dollars in the implementation of community-driven development, which emphasizes the participation of beneficiaries in decision-making around development projects (Mansuri & Rao, 2004). Forms of direct democracy, such as participatory budgeting, have spread all over the world (Ganuza & Baiocchi 2012, Goldfrank 2012), from its origins in South America to various countries including South Korea (Hong & Cho 2018), Indonesia (Grillos 2017), and Kenya (Wampler & Touchton 2017). Public participation in decision-making is now formally encouraged in several national constitutions. Deliberative democracy has even been proposed as an ideal for science communication (Dietz 2013) and global environmental governance (Berg & Lidskog 2018).

While there are compelling, normative reasons to encourage more inclusive forms of decision-making independent of results, it is important to recognize that participation also imposes costs on participants. These costs may be particularly burdensome to the poor, who are already constrained in both time and material resources. Many scholars point to positive benefits of participation relative to none, particularly with respect to the resolution of collective action problems, such as environmental resource management (Ostrom 1990, Agrawal 2005, Brooks et al. 2012). However, forms of participation vary greatly in

practice, differing along several key dimensions (Fung 2006b). Some forms of participation are more costly and time-intensive than others, and so the particular design of participatory institutions should ideally be justified with demonstrated benefits of one form over another. A crucial open question in this line of literature is: which exact forms of participation improve which outcomes through what mechanisms?

Political theory gives us ample reason to believe that one particular form of participation, deliberative argumentation, should be expected to produce a host of positive outcomes. Yet studies of deliberation have largely shied away from an explicit examination of outcomes, instead focusing on procedural factors (Landemore 2017). Empirical work has established that deliberation results in shifts of opinion (Luskin, Fishkin, and Jowell 2002, Barabas 2004, Fishkin & Luskin 2005, Farrar et al. 2010), but has had very little to say about whether changes produced by deliberation are actually “for the better” (Neblo 2007). This is largely due to the fact that forms of deliberation commonly used in practice are often value-laden, making it impossible to objectively assess end results. Scholars have called for the elaboration and testing of specific, falsifiable hypotheses that follow from deliberative theory (Mutz 2008), and the literature on deliberative democracy has recently begun to embrace the experimental method, but has thus far provided limited exploration of effects on decision quality and behavior change.

There are two major impediments to establishing these causal relationships with observational data: First, forms of participation vary greatly in practice, and so it is very difficult to isolate particular dimensions of participation (and related causal mechanisms) that could be driving observed effects. Second, outcomes like the “collective good”, “decision quality,” and “socially desirable behavior” – all theoretically reasonable as hypothesized outcomes of participation and deliberation – are nearly impossible to define in real-world deliberative situations without imposing the values of the researcher.

Participation scholars rarely isolate a single aspect of participation as instrumental to observed positive effects, to help justify its inclusion in institutional design. Meanwhile, political theorists extol the virtues of deliberative discussion, but empirical work on deliberation has largely shied away from a direct causal examination of collective outcomes. Both bodies of research could benefit from greater dialogue with each other, and also from the use of controlled variation, which is uniquely provided by laboratory experiments (Falk & Heckman 2009). I use a laboratory experiment to test for theoretically plausible impacts of deliberative participation in a controlled context that allows me to avoid the aforementioned obstacles of unspecified treatment components and undefined quality of outcomes.

In this paper, I briefly discuss the prior literature on participation and collective action, on the one hand, and deliberative democracy, on the other, suggesting that each has a lot to offer the other. I then present a conceptual framework in which I posit that there are two core pathways through which design features of a decision process may influence outcomes. They may (i) improve the quality of resulting decisions or they may (ii) have effects on the participants themselves that could lead to socially desirable behavior change. I also discuss various underlying causal mechanisms through which decision quality and behavior change could be achieved. Next, I describe the study setting, and present an experimental design with controlled variation in the form of particular aspects of participatory decision-making to which participants are randomly assigned, and in which participants are presented with a decision that (i) has direct implications for the collective good, and (ii) has an objectively calculable, socially optimal solution. I use this experimental approach to isolate the causal effect of participation on both decision outcomes and human behavior.

PARTICIPATION & COLLECTIVE ACTION

Overcoming the social dilemma implicit in collective action problems is one of the key challenges facing political scientists (Ostrom 1998). One of Ostrom's original design principles for successful collective action was that those affected by rules participate in decision-making related to them (1990). But these referred to endogenously created institutions, whereas the use of participatory processes in public policy more often involves efforts to engage citizens or stakeholders in externally organized decision-making processes, which range from mere consultation with stakeholders to direct democracy through voting to extensive discussion resulting in consensus.

Despite long-standing claims about the beneficial effects of participation, many have observed that the empirical evidence in support of participatory decision-making is inconsistent (Duit & Hall, 2014, Koontz & Thomas, 2006; Birnbaum, 2016). In addition to positive cases (Agrawal 2005, Brooks et al. 2012), there are instances where participation has failed to yield anticipated results on cooperation (Lubell 2004). Many studies take the form of either individual case studies or meta-analyses and literature reviews attempting to synthesize those case studies (Beierle & Cayford, 2002; Irvin & Stansbury, 2004; Koontz & Thomas, 2006; Reed, 2008). But efforts to systematize this body of work are complicated by the myriad forms that participatory institutions take in practice. When studies demonstrate that some form of participation has led to superior outcomes, it is often difficult for the researcher to pinpoint which design features of the decision-making process are actually doing the work.

Effects of participation are likely to vary depending on the particulars of the methods used in practice. If the goal is simply to ensure that local preferences are taken into account, then mere consultation with beneficiaries or direct voting procedures should be sufficient to achieve results. If the motivation is to change perceptions alter behavior, on the other hand, more intensive (and therefore costly) forms of participation may be warranted.

Some prior experimental work has specifically examined the effects of group decision-making on collective action behavior. In a field experiment, voting on selection of a public good increased self-reported willingness-to-pay for maintenance of that good (Olken 2010), but actual investments were not observed. Similarly, in a laboratory setting, involvement in decision-making also increased self-reported intentions to promote sustainability (Aga, et al., 2017). In addition, election of a leader who can enforce sanctions increased cooperation in a public good game (Grossman & Baldassarri 2012). But none have looked specifically at the effect of deliberation on decision quality and behavior change.

Laboratory experiments have contributed greatly to the study of collective action (Ostrom 2005, Poteete, Janssen & Ostrom 2010). One of the most consistent findings in this body of literature is that face-to-face communication increases the likelihood of groups to achieve collective action (Bornstein 1992, Dawes et al. 1990, Ostrom et al. 1994, Sally 1995, Ledyard 1995). However, the effect of communication varies with contextual factors (Cardenas 2004) and with the content of the communication (Lopez and Villamayor-Tomas 2016), and researchers still lack a complete understanding of how exactly communication improves cooperation. Deliberative theory promotes a particular kind of discursive communication that involves rational argumentation (Bächtiger et al. 2010) and has the potential to lead to more socially optimal decision-making and to more socially optimal behaviors. But the literatures on participatory institutions and local collective action have rarely engaged directly with political theory on deliberative democracy.

DELIBERATIVE DISCUSSION

Deliberative discussion is believed to be more transformative than a mere aggregation of individual preferences (Elster 1986, Chambers 2003). One of the central elements of deliberation is that it involves “reason-based decision-making,” in which participants try to persuade each other of a course of action using reasons that appeal to others, such as fairness, group-mindedness or logic regarding effectiveness (Fung & Wright 2003, Gutman & Thompson 2004, Neblo 2005, Thompson 2008). While the foundational work in deliberative democracy is largely normative in nature (Manin 1987, Habermas 1989, 1996, Cohen 1989, Dryzek 1994) and thus difficult to test empirically, scholars have called for middle-range theories that try to connect concrete aspects of deliberative communication to desirable outcomes (Mutz 2008, Bächtiger et al. 2010).

There has been an ‘empirical turn’ in deliberative democracy, resulting in a growing body of literature examining its effects (Carpini et al. 2004, Ryfe 2005, Bächtiger et al. 2010). The “science of deliberation” (Dryzek et al. 2019) has now provided convincing evidence that people are willing (Esterling, Neblo & Lazer 2011, Neblo, Esterling & Lazer 2018) and able (Gerber et al. 2018) to engage in quality deliberations and that deliberative encounters increase political discussion beyond the formal event (Lazer et al. 2015). The use of “deliberative contestation” in participatory processes may help marginalized groups to challenge capture by narrow interests (Gibson & Woolcock 2008) and help to overcome clientelism (Fujiwara & Wantchekon 2013). Citizens engaging in deliberation are able to recognize and subvert elite manipulation of discourse (Niemeyer 2011, Druckman & Nelson 2003), and are less prone to populist mobilization (Fishkin 2018). Deliberation may also decrease political polarization and ideological division (Grönland et al. 2015, Ugarizza & Caluwaerts 2014). Thus far, however, no experiments have directly assessed whether deliberation actually improves outcomes (whether through decision quality or behavior).

Several scholars have demonstrated that deliberation can result in shifts of opinion (Luskin, Fishkin, and Jowell 2002, Barabas 2004, Fishkin & Luskin 2005, Farrar et al. 2010). However, the empirical literature has had much less to say about the quality of the resulting opinions (Fung 2006a). This is in part due to a reluctance on the part of some deliberative democrats to embrace the existence of an objective ‘truth’ in political disagreements (Rawls 1993), while others have argued that work on deliberation must acknowledge and test its epistemic benefits for truth-seeking (Estlund 1998, Cohen 2009, Landemore 2017). Observational data suggests that deliberation at least leads to decisions that are more consistent with an individual’s own underlying values (Niemeyer 2011) and more rational at the collective level, in the sense of increased single-peakedness (List et al. 2013). Theorists suggest that deliberation can result in more socially oriented decision-making (and presumably behavior) by constraining self-interest (Ackerman & Fishkin 2002, Mansbridge et al. 2010), but empirical work demonstrates that consultation through open discussion may be more prone to elite capture than simple voting procedures (Beath et al 2012). Some scholars argue that group discussion can lead to further polarization and worse decisions than individuals deciding on their own (Sunstein 2000, 2002, Sunstein & Hastie 2015), though this work does not impose the requirement of deliberative argumentation. We still have very little causal evidence about whether or not deliberative argumentation, in particular, improves decision quality (Neblo 2007), as this is of course rather difficult to assess in a field setting without imposing the values of the researcher.

The empirical work on deliberation also lacks evidence about whether shifts in opinion result in behavioral change. The idea that participation in decision-making processes may change people’s beliefs, preferences and behaviors in ways that benefit society has a very long history in political theory (Mansbridge 1999). If people’s opinions are truly changed through deliberation, are they also more likely to value outcomes of the decision-making, and thereby more likely to invest in, maintain or comply with those outcomes?

Does deliberation make them more likely to act in favor of the collective good? As Mansbridge (1999) has posed, does participation in fact “make better citizens?”

CONCEPTUAL FRAMEWORK

There are various reasons to expect that shifting from centralized decision-making to a particular form of collective decision-making (for example, deliberative discussion) may influence outcomes related to the collective good. It may do so through two main pathways: it may improve the quality of the decisions that are made, or it may influence the behavior of the individuals who participated in that decision (Grillos 2015). Either of these pathways may be achieved through various categories of causal mediators. Effects may be merely aggregative (in line with the expectations of rational choice theory) or more transformative (in line with sociological or discursive institutionalism). I describe several categories of mechanisms below.

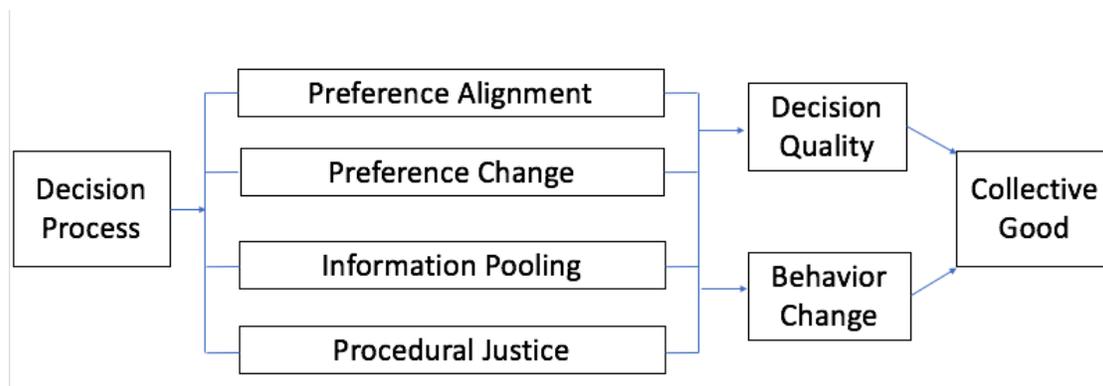


Figure 1: Conceptual Framework

Preference Alignment

The simplest mechanism involves preference alignment. Because the decision procedure now takes into account the preferences of local participants, the decision outcome is more likely to be something they preferred and placed value on in the first place. By considering the aggregated preferences of the individuals involved, collective decision-making allows for the selection of the outcome that most people valued the most to begin with. Regardless of whether the decision outcome is the best from the perspective of the collective good, it will be the best from the perspective of those individuals who selected it. If a majority of individuals value that decision outcome enough, this may be sufficient to induce increased collective action, such as investments in a resulting public good. Notably, this mechanism requires no change to individuals’ knowledge or preferences – it merely requires that pre-existing individual preferences be taken into account in the decision-making process. Simple majority rule mechanisms should be sufficient to activate this mechanism.

Preference Change / Procedural Utility

It is also possible that individual preferences are actually transformed through the process of decision-making, such that participants place greater value on the same exact decision outcome as a result of the participatory process. Work on procedural utility has confirmed that individuals value not only outcomes, but also the processes that lead to them (Frey and Stutzer 2004) and may value the same outcome more if they themselves participate in creating it (Norton et al 2011). They may therefore also be more likely to invest in, maintain or comply with those outcomes over the long-run. Self-determined choice can, under the right circumstances, improve motivation, effort, and task performance, even when choice is irrelevant to task difficulty (Patall et al. 2008; Leotti et al 2010). In addition to enhancing performance, decision-

making is associated with increased activity in a brain region (Murayama et al 2013) that is also linked to the valuation of goods and willingness to pay for them (Plassman et al 2007; Chib et al 2009).

However, the aforementioned studies involve *individual* decision-making in which the participant has full authority over the outcome of the decision. Self-determination theory (Deci, 1980; Ryan & Deci, 2000) may explain such findings, as it suggests that an individual's motivation to act is strongest when their perceived autonomy remains intact. In group decision-making, however, no individual has full control over the outcome, and no individual is guaranteed to receive their preferred outcome. It is thus unclear whether the psychological benefits of choice and autonomy extend to collective decisions.

Information Pooling

Another way that deliberation may improve collective outcomes is through information sharing. Where there are asymmetries of information, the aggregation of knowledge may accrue to each individual through a learning process, potentially changing their willingness to accept particular decision outcomes. Assuming the decision process allows for some discussion to occur, individuals may reveal information that was previously only available to them privately. Since decision-making may not immediately result in desired outcomes, but instead involves a path toward achieving those outcomes, individuals likely have imperfect information about how their values will be met by different decisions. Deliberative processes in particular may allow individuals to gain new information that causes them to update pre-existing beliefs or gain new perspectives (Martí 2006, Caluwaerts & Ugarizza 2012). This could lead them to value the decision outcomes differently, even if their underlying preferences have not changed.

The role of knowledge may not be restricted to the sharing of new information. Even if deliberation stops short of changing people's preferences, and even if no new information is introduced through the deliberative process, it may still succeed in changing the decision criteria through which people translate preferences and beliefs into a decision. For example, the deliberative process may help participants to overcome cognitive biases and acknowledge previously missed logical implications of existing knowledge (Hafer & Landa 2007, Landa 2015).

Procedural Justice / Constrained Self-Interest

Deliberation may also succeed in altering decision criteria by limiting people's ability to make decisions based purely on their own preferences. Individuals may be induced to accept an outcome that they do not prefer, whether through social pressure or self-imposed normative constraints. The procedural justice literature argues that people are more willing to accept a decision if they believe it was fairly determined (Tyler, 1990; Lind & Tyler, 1988; Tyler & Blader, 2000) – even if their personal preference over the outcome has not changed. Providing an individual with “voice” (understood as an opportunity to share their views on a decision) is associated with greater procedural justice and therefore acceptance of related outcomes (Thibaut & Walker, 1975; Price et al., 2001), even compared with direct “choice” - providing direct control over a decision (Earley & Lind, 1987). While there is an emphasis on preference change in deliberative theory, it is important to note that the deliberative ideal does not require the absence of self-interest, but rather suggests that self-interest can be constrained by deliberation (Mansbridge et al. 2010). Individuals may yield to the “forceless force” of the better argument (Habermas 1981) – leading to collective decisions that are more in line with the collective good even if competing individual preferences remain intact.

STUDY SETTING: KENYA

Kenya ratified a new constitution in 2010 by popular referendum. The new constitution sets up a federal system, granting greater authority to the county governments, and it includes requirements for citizen participation in government decision-making: “A county assembly shall... facilitate public participation and involvement in the legislative and other business of the assembly and its committees (Constitution of Kenya, p.119).” However, the specific method of public engagement was left largely to the discretion of the county governments, and many were still struggling to develop a public participation plan when I conducted key informant interviews with county officials in 2014 (Grillos 2018). The World Bank has provided support for the implementation of participatory budgeting, which typically takes place at the ward level and relies on consensus-based decision-making (Wampler & Touchton, 2017).

Kenya is also a setting in which local collective action is extremely salient. In Kenya, the provision and maintenance of public goods are largely dependent on local contributions, often through fundraisers known as *harambee* (Wilson 1992, Miguel & Gugerty 2005). This degree of community responsibility is fairly common throughout the developing world, where governments are often minimally responsive to marginalized communities. See, for example, Habyarimana et al. (2009)’s discussion of local public good provision in neighboring Uganda. They describe how decentralization of public good provision left more and more responsibility in the hands of local citizens. Even when public goods are directly provided by the government or international organizations, maintenance of those goods is often left up to local communities. This requires not only willingness to invest effort on behalf of the public good but also relevant knowledge and ability on the part of community members.

When top-down provision of public goods requires bottom-up investments to maintain those goods, participatory decision-making is often invoked as a method for increasing a sense of ownership among community members. Common approaches to development practice take this as a given. Then-president of the World Bank, James Wolfensohn, once wrote of in-country stakeholders: “I am convinced that their involvement and collaboration can not only make our development efforts more effective and sustainable, but can also foster ownership and a sense of belief in the relevance and value of our programs right down to the community level.”¹

There is still minimal experimental evidence on deliberative decision-making and its effects on collective outcomes through (i) the quality of decisions and (ii) participant behavior, especially related to collective action. When decision-making concerns choices related to some sort of collective or valence good (as opposed to more value-laden decisions such as choosing between different goods that individuals may value differentially), there is an outcome (the state of the good itself) that can be objectively assessed as being in the public interest. This study provides a rigorous experimental test of the hypothesis that participation in deliberative decision-making improves outcomes related to a collective good. I also provide some exploratory analyses aimed at distinguishing between the causal pathways underlying that hypothesized effect. The experimental design described below was registered as a Pre-Analysis Plan in the EGAP (Evidence in Governance and Politics) repository².

¹ Foreword to *The World Bank Participation Sourcebook*, February 1996. Available for download at: <http://documents.worldbank.org/curated/en/289471468741587739/pdf/multi-page.pdf>

² The pre-analysis plan for this project has EGAP ID number 20180720AA, available at: <https://egap.org/registration/4963>

METHODS

I take advantage of a controlled laboratory setting to assess differences between commonly used approaches to group decision-making. I differentiate groups that rely on a simple majority rule (expected to successfully aggregate preferences, but unlikely to achieve more transformative change) vs. a deliberative discussion resulting in a consensus-based decision (a more intensive and time-consuming form of decision-making, but one perhaps more likely to influence outcomes). I also differentiate between processes that make use of secret ballot voting vs. those that use a public vote procedure.

Laboratory experiments have been criticized for lacking generalizability across cultures and contexts (Levitt and List 2007, Henrich et al. 2010). Others argue that these concerns are overstated and that many common laboratory findings are indeed replicable across cultures (Klein et al. 2018), and that generalizability is a problem common to all research methods, not just lab experiments (Falk & Heckman 2009). Furthermore, the “realism” of an experimental context should not be judged by the context but rather by how well the experimental design approximates the real world experience it is meant to simulate (Falk & Heckman 2009). Deliberative scholars have argued that experiments on deliberation can be considered valid to the extent that they involve representation of varied viewpoints on a perspective of public concern (Caluwaerts & Ugarizza 2012). My experiment meets these conditions by simulating decisions related to a public good, on which there is no pre-treatment consensus.

I further mitigate concerns by recruiting subjects from a relevant, developing country context. In addition, a review article concluded that the psychological desirability of decision-making has a biological basis, rather than being learned through cultural influence (Leotti et al 2010). This implies that findings may be generalizable to different settings, making their extension to group decision-making processes an ideal topic for experimental research. An important advantage of laboratory research is that I can move beyond self-reported intentions to measure actual behavior within an incentivized behavioral game. In addition, the quality of decisions, which are difficult to specify in real-world situations without imposing the values of the researcher, can be more easily assessed in the controlled laboratory setting where the socially optimal decision is easily calculated.

The experiment was conducted at Busara Behavioral Lab in Nairobi, Kenya. Busara staff (Kenyan citizens) implemented the experimental protocol in Swahili, the national language. We conducted all aspects of the experimental protocol using z-Tree software (Fischbacher, 2007), except for the decision-making treatments. For the decision process itself, each team met separately in a smaller room outside of the computer lab, with facilitators following a protocol pre-programmed into Qualtrics. All sessions occurred between July 27th and August 23rd of 2018.

Participants

Busara draws its research subjects primarily from the Kibera slum, a low-income population for whom local collective action may be a prominent source of public good provision. This experiment included 570 participants, spread across four treatments and one control group.³ Sixty-three percent of the participants

³ A total of 587 participants originally participated in the study, with two repeat attendees (yielding a total of 589 observations). When the repeat participants were identified, all 15 observations from the session that included the repeat participants were dropped, over concerns that the repeat attendance could bias responses for their respective groups. Another team of 4 participants was dropped because it had fewer than 5 participants (which changes the earning incentives in a way that makes their behavior not comparable to others). Replacement sessions were then run (using the same treatment that had been randomly assigned to the dropped session). This left a final sample of 570 observations from 570 unique individuals

were female, 35% had never been married, and 74% report having engaged in some sort of real-world collective action in their communities within the past month. The participants were, on average, 34 years old with 2 children and with 10 years of education (the equivalent of having completed some secondary school). (See Appendix A for a full table of descriptive statistics.) A larger percentage of participants were assigned to the control group in order to increase statistical power. Of the 570 total participants, 210 (~37%) were assigned to the control group and the remainder were split evenly across the four treatment groups, with 90 participants in each of the four treatments. (See Appendix B for a discussion of statistical power calculations.)

Experimental design

This experiment involved a modified public good game in which contributions to the public good were determined by performance of individual team members in an effort task. In contrast to the commonly used voluntary contribution mechanism (VCM) game, this modified version acknowledges the fact that public good maintenance often faces additional barriers beyond willingness to cooperate (such as strategic decision-making regarding inputs, and willingness and ability to invest time and effort). Individual performance on the effort task earned money for a team fund which was then divided evenly across all team members, irrespective of individual performance.

Participants were first randomly assigned to a session which was randomly assigned to a treatment group. On the day of the experiment, participants engaged in a collective decision-making process with their team that determined which of several real-effort tasks they would participate in to earn money toward the public good (the team fund). The form that this decision-making process took depended on the treatment group assignment. There were four overlapping treatment groups and a control group. Depending on the randomly assigned treatment group, the decision was made through either a private or public vote and using either a majority rules voting procedure or through deliberative discussion requiring full consensus.

Table 1: Overlapping Treatment Groups

	Majority Rule	Deliberation & Consensus
Private Vote (Secret Ballot)	Treatment Group 1	Treatment Group 3
Public Vote (Show of Hands)	Treatment Group 2	Treatment Group 4

The experiment took place in several stages. These are depicted visually in Figure 2 and also summarized in the text below. The full experimental protocol documents are attached as Appendix K & L.

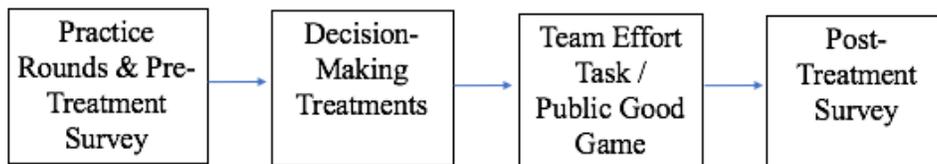


Figure 2: Stages of Experimental Design

(the same as was set forth in the pre-analysis plan). However, the main results presented in this paper are robust to the inclusion of the dropped observations.

1. Practice Rounds & Pre-Treatment Survey

Prior to treatment, I captured some information about the real-world activities of participants. Next, all participants were given an opportunity to briefly test out each of three effort tasks: the Letter Counting Task, the Sliders Task, and the Stroop Task. All are adaptations of previously vetted real effort tasks. In the Letter Counting task, participants are shown a string of letters and numbers and asked to count the number of times a particular letter appears in the sequence (adapted from Rey-Biel, Sheremeta & Uler 2011). In the Sliders task, participants are given a target number between 0 and 100 and asked to move an on-screen slider to that number (adapted from Gill & Prowse 2012). In the Stroop task, an arrow appears on screen and participants must tap the side of the screen that the arrow points to or the side of the screen that the arrow points from, depending on the color of the arrow (adapted from Stroop 1935).

The practice rounds were incentivized based on individual performance, so that they could be used as a measure of ability. Participants were told that they would later be given an opportunity to participate in one of these activities in order to earn money as a team. They were then asked to fill out a survey, which asked them to rate the three effort tasks according to enjoyment, difficulty, and overall preference. This survey allows me to compare results across individuals who had the same initial preferences.

2. Decision-Making Treatments

Individuals were randomly assigned to teams of five individuals each. Each team then engaged in one of five decision processes (four treatments and a control) which would determine which of several effort tasks their team would work on. During the decision-making stage of the experiment, the facilitator took notes about the results of the decision-making process and noted whether any team members knew each other before the experiment. Table 2 below depicts the four treatment groups, determined by two overlapping treatment variations.

In all treatment groups, including the control group, teams met face-to-face in a separate room and had an opportunity to introduce themselves. Facilitators then explained how the public good effort game would work (participants would earn points for a shared, team fund through performance on one of three effort tasks). In the control group, the task was then selected using a random number generator in Qualtrics, the team was informed of the selection, and they returned to the computer lab to play. In the decision-making treatments, in contrast, the team was told that they would select the task and the decision rule was explained.

In the deliberative consensus groups, each participant was asked to state their preference and give the reason why they preferred that task (to induce some form of reason-based argumentation). Then they were given time to discuss the options in order to arrive at a consensus decision. Once the team believed they had arrived at a decision, the consensus was then confirmed via a vote (which could be either a public show of hands or a secret ballot, depending on the overlapping treatment). If the vote revealed that there was not yet a full consensus, the team was asked to repeat the deliberation until they felt they could all agree on a single task. If a team could still not arrive at a consensus after a second round of deliberation, then the facilitators were instructed to default to the majority rule, but this situation never occurred.

In the majority rule groups, in contrast, individuals were not required to discuss anything in particular and were simply asked to vote for their preferred outcome. In the case of a tie, the task with the fewest votes was removed as an option, and the team was asked to engage in a second round of voting to break the tie. With five team members and three options, it was impossible for the second round of voting to result in a

tie. The treatment groups also varied according to whether vote tallying took place through secret ballot or a public show-of-hands.

3. Effort-based Public Good Game

Teams then had an opportunity to earn more compensation by performing the task selected in Step 2. Participants were invited to complete as many iterations of the activity as they could within 10 minutes, and they would earn 5 shillings⁴ per point earned (which would go to a team pot). Before the task began, we also asked participants to guess how well their team would perform on the task – a measure of expectations regarding the contributions of others. At the end of the activity, the money earned by the team would be evenly split across the team members. This portion of the experiment resembled a public good game, in which individual task performance earned money for an aggregate team pot, such that there was an opportunity for free-riding off the effort of others.

To make the opportunity for free-riding more realistic, we reminded participants that they were not required to participate in the task, and we provided an activity sheet and a short story to allow for a realistic alternative to participation.⁵ In the final dataset, 11 individuals completed no tasks at all during the public good game, and around 13% completed fewer than 20 tasks, compared with a mean of 47 and a maximum of 103. Even after adjusting for ability, task completion varies from 0 to 63, suggesting that there is meaningful variation in effort.

4. Post-Treatment Survey

After the activity was completed, participants were told how well they performed, how much money the team earned in total, and what their share of the winnings was. They were then asked to once again rate the enjoyment and difficulty of the task, and they were also asked questions to assess their perceptions regarding autonomy, fairness, agreement with the outcomes, willingness to work with the team again, etc. Finally, they were also given an opportunity to “reinvest” some of their earnings back into the team pot, which would be multiplied before being redistributed evenly across the team, as in a more traditional public good game. This was intended to be a measure of complementary investment.

Analytic Methods

My main analyses use individual performance on the effort task (contribution to the public good) as the primary dependent variable, with the decision-making treatments as the key explanatory variables. This main analysis takes the form of a linear model with robust standard errors clustered by team (the randomly assigned five-person group with which individuals generated their public good).⁶ After testing for differences between the coefficients of the public and private vote groups, my final model collapses the

⁴ The exchange rate is approximately 100 Shillings = 1USD.

⁵ This aspect of the design does not appear in the original Pre-Analysis Plan, but was added after initial piloting suggested that little free-riding was occurring. The rest of the study protocol is identical to that included as an appendix to the Pre-Analysis Plan.

⁶ This differs slightly from what was specified in the Pre-Analysis Plan, which called for poisson models (to account for the bounded nature of the outcome variable). Since the Pre-Analysis Plan was written before any data had been collected, I could not test for the appropriateness of this modeling approach at the time. Data in hand, it was apparent that (1) the mean of the outcome variable did not equal the variance, rendering a poisson model inappropriate. However, for the sake of transparency, I include the original pre-registered version of the analysis in Appendix D, which shows findings consistent with what is presented in the main text of the paper.

treatments into only three groups: control, majority rule and deliberation. My main model also includes several pre-registered control variables that are known to influence cooperation in a public goods setting.⁷

As shown in the balance table in Appendix A, the individuals assigned to the various treatments were not statistically different ($\alpha=0.05$) from those in the control group with respect to any of the pre-treatment demographic variables, nor with respect to pre-treatment preferences or ability on the tasks. The only differences that appear across groups emerge subsequent to the treatment (with respect to the decisions made as a team and the effort individuals exerted on the team task). Since the randomization successfully balanced the treatment groups on all observable covariates, control variables are arguably unnecessary, but their inclusion (i) does not change the direction or statistical significance of my main results and (ii) results in a higher R^2 than excluding them, allowing for more precise estimates of the treatment effects.

The standard session included 15 participants randomly assigned to one of three teams, but in a few cases, too few participants showed up at the scheduled time, and sessions had to be conducted with 10 participants instead (with make-up sessions of 5 participants conducted later to reach the desired sample size). To account for potential bias resulting from this unplanned variation across sessions, I also include a control for the size of the session.

In addition to this primary analysis, I also examine several pre-registered secondary and intermediate outcomes. First, as outlined in the conceptual framework, I wished to distinguish between improved outcomes as a result of behavior change as opposed to decision quality. The variable *task-specific ability* serves primarily as a way to distinguish between effects through decision quality (choosing the task that the team was already collectively best at) and effects through behavior change (investing more effort on behalf of the collective good, irrespective of ability). *Task-specific ability* is measured through the individually incentivized version of the effort task that was conducted during the pre-treatment survey. Since average ability on the tasks did not vary across treatment groups prior to the treatment, task-specific ability should only mediate outcomes through better task selection. This variable thus helps me to establish whether the effect has occurred through better decision quality, rather than through increased effort (behavior change).

To further explore the sub-mechanisms through which any improved collective outcome is achieved, I also analyze several additional explanatory variables, which serve as potential causal mediators. These are presented in These additional variables include: (1) preference alignment, (2) preference change, (3) procedural justice, and (4) acceptance of the team decision.

In the pre- and post-treatment surveys, individuals were asked to name the task they would choose if given the option, and they were also asked to rank the tasks according to difficulty. The *preference alignment* variable indicates whether the initial individual task preference identified in the pre-treatment survey matches with the ultimate task selected by that individual's team. If preference alignment mediates the outcome, it could mean that either people tended to prefer the task they were best at or that people invest more effort into the task because they got what they wanted in the first place. *Preference change* measures whether an individual who did not initially prefer their team's selection later changes their mind, selecting the team choice in the post-treatment survey. This suggests that they were persuaded to agree with their teammates about the task selection, which could lead them to invest more effort into the task they have now been convinced to prefer.

⁷ These control variables include: gender, age, education, marital status, number of children, whether the participant engaged in any real world collective action within the past month, whether the participant knew others in their randomly assigned team, the proportion of women in the team, and the number of co-ethnics on the team.

In the post-treatment survey, participants were asked the following questions: “There were three tasks presented to you earlier, but only one was chosen for your team. How much did you agree with the final decision?” and “How fair do you think the decision was to choose a task for your team?” Each question was scored using a 5-point likert scale. Responses to the latter is interpreted as a measure of *perceived fairness*. *Acceptance of decision* is assessed based on responses to the former question, with an additional control for whether the individual’s final preference directly matched the team decision. These questions are meant to assess the notion of procedural justice, which implies that a perception of legitimacy in the process of decision-making could lead people to comply with decision-outcomes, even absent preference alignment or preference change.

Table 2: Summary of Outcome Variables Used in Analyses

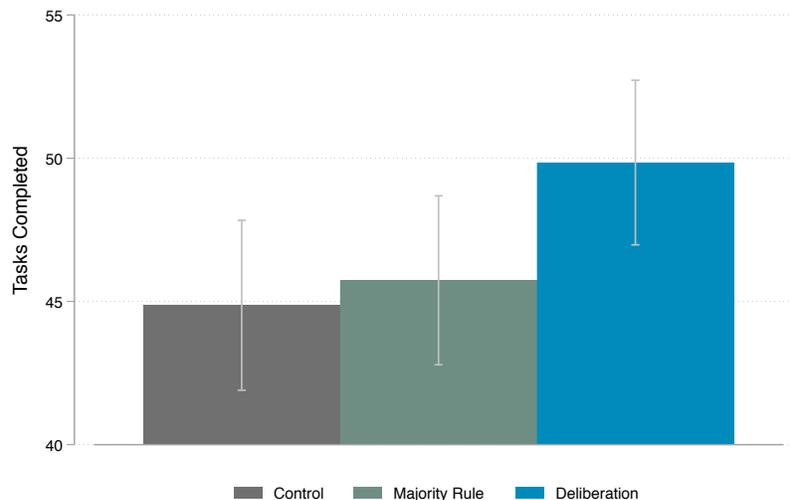
Primary Outcome Variables	
Contributions to the public good	Individual performance on the team effort task, which takes the form of a public good. Continuous, bounded at zero.
Effort	Same as contributions, but controlling for ability. Continuous, bounded at zero.
Ability	Pre-treatment individual performance on the individually incentivized version of the same effort task chosen by the team during the decision-making treatment. Continuous, bounded at zero.
Potential Causal Mediators	
Preference Alignment	Binary variable indicating whether an individual’s initial task preference was that selected by his/her team.
Preference Change	Binary variable indicating whether an individual shifted from pre-treatment preference that differed from team choice to a post-treatment preference that matched the team choice
Perceived Fairness	Response to a post-treatment likert-scale question (1-5) asking “How fair do you think the decision was to choose a task for your team?”
Acceptance of Decision	Response to a post-treatment likert-scale question (1-5) asking “How much did you agree with the final decision?”

Table 2 summarizes all of the outcome and mediating variables that were used in the analysis. In subsequent analyses using the mediators, I first tested for the effect of the treatments on each potential mediator, and then I tested for the effect of each potential mediator on the main outcome variables (individual performance on the team effort task, both with and without additionally controlling for individual ability). These regression analyses follow the same general approach described above (robust, clustered standard errors by team and the inclusion of pre-registered control variables). However, they do differ based on model type, with probit models being used in the case of binary outcome variables.

Finally, where the prior analyses indicate that it is appropriate, I use causal mediation analysis to test whether there is a mediated effect of the treatment through the hypothesized mediating variables.⁸ For this part of the analysis, I used the stata package “mediation” (Hicks & Tingley 2011), which implements the algorithm described in Imai, Keele & Tingley 2010 and Imai, Keele & Yamamoto 2010. This approach requires the assumption of sequential ignorability (Imai et al. 2011) which could be violated if there is an unobserved confounder that influences both an individual’s propensity for preference change (or other intermediate variable) and his/her willingness to exert effort. To address this, in the model predicting the effect of the mediator on the final outcome, I also include a suite of pre-treatment control variables that are known to influence cooperation in public good games. I also conduct sensitivity analyses to provide a sense for the likelihood of the assumption being violated. Additionally, I use individual task-specific ability as a control variable in cases where I wish to isolate the effect of a different mediator on the effort an individual exerted (as opposed to ability through superior task selection). In the mediation analyses, I use a binary treatment variable indicating whether an individual belonged to either of the deliberation treatments. This means that in contrast to the main regression models, the comparison group includes the majority rule treatments. Thus, these could be considered particularly conservative tests of mediated effects.

RESULTS

Figure 3: Mean Contribution to the Public Good
(Individual Output on Team Effort Task by Treatment Group)



Deliberation leads to better collective outcomes (higher earnings for the team)

Individual performance on the team effort task is significantly better, on average, in the deliberation treatments than in the control group. Since better performance translates directly into higher earnings for the team, this is objectively the superior outcome for all involved in this controlled setting. This is apparent even from basic descriptive statistics showing mean contributions across the treatment groups (See Figure 3), and is corroborated by regression analysis (See Appendix C). On average, individuals in the deliberation treatment completed five more tasks correctly (earning an additional 25 Shillings for their

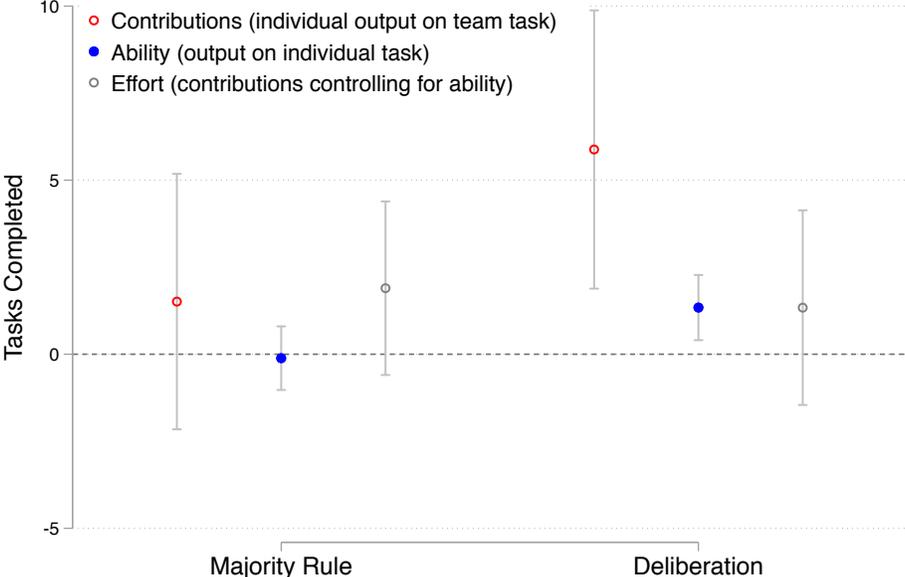
⁸ While the experimental design, primary analyses and intermediate outcomes were pre-registered, I did not pre-register the use of mediation analysis as described in Imai et al. 2010. However, these analyses are firmly grounded in the theory described earlier.

team) as compared with the control group, whereas the majority rules treatment groups do not differ significantly from the control group. Contributions are defined here as the number of tasks completed by an individual on the team-level incentivized effort task. Results from the main analysis suggest that participatory decision-making (when that decision-making takes the form of reason-based deliberative argumentation leading to a consensus-based decision) does result in better collective outcomes (increased investments in the public good and greater earnings for members of the team).

Deliberation improves decision quality (leads to more socially optimal decisions)

While individuals in the deliberation group perform better on the team effort task (higher contributions), they also have higher task-specific ability (performance on the individually incentivized version of the same task chosen by the team). When I control for task-specific ability in the main model (See Appendix E), the statistical significance of the effect on contributions disappears (although the coefficient is still positive). Figure 4 shows treatment effects on contributions (individual performance on the team effort task), ability (performance on the individually incentivized version of that same effort task), and effort (contributions, controlling for ability). None of the treatment effects for the majority rule treatment are statistically significantly different from zero, meaning that outcomes were no different from those in the control group. Deliberation, on the other hand, has a statistically significant effect on contributions and task-specific ability, but not on effort.

Figure 4: Treatment Effects - Decision Quality vs. Behavior Change

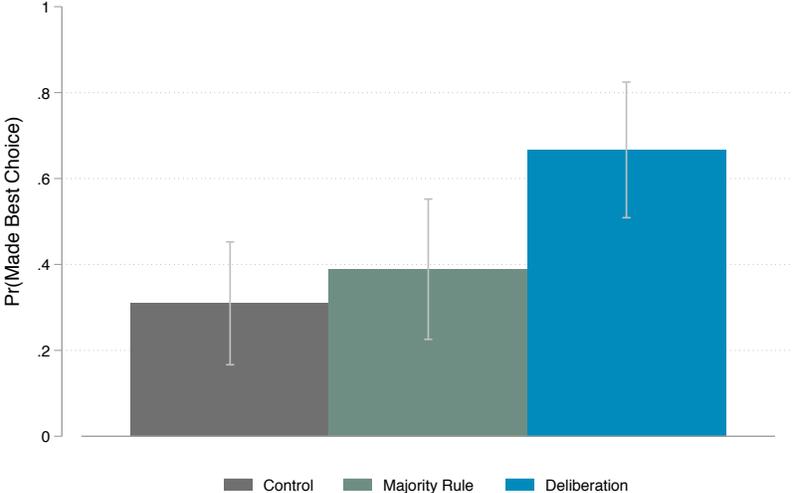


Balance tests (Appendix A) confirm that individuals in the deliberation groups were not, by random chance, more skillful at any of the three tasks compared with people in other groups. The effect of deliberation on ability can thus only have occurred through task selection. The teams in the deliberation groups were more likely to strategically choose the most socially optimal task for their particular team (i.e., the task at which they were most skillful as a collective).

To validate this further, I generated a new (not pre-registered) team-level variable, called “Chose Best Option for Team,” that indicates whether a team made the best decision for its members. By “best decision,” I mean the socially optimal decision for which they would have earned the most money based

on their collective ability in the individually compensated practice rounds. Being in one of the deliberation treatments is a statistically significant predictor of making the best team-level decision (See Appendix E). In terms of raw data, more than 65% of deliberation treatment teams choose the task that their team members performed collectively best on during the practice rounds (See Figure 5). Unsurprisingly, in the control group, where the tasks were randomly assigned, only around 30% of teams happened to be assigned the task that they were best at. However, in the majority rule treatment groups, teams chose the task they were best at in fewer than 40% of cases. Deliberation served the purpose of allowing teammates to share relevant information about ability and strategy and therefore make better decisions than they would have through individual calculation.

Figure 5: Decision Quality
(Proportion of Teams Choosing Most Profitable Task)



My data strongly supports the conclusion that the deliberation treatment led teams to more frequently choose the most appropriate task for their team – the task that was most profitable for their own mix of abilities. This suggests that, while individuals in the deliberation treatments do generate more contributions to the public good, the difference is mostly attributable to improved task selection (better decision-making), as opposed to effort (behavior change). It is important to note here that this does not entirely rule out the possibility of an effect on behavior, but it suggests that such an effect, if it exists, is too small across the full sample to be statistically significant. In particular, it must be smaller than 0.2 standard deviations, according to my power calculations (See Appendix B). However, it leaves open the possibility for heterogeneous treatment effects or effects mediated by an infrequent mediator.

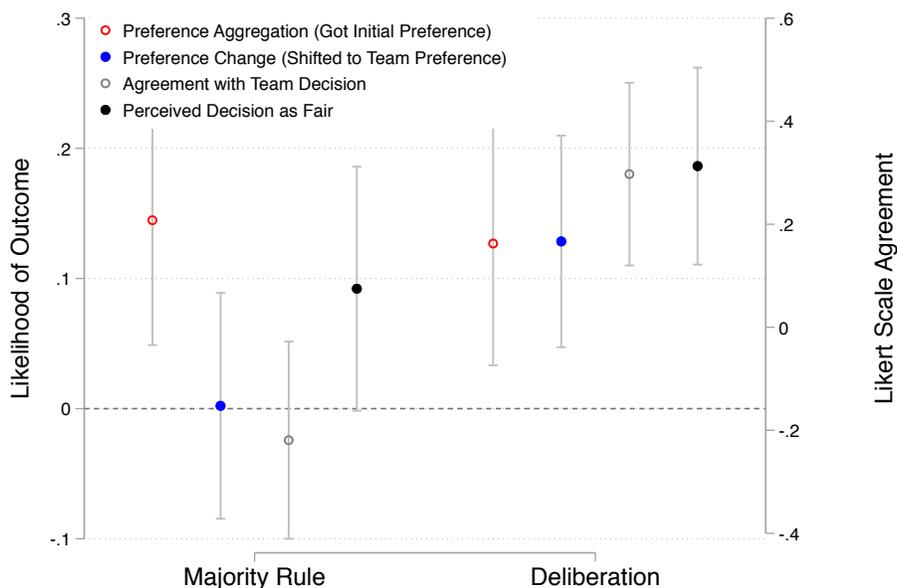
Deliberation increased preference change, perceived fairness and acceptance of outcomes

In addition to better decisions, deliberation was also associated with higher preference alignment, preference change, perceived fairness, and agreement with outcomes, as compared with the control group. Figure 6 shows the treatment effects of both the majority rule and deliberation treatments on each of these intermediate outcomes. (See Appendix F for the regression output.)

Individuals in the majority rule and deliberation treatments were both more likely to receive their initial preferred outcome as compared with the control group (where the outcomes were randomly assigned). However, this effect is larger in the majority rule treatment than in the deliberation groups (Appendix G). In the control group, about 36% of individuals happened to be randomly assigned the task that was their top choice according to the pre-treatment survey. In the deliberation treatment, about 48% of individuals

belonged to a team that ultimately selected their first choice of task, whereas in the majority rule groups, just over 51% of individuals got their way. The difference between the majority rule and deliberation groups on this outcome is not statistically significant, so preference alignment alone cannot explain superior outcomes in the deliberation group.

Figure 6: Treatment Effects – Intermediate Outcomes



Despite slightly more people in the majority rule group getting the outcome they initially preferred, individuals in the deliberation groups were more likely to perceive the process as having been fair, whereas in the majority rule treatments, the answer to this question was not significantly different from the control group (See Appendix G). Furthermore, individuals in the deliberation treatment were more likely to say that they agreed with the decision that was made, whereas participants in the majority rule treatment were actually less likely to agree with the decision than those surveyed in the control group (where the task was assigned using a random number generator). In the deliberation group, more than 80% report that they “completely agree” with the decision (the most extreme response on a 5-point likert scale), whereas in the majority rule group fewer than 60% select this response.

I also find that those in the deliberation groups are more likely to change their self-reported individual preferences over the tasks after the treatment. Those who have experienced the deliberation treatment are significantly more likely to change their preference to match the choice taken by their team. Figure 6 compares the match between the individual and team preference across the treatments. The left-hand figure shows, for each treatment group, the proportion of individuals whose initial preference was that chosen by their team, while the right-hand side shows the proportion of individuals whose final, stated preference on the post-treatment survey matches with what their team selected. In the deliberation group, 57 individuals (about 31%) changed their preference in favor of the collective team decision. This proportion is actually quite striking when we consider that almost 50% of them already preferred that task and thus could not possibly have shifted their preference in favor of it. In 30 of the 36 teams assigned to deliberation (83%), at least one individual was persuaded to favor the task chosen by collective decision.

In contrast, fewer than 20% of individuals shifted their preference to the team choice in the control and majority rule groups.

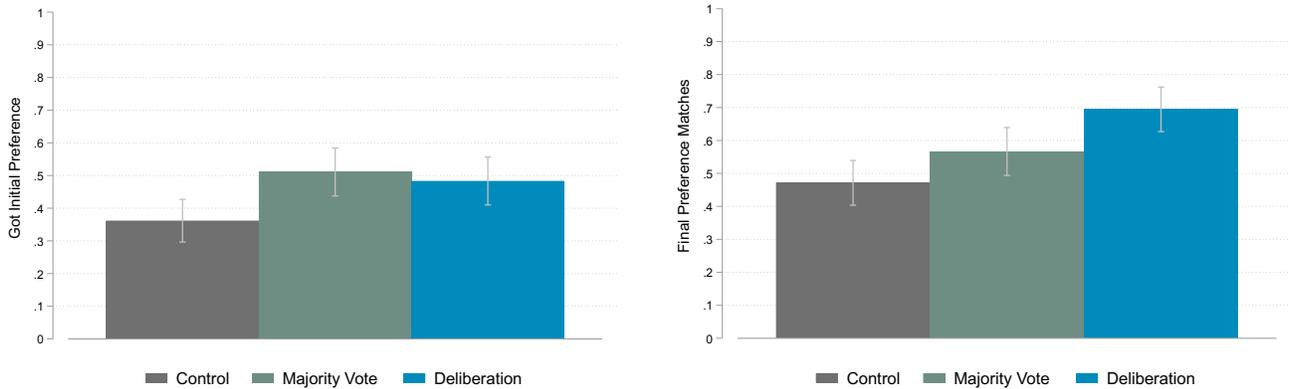


Figure 7: Preference Alignment & Change

(Left: Pre-Treatment Preference is Team Decision; Right: Post-Treatment Preference is Team Decision)

Deliberation may indirectly increase effort through preference change

Table 3 shows the estimated effect of each intermediate variable on my final outcomes (contributions, ability and effort). Getting one’s initial preference and self-stated agreement with the decision at endline are both associated with greater contributions and ability, but neither of these variables predicts effort. In fact, getting one’s initial preference predicts a *reduction* in effort, on average (though this result is not statistically significant), while those who were persuaded to prefer a task that they did not initially prefer actually tried harder during the team task. The only intermediate variable that is a statistically significant predictor of effort is preference change.

Table 3: Effects of intermediate variables on primary outcome pathways
(Estimated treatment effects and p-values)

	Contributions (decision quality)	Effort (behavior change)
Preference Alignment	6.520*** (p=0.000)	-1.637 (p=0.152)
Preference Change	0.0232 (p=0.991)	3.249* (p=0.026)
Perceived Fairness	1.068 (p=0.184)	0.433 (p=0.542)
Acceptance of Decision	1.645* (p=0.021)	0.874 (p=0.101)

For the most promising candidates for causal mediation (preference alignment on contributions and preference change on effort), I evaluated them using causal mediation analysis as well. With respect to preference alignment, there is a statistically significant average causal mediation effect (0.37) on ability when the deliberation treatment is compared to the control group, but when the comparison group includes the majority rule treatment, the statistical significance disappears. Thus, preference alignment does increase contributions to the public good through increasing the match between individual ability and the task selected, but it does not explain why the deliberation group outperforms the majority rule treatment with respect to decision quality. Though I cannot test this directly, it seems likely that this occurred through the information pooling mechanism.

In my main analysis presented earlier, there was no statistically significant effect of deliberation on behavior change (effort) but only on performance (through better collective decisions related to task ability). However, causal mediation does not necessarily require statistical significance in the main effect (Baron & Kenny 1986). Since preference change both is influenced by deliberation ($p < 0.01$) and also has an impact on effort ($p < 0.05$), this implies that there may be an indirect effect of deliberation on behavioral change, mediated by preference change. Even though the direct effect of deliberation on effort could not be clearly established, inconsistent mediation could still allow for a statistically significant indirect effect through the mediator (Zhao, Lynch & Chen 2010).

The causal mediation analysis suggests that the average causal mediation effect of preference change on effort is 0.42 (with the 95% confidence interval ranging from 0.03 to 1.00), accounting for around 23% of the total effect. This is statistically different from zero (even though the comparison group includes the majority rule treatment), so we can reject the null hypothesis that there is no mediated effect of deliberation on effort, as long as the sequential ignorability assumption holds. Sensitivity analysis demonstrates that the average causal mediation effect remains positive as long as the correlation between error terms in the two equations (predicting the mediator and the outcome, respectively) is less than 0.1.

Deliberation, when it manages to change preferences, also leads to changes in behavior. In particular, individuals put more effort into public good contributions if they have been persuaded by reason-based argumentation to change their minds about the best way to achieve those contributions. Since the causal mediator is not randomly assigned, this portion of the analysis lacks the full causal identification afforded by the experimental design. However, as changes to an individual's own preferences cannot credibly be randomly assigned, this is among the strongest empirical support possible in favor of altered preferences resulting in behavior change.

DISCUSSION

The experimental results demonstrate that individuals who engage in deliberative discussion involving reason-based argumentation achieve better collective outcomes. This occurs through two pathways. First, the deliberative teams make better strategic decisions regarding the collective good. Second, deliberation has an indirect effect through preference change that leads some individuals to invest more effort into contributing to the collective good.

The improvement in decision quality is not driven by mere preference aggregation. Individuals in the majority rules treatment are more likely to receive a task assignment that matches their pre-treatment preferences, but they are less effective at increasing contributions. Preference transformation does not fully explain these results either. While teams with more individuals who shifted preferences are somewhat more likely to make the best choice, controlling for preference change does not eliminate the effect of deliberation on decision quality. Even where individual preferences remain stable, teams manage to make the best collective choice through deliberation.

The quality of decisions was likely improved through information effects: either the aggregation or transformation of knowledge. Either is plausible in this experimental setting. All participants had some information about how well they performed on each task during the practice rounds, but they did not know how well their team members performed. If this information was revealed during deliberation, then this new information could have led to improved decision-making (aggregation of knowledge). It is also possible that, although participants had information about their relative level of ability on each task, they

were not basing their decision on that information prior to deliberation. They could, instead, have been persuaded through deliberation that this was the correct way to make the decision (transformation of knowledge) without necessarily altering their underlying personal preference for a given task.

Regarding the other pathway to improved collective outcomes, behavior change, the main analysis did not provide any direct support for increased effort as a result of deliberation. However, deliberation has a statistically significant indirect effect on behavior through preference change. After being persuaded through reason-based argumentation to a certain course of action, individuals put more effort into that action. This suggests that involving people in deliberative decisions that affect their lives may, irrespective of the decision that is ultimately made, change people's relationship to the decision outcomes.

Some specific features of the study context may have limited my ability to detect effects with respect to effort. First, although I made efforts in the study design to make the opportunity for freeriding more realistic, there is arguably very little opportunity cost to an individual's participation in the effort task. Since they have already planned to spend this time in the lab setting, the additional cost of participating in the activity may be negligible to them. (Relatedly, only about 2% of participants contribute nothing at all to the public good and the majority of these also failed to complete any tasks during the individually compensated task.)

Second, the study participants are a highly cooperative sample. More than 70% report having been involved in some form of real world collective action (participating in a community project or fundraiser event) within the past month. In my measure of complementary effort – the standard voluntary contribution mechanism version of the public good game – more than 70% of participants contribute at least half of their endowment, despite a more realistic opportunity cost (because the portion of the individual endowment not contributed may be kept).

If the majority of participants are already prone to exert their maximum effort, then that limits my ability to observe meaningful variation on effort across the treatment groups. This would imply that ability, not willingness to contribute, is the main constraint on public good provision in this context. However, this may actually be true of many resource-constrained communities in developing countries, where social capital is high but the successful collective action remains limited. If so, the finding that deliberation allows for improvements through better strategic decision-making may have direct relevance for policy-makers. Future research should aim to replicate this experimental design to the extent possible in a field setting in multiple contexts which vary based on pre-study predisposition to cooperation.

Given the cooperative nature of the study context, I would consider this a particularly hard test of the effects of deliberation on behavior change. Thus, the observable impact on net contributions through improved decision-making and preference transformation is noteworthy. Combined with the statistically significant secondary effects on perceptions of fairness and agreement with outcomes, I view these results as supportive of further research into the potentially transformative effects of deliberation as applied to collective action problems.

This result has major policy implications, corroborating the hunch of many a grassroots development practitioner. The study outcomes are very supportive of the use of deliberative processes in the decisions leading up to the creation of a public good. However, the author cautions that those planning to design participatory decision processes such as these should still carefully weigh both the costs and benefits of the process in their design. This experiment validates the existence of hypothesized effects but it cannot

speak to whether the magnitude of the effects would provide a benefit outweighing the opportunity cost of participants' time in a more realistic field setting.

One contribution of this work is to bring the public policy literature on participatory decision-making into dialogue with democratic theory and experimental work from political psychology and behavioral economics. I demonstrate here that deliberative discussion leads to decision-making that is more than the sum of its parts. Collective rationality is perhaps less bounded than that of any one individual. In addition, there is experimental evidence for the notion that preferences may be transformed through the process of deliberation and suggestive evidence that those who experience this transformation may indeed become 'better citizens' in the sense that they behave in more socially optimal ways.

References

- Ackerman, B., & Fishkin, J. S. (2002). Deliberation Day. *Journal of Political Philosophy*, 10(2), 129-152.
- Aga, D. A., Noorderhaven, N., & Vallejo, B. 2017. Project Beneficiary Participation and Behavioural Intentions Promoting Project Sustainability: The Mediating Role of Psychological Ownership. *Development Policy Review*.
- Agrawal, A. (2005). Environmentalism: Community, intimate government, and the making of environmental subjects in Kumaon, India. *Current anthropology*, 46(2), 161-190.
- Arnstein, S. R. 1969. A ladder of citizen participation. *Journal of the American Institute of planners*, 35(4), 216-224.
- Bächtiger, A., Niemeyer, S., Neblo, M., Steenbergen, M. R., & Steiner, J. (2010). Disentangling diversity in deliberative democracy: Competing theories, their blind spots and complementarities. *Journal of Political Philosophy*, 18(1), 32-63
- Barabas, J. (2004). How deliberation affects policy opinions. *American Political Science Review*, 98(4), 687-701.
- Baron, R. M., & Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic and statistical considerations. *Journal of Personality and Social Psychology*, 51, 1173-1182.
- Beard, V. A., Miraftab, F., & Silver, C. (Eds.). 2008. *Planning and decentralization: contested spaces for public action in the global South*. Routledge.
- Beath, A., Christia, F., & Enikolopov, R. (2017). Direct democracy and resource allocation: Experimental evidence from Afghanistan. *Journal of Development Economics*, 124, 199-213.
- Beierle, T. C., & Cayford, J. 2002. *Democracy in practice: Public participation in environmental decisions*. Resources for the Future.
- Berg, Monika & Rolf Lidskog (2018) Deliberative democracy meets democratised science: a deliberative systems approach to global environmental governance, *Environmental Politics*, 27:1, 1-20, DOI: 10.1080/09644016.2017.1371919
- Biggs, S. D. 1989. Resource-poor farmer participation in research: A synthesis of experiences from nine national agricultural research systems. *OFCOR-Comparative Study (Países Bajos)*.
- Birnbaum, S. 2016. Environmental co-governance, legitimacy, and the quest for compliance: when and why is stakeholder participation desirable?. *Journal of Environmental Policy & Planning*, 18(3), 306-323

- Bornstein, G. (1992). The free-rider problem in intergroup conflicts over step-level and continuous public goods. *Journal of Personality and Social Psychology*, 62(4), 597.
- Brooks, J. S., Waylen, K. A., & Mulder, M. B. (2012). How national context, project design, and local community characteristics influence success in community-based conservation projects. *Proceedings of the National Academy of Sciences*, 109(52), 21265-21270.
- Caluwaerts, D., & Ugarriza, J. (2012). Favorable conditions to epistemic validity in deliberative experiments: A methodological assessment. *Journal of Public Deliberation*, 8(1).
- Cardenas, J. C. (2004). Norms from outside and from inside: an experimental analysis on the governance of local ecosystems. *Forest Policy and Economics*, 6(3), 229-241.
- Carpini, M. X. D., Cook, F. L., & Jacobs, L. R. (2004). Public deliberation, discursive participation, and citizen engagement: A review of the empirical literature. *Annu. Rev. Polit. Sci.*, 7, 315-344.
- Chambers, S. (2003). Deliberative democratic theory. *Annual review of political science*, 6(1), 307-326
- Chib, V. S., Rangel, A., Shimojo, S., & O'Doherty, J. P. (2009). Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex. *The Journal of neuroscience*, 29(39), 12315-12320.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). New Jersey: Lawrence Erlbaum.
- Cohen, Joshua. 1989. "Deliberation and Democratic Legitimacy." In *The Good Policy*, ed. Alan Hamlin and Philip Pettit. Oxford: Basil Blackwell, 17–34.
- Cohen, Joshua. 2009. "Truth and Public Reason." *Philosophy and Public Affairs* 37 (1): 2–42.
- Dawes, R. M., Van de Kragt, A. J., & Orbell, J. M. (1990). Cooperation for the benefit of us—Not me, or my conscience. In *Beyond Self-Interest*, ed J Mansbridge. University of Chicago Press.
- Dietz, T. (2013). Bringing values and deliberation to science communication. *Proceedings of the National Academy of Sciences*, 110(Supplement 3), 14081-14087.
- Druckman, J. N., & Nelson, K. R. (2003). Framing and deliberation: How citizens' conversations limit elite influence. *American Journal of Political Science*, 47(4), 729-745.
- Dryzek, John S., André Bächtiger, Simone Chambers, Joshua Cohen, James N. Druckman, Andrea Felicetti, James S. Fishkin, David M. Farrell, Archon Fung, Amy Gutmann, Hélene Landemore, Jane Mansbridge Sofie Marien, Michael A. Neblo, Simon Niemeyer, Maija Setälä, Rune Slothuus, Jane Suiter, Dennis Thompson and Mark E. Warren. 2019. The crisis of democracy and the science of deliberation. *Science* 363 (6432), 1144-1146.
- Dryzek, J. S. (1994). *Discursive democracy: Politics, policy, and political science*. Cambridge University Press
- Duit, A., & Hall, O. (2014). Causes and Consequences of Stakeholder Participation in Natural Resource Management: Evidence from 143 Biosphere Reserves in Fifty-Five Countries. *State and environment: The comparative study of environmental governance*, 293.
- Elster, J. 1986. The Market and the Forum. Pp. 103–32 in J. Elster and A. Hylland (eds), *Foundations of Social Choice Theory*. Cambridge: Cambridge University Press.
- Esterling, K. M., Neblo, M. A., & Lazer, D. M. (2011). Means, motive, and opportunity in becoming informed about politics: A deliberative field experiment with members of Congress and their constituents. *Public opinion quarterly*, 75(3), 483-503.
- Estlund, David. 1998. "The Insularity of the Reasonable: Why Political Liberalism Must Admit the Truth." *Ethics* 108 (2): 252–275.
- Falk, A., & Heckman, J. J. (2009). Lab experiments are a major source of knowledge in the social sciences. *Science*, 326(5952), 535-538.

- Farrar, C., Fishkin, J. S., Green, D. P., List, C., Luskin, R. C., & Paluck, E. L. (2010). Disaggregating deliberation's effects: An experiment within a deliberative poll. *British journal of political science*, 40(2), 333-347.
- Farrington, J. (1998). *Organisational roles in farmer participatory research and extension: lessons from the last decade*. London: Overseas Development Institute.
- Fischbacher, Urs. (2007): z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics* 10(2), 171-178.
- Fishkin, J. S. (2018). *Democracy when the people are thinking: revitalizing our politics through public deliberation*. Oxford University Press.
- Fishkin, J. S., & Luskin, R. C. (2005). Experimenting with a democratic ideal: Deliberative polling and public opinion. *Acta politica*, 40(3), 284-298.
- Frey, B. S., Benz, M., & Stutzer, A. (2004). Introducing procedural utility: Not only what, but also how matters. *Journal of Institutional and Theoretical Economics JITE*, 160(3), 377-401.
- Fujiwara, T. & Wantchekon, L. (2013). Can informed public deliberation overcome clientelism? Experimental evidence from Benin. *American Economic Journal: Applied Economics* 5(4)241-55.
- Fung, Archon. 2006a. "Democratizing the Policy Process." in *Oxford handbook of Public Policy*.
- Fung, A. 2006b. Varieties of participation in complex governance. *Public administration review* 66(s1) 66-75.
- Fung, A., & Wright, E. O. (2001). Deepening democracy: innovations in empowered participatory governance. *Politics and society*, 29(1), 5-42.
- Fung, A., Wright, E. O., & Abers, R. (2003). *Deepening democracy: Institutional innovations in empowered participatory governance* (Vol. 4). Verso Books.
- Ganuzza, E., & Baiocchi, G. (2012). The power of ambiguity: How participatory budgeting travels the globe. *Journal of Public Deliberation*, 8(2), 8.
- Gerber, M., Bächtiger, A., Shikano, S., Reber, S., & Rohr, S. (2018). Deliberative abilities and influence in a transnational Deliberative Poll. *British journal of political science* 48(4)1093-1118.
- Gibson, C., & Woolcock, M. (2008). Empowerment, deliberative development, and local-level politics in Indonesia: Participatory projects as a source of countervailing power. *Studies in comparative international development*, 43(2), 151-180.
- Gill, D., & Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. *American Economic Review*, 102(1), 469-503.
- Glennerster, R., & Takavarasha, K. (2013). *Running randomized evaluations: A practical guide*. Princeton University Press.
- Goldfrank, B. (2012). The World Bank and the globalization of participatory budgeting. *Journal of Public Deliberation*, 8(2), 7.
- Grillos, T. (2015). *Participation, Power & Preferences in International Development*. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.
- Grillos, T. (2017). Participatory budgeting and the poor: Tracing bias in a multi-staged process in Solo, Indonesia. *World Development*, 96, 343-358.
- Grillos, T. (2018). Women's participation in environmental decision-making: Quasi-experimental evidence from northern Kenya. *World Development*, 108, 115-130.
- Grönlund, K., Herne, K., & Setälä, M. (2015). Does enclave deliberation polarize opinions?. *Political Behavior*, 37(4), 995-1020.
- Grossman, G., & Baldassarri, D. (2012). The impact of elections on cooperation: Evidence from a lab-in-the-field experiment in Uganda. *American journal of political science*, 56(4), 964-985.
- Gutman, A., & Thompson, D. (2004). *Why deliberative democracy*. New Jersey: Princeton University.
- Habermas, J. (1981). *Theorie des kommunikativen Handelns: band 1-handlungsrationalität und gesellschaftliche Rationalisierung*. Suhrkamp Verlag.

- Habermas, Jurgen. (1989). " *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. Cambridge, MA: MIT Press.
- Habermas, J. (1996). *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*.
- Habyarimana, J., Humphreys, M., Posner, D. N., & Weinstein, J. M. (2009). Chapter 2: Public Goods Provision in Kampala. *Coethnicity: Diversity and the dilemmas of collective action*. Russell Sage Foundation
- Hafer, C., & Landa, D. (2007). Deliberation as self-discovery and institutions for political speech. *Journal of theoretical Politics*, 19(3), 329-360.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Beyond WEIRD: Towards a broad-based behavioral science. *Behavioral and Brain Sciences*, 33(2-3), 111-135.
- Hicks, R., & Tingley, D. (2011). Causal mediation analysis. *The Stata Journal*, 11(4), 605-619.
- Hong, S., & Cho, B. S. (2018). Citizen participation and the redistribution of public goods. *Public Administration*, 96(3), 481-496.
- Imai, K., L. Keele, and D. Tingley. (2010). A general approach to causal mediation analysis. *Psychological Methods* 15: 309–334
- Imai, K., Keele, L., Tingley, D., & Yamamoto, T. (2011). Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *American Political Science Review*, 105(4), 765-789.
- Imai, K., L. Keele, and T. Yamamoto. 2010. Identification, inference, and sensitivity analysis for causal mediation effects. *Statistical Science* 25: 51–71.
- Irvin, R. A., & Stansbury, J. (2004). Citizen participation in decision making: Is it worth the effort?. *Public administration review*, 64(1), 55-65.
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams Jr, R. B., Alper, S., ... & Batra, R. (2018). Many Labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443-490.
- Koontz, T. M., & Thomas, C. W. (2006). What do we know and need to know about the environmental outcomes of collaborative management?. *Public administration review*, 66(s1), 111-121.
- Landa, D. (2015). Behavioral political economy, argumentation, and democratic theory. *The Good Society*, 24(1), 86-97.
- Lawrence, A. (2006). 'No personal motive?' Volunteers, biodiversity, and the false dichotomies of participation. *Ethics Place and Environment*, 9(3), 279-298.
- Lazer, D. M., Sokhey, A. E., Neblo, M. A., Esterling, K. M., & Kennedy, R. (2015). Expanding the conversation: Multiplier effects from a deliberative field experiment. *Political Communication*, 32(4), 552-573.
- Ledyard, J. O. (1995). Public goods: A survey of experimental research. In J. Kagel & A. Roth (Eds.), *The Handbook of Experimental Economics* (pp. 111-194.). Princeton: Princeton University Press.
- Leotti, L. A., & Delgado, M. R. (2011). The inherent reward of choice. *Psychological Science*.
- Leotti, L. A., Iyengar, S. S., & Ochsner, K. N. (2010). Born to choose: The origins and value of the need for control. *Trends in cognitive sciences*, 14(10), 457-463.
- Levitt, S. D., & List, J. A. (2007). What do laboratory experiments measuring social preferences reveal about the real world?. *Journal of Economic perspectives*, 21(2), 153-174.
- Lind, E. A., & Tyler, T. R. (1988). *The social psychology of procedural justice*. Springer Science & Business Media.
- List, C., Luskin, R. C., Fishkin, J. S., & McLean, I. (2012). Deliberation, single-peakedness, and the possibility of meaningful democracy: evidence from deliberative polls. *The journal of politics*, 75(1), 80-95.

- Lopez, M. C., & Villamayor-Tomas, S. (2017). Understanding the black box of communication in a common-pool resource field experiment. *Environmental Science & Policy*, 68, 69-79
- Lubell, M. (2004). Collaborative environmental institutions: All talk and no action?. *Journal of policy analysis and management*, 23(3), 549-573.
- Luskin, R. C., Fishkin, J. S., & Jowell, R. (2002). Considered opinions: Deliberative polling in Britain. *British Journal of Political Science*, 32(3), 455-487.
- Manin, B. (1987). On legitimacy and political deliberation. *Political theory*, 15(3), 338-368.
- Mansbridge, J. (1999). On the idea that participation makes better citizens. *Citizen competence and democratic institutions*, 291-325.
- Mansbridge, J., Bohman, J., Chambers, S., Estlund, D., Føllesdal, A., Fung, A., & Manin, B. (2010). The Place of Self-Interest and the Role of Power in Deliberative Democracy. *Journal of political philosophy*, 18(1), 64-100.
- Mansuri, G., & Rao, V. (2004). *Community-based (and driven) development: A critical review* (Vol. 3209). World Bank.
- Mansuri, G., & Rao, V. (2012). *Localizing development: does participation work?* World Bank.
- Martí, J. L. (2006). The epistemic conception of deliberative democracy defended reasons, rightness and equal political. *Deliberative democracy and its discontents*, 27.
- Miguel, E., & Gugerty, M. K. (2005). Ethnic diversity, social sanctions, and public goods in Kenya. *Journal of public Economics*, 89(11-12), 2325-2368.
- Murayama, K., Matsumoto, M., Izuma, K., Sugiura, A., Ryan, R. M., Deci, E. L., & Matsumoto, K. (2013). How self-determined choice facilitates performance: a key role of the ventromedial prefrontal cortex. *Cerebral Cortex*, bht317.
- Mutz, D. C. (2008). Is deliberative democracy a falsifiable theory?. *Annu. Rev. Polit. Sci.*, 11, 521-538.
- Neblo, M. (2005). Thinking through democracy: between the theory and practice of deliberative politics. *Acta politica*, 40(2), 169-181.
- Neblo, M. A. (2007). Change for the better? Linking the mechanisms of deliberative opinion change to normative theory. *Common voices: The problems and promise of a deliberative democracy*.
- Neblo, M. A., Esterling, K. M., & Lazer, D. M. (2018). *Politics with the people: Building a directly representative democracy*(Vol. 555). Cambridge University Press.
- Niemeyer, S. (2011). The emancipatory effect of deliberation: empirical lessons from mini-publics. *Politics & Society*, 39(1), 103-140.
- Norton, M. I., Mochon, D., & Ariely, D. (2012). The IKEA effect: When labor leads to love. *Journal of Consumer Psychology*, 22, 453-460.
- Olken, B. A. (2010). Direct democracy and local public goods: Evidence from a field experiment in Indonesia. *American Political Science Review*, 104(02), 243-267.
- Ostrom, E., Gardner, R., & Walker, J. (1994). *Rules, games, and common-pool resources*. University of Michigan Press.
- Ostrom, E. (2005). *Understanding institutional diversity*. Princeton university press.
- Ostrom, E. (1990). Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*.
- Ostrom, E. (1998). A behavioral approach to the rational choice theory of collective action: Presidential address, APSA, 1997. *American political science review*, 92(1), 1-22.
- Plassmann, H., O'Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *The Journal of neuroscience*, 27(37), 9984-9988.
- Pretty, J. N. (1995). Participatory learning for sustainable agriculture. *World development* 23(8)1247-63.
- Patall, E. A., Cooper, H., & Robinson, J. C. (2008). The effects of choice on intrinsic motivation and related outcomes: a meta-analysis of research findings. *Psychological bulletin*, 134(2), 270.

- Price, K. H., Hall, T. W., van den Bos, K., Hunton, J. E., Lovett, S., & Tippett, M. J. (2001). Features of the value function for voice and their consistency across participants from four countries: Great Britain, Mexico, the Netherlands, and the United States. *Organizational Behavior and Human Decision Processes*, 84(1), 95-121.
- Reed, M. S. (2008). Stakeholder participation for environmental management: a literature review. *Biological conservation*, 141(10), 2417-2431.
- Rey-Biel, P., Sheremeta, R.M., & Uler, N. (2011). (Bad) luck or (lack of) effort?: Comparing social sharing norms between US and Europe. ESI Working Paper 11-11. Retrieved from http://digitalcommons.chapman.edu/esi_working_papers/97
- Rowe, G., & Frewer, L. J. (2000). Public participation methods: A framework for evaluation. *Science, technology & human values*, 25(1), 3-29.
- Ryfe, D. M. (2005). Does deliberative democracy work?. *Annu. Rev. Polit. Sci.*, 8, 49-71.
- Sally, D. (1995). Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. *Rationality and society*, 7(1), 58-92.
- Sanyal, Paromita and Vijayendra Rao. 2018. *Oral Democracy*. Cambridge University Press.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of experimental psychology*, 18(6), 643.
- Sunstein, C. R. (2000). Deliberative trouble? Why groups go to extremes. *The Yale Law Journal*, 110(1), 71-119.
- Sunstein, C. R. (2002). The law of group polarization. *Journal of political philosophy*, 10(2), 175-195.
- Sunstein, C. R., & Hastie, R. (2015). *Wiser: Getting beyond groupthink to make groups smarter*. Harvard Business Press.
- Thibaut, J. W., & Walker, L. (1975). *Procedural justice: A psychological analysis*. L. Erlbaum Ass.
- Thompson, D. F. (2008). Deliberative democratic theory and empirical political science. *Annu. Rev. Polit. Sci.*, 11, 497-520.
- Tyler, T. R. (2006). *Why people obey the law*. Princeton University Press.
- Ugarizza & Caluwaerts. (2014) *Democratic Deliberation in Deeply Divided Societies: From Conflict to Common Ground*.
- Wampler, B. and Touchton, M. (2017) Participatory budgeting: adoption and transformation, Making All Voices Count Research Briefing, Brighton: IDS
- Wilson, L. S. (1992). The Harambee movement and efficient public good provision in Kenya. *Journal of Public Economics*, 48(1), 1-19.
- Wondolleck, J. M., & Yaffee, S. L. (2000). *Making collaboration work: Lessons from innovation in natural resource management*. Island Press.
- Zhao, X., Lynch Jr, J. G., & Chen, Q. (2010). Reconsidering Baron and Kenny: Myths and truths about mediation analysis. *Journal of consumer research*, 37(2), 197-206.