

Deliberation improves collective decisions: Experimental evidence from Kenya

Tara Grillos
Department of Political Science
Purdue University

ABSTRACT

Citizen participation in decision-making has been widely lauded as a method for improving outcomes in international development. While there are normative reasons to encourage more inclusive decision-making processes, costly and time-intensive group decision-making processes are often justified on the grounds that they may also improve outcomes. Deliberative discussion, in particular, is believed to be more transformative than a mere aggregation of individual preferences, leading perhaps to more socially optimal decision making and subsequent behavior. Prior work confirms that deliberation results in shifts of opinion, but it has had little to say about the quality of the resulting decisions, which are difficult to assess in a field setting. I report the results from a laboratory experiment with 570 subjects in Nairobi, directly testing the effect of participation in deliberative group decision-making on collective outcomes. Participants are asked to engage in a group effort task to earn compensation toward a shared group fund. Randomly assigned treatments vary according to whether decision-making over the task to be completed occurs through (1) external assignment, (2) a majority vote, or (3) consensus through deliberative discussion. I find that participation in group decision-making involving deliberation (but not a simple majority vote) does improve collective outcomes. This effect is achieved primarily through better strategic decision making that minimizes the costs associated with contributions. Deliberation is also associated with changes in preferences, greater levels of agreement with decision outcomes, and greater perceived fairness. Evidence for behavior change is weaker, but there may be a positive effect mediated by preference change.

This is an unpublished manuscript. Recommended temporary citation for this working paper:

Grillos, Tara. 2019. Deliberation improves collective decisions: Experimental Evidence from Kenya. *Prepared for presentation at the conference of the American Political Science Association.*

Citizen participation in decision-making 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 nearly every global region, including East Asia (Hong & Cho 2018), SE Asia (Grillos 2017), Africa (Wampler & Touchton 2017) and North America (Lerner & Secondo 2012). Public participation is now formally encouraged in several national constitutions.

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 2006). 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 theorists extol the virtues of a particular form of citizen participation: deliberative democracy (Dryzek et al. 2019). Apart from normative benefits related to democratic principles of inclusion, deliberation, it is argued, serves at least two other core functions: an *epistemic* function, by improving informational quality and leading to better decisions, and also an *ethical* function, by promoting mutual respect (Mansbridge et al. 2012). The claim that deliberation may in fact make “better citizens” (Mansbridge 1999) is closely related to these epistemic and ethical functions, with the specific implications being that deliberation may lead citizens to (i) make more socially optimal decisions and (ii) engage in more socially optimal behaviors.

Yet empirical studies of deliberation have largely shied away from an explicit examination of outcomes, instead focusing on procedural factors (Landmore 2017). 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 either decision quality or behavior change. 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).

When decision-making concerns the allocation of resources to maximize some collectively valued outcome (as opposed to choosing between different possible goals), the outcome can be objectively assessed as being in the public interest or not. This study provides a rigorous experimental test of the hypothesis that participation in deliberative decision making improves collective outcomes. Using random assignment to different decision-making processes in a controlled laboratory setting, I assess whether

participation in collective decision making leads to socially optimal decisions and/or behavior. I specifically isolate the use of deliberative argumentation, as defined by deliberative theorists, as opposed to another commonly used form of collective decision making: preference aggregation through majority rule voting procedures.

I find strong experimental evidence in support of the epistemic benefits of deliberation. That is, deliberative argumentation leads to better decisions, which are more in line with the socially optimal outcome. Evidence is weak in support of related behavior change (investment of effort on behalf of collective outcomes). However, I offer some suggestive evidence that behavior change may occur in certain cases, specifically for individuals whose preferences have been changed through the process.

Related Literature

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). Efforts to systematize the 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. 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 specific features relative to their additional costs.

I follow previous scholars in defining participation as involvement in decision-making by those actors affected by the decisions in question (Reed, 2008). However, this notion of participation is still quite broad

and encompasses a number of very different group decision-making processes. Various scholars have offered typologies to help distinguish between different varieties of participation in decision making (Arnstein 1969, Biggs, 1989; Pretty 1995; Farrington, 1998; Lawrence 2006, Rowe & Frewer 2000). Fung (2006) argues that forms of participation vary along three key dimensions: the inclusiveness of participant selection, the level of authority given to decision outcomes, and the method of communication used to arrive at decisions.

Deliberative theory promotes a particular method of communication that involves rational argumentation (Bächtiger et al. 2010). Deliberative discussion is believed to be more transformative than a mere aggregation of individual preferences (Elster 1986, Chambers 2003), and it has the potential to lead to more socially optimal decision-making and to more socially optimal behaviors. 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). Deliberation has been proposed as an ideal for science communication (Dietz 2013) and global environmental governance (Bäckstrand 2010, Dryzek & Pickering 2017, Berg & Lidskog 2018.) But the literatures on participatory institutions and local collective action have only recently begun to engage directly with political theory on deliberative democracy (Heller & Rao 2015).

Deliberation may improve decision outcomes through several pathways. First, it may result in more socially oriented decision-making by constraining self-interest (Ackerman & Fishkin 2002, Mansbridge et al. 2010). Social pressure may make it difficult to rely on purely self-regarding arguments during the process of deliberation. Individuals may yield to the “forceless force” of the better argument (Habermas

1975, 1984) – leading to group decisions that are more in line with the collective good even if competing individual preferences remain intact. Second, deliberation may improve decisions by correcting information asymmetries. 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. Finally, the deliberative process may change 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).

There is also reason to expect that deliberation may lead to more socially optimal behavior, such as investments of effort toward the achievement of collective outcomes. 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. The procedural justice literature argues that people are more willing to behave in compliance with 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. While simpler forms of participation, such as majority rules voting procedures, may be sufficient to activate these mechanisms, one might expect that more intensive forms of engagement, such as deliberative discussion, would be more effective at doing so. Furthermore, if constrained self-interest is truly activated in the decision-making process, as argued above, then this constraint could also be internalized as a norm and thus nudge participants toward more pro-social behavior in the future as well.

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).

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). However, we still have very little causal evidence about whether or not deliberative argumentation, in particular, improves decision quality (Neblo 2007; Landemore 2017).

There are several major impediments to establishing these causal relationships with observational data: First, complex variation in program design makes it very difficult to isolate particular dimensions of participation (and related causal mechanisms) that could be driving any observed effects. Second, there is self-selection in both the creation of participatory institutions by policymakers and engagement in participatory processes by citizens, which creates concerns over reverse causality. Finally, socially desirable decisions and behavior – while theoretically reasonable as hypothesized outcomes of participation and deliberation – are nearly impossible to define in many real-world deliberative situations without imposing the values of the researcher. This area of research could thus benefit from more experimental research, which allows for the use of controlled variation (Falk & Heckman 2009).

Laboratory experiments have contributed greatly to the study of socially beneficial behavior in the form of cooperation for 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 cooperative behaviors (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. Some prior experimental work has examined the effects of group decision-making processes on cooperation (Olken 2010, Hamman, Weber & Woon 2011, Grossman & Baldassarri 2012, Aga et al. 2017), but none have looked specifically at how decision making that uses deliberative communication affects decision quality and behavior change.

Research Methods

I take advantage of a controlled laboratory setting to assess differences between commonly used approaches to group decision-making. This experiment involved a randomly assigned decision-making procedure, which in turn determined which of three effort tasks participants would receive. Individual performance on the effort task earned money toward a collectively shared outcome: a team fund which was then divided evenly across all team members, irrespective of individual performance. I differentiate between groups that rely on a simple majority rule voting procedure 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).

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). My experiment simulates the experience of groups of individuals coming together to decide how to allocate scarce resources (in this case, their time and effort) toward the achievement of collective outcomes.

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. 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 shared collective outcomes, on which there is no pre-treatment consensus. I further mitigate concerns by recruiting subjects from a relevant, developing country context.

Study Setting

The experiment was conducted at Busara Behavioral Lab in Nairobi, Kenya. Nairobi was chosen as the site of this research because it is located in a country where participatory institutions are currently being designed, and also because its people face scarcity of resources and thus would reasonably be concerned about the additional demands placed upon them by intensive forms of group decision making.

Kenya ratified a new constitution in 2010 by popular referendum. The new constitution includes requirements for citizen participation in government decision-making, but the specific method of public engagement was left largely to the discretion of the newly formed 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 the allocation of scarce resources to achieve collectively valued outcomes is extremely salient. In Kenya, community fundraisers known as *harambee* are a common form of local

collective action to provide local public goods and services (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.

The experimental protocol was conducted using z-Tree software (Fischbacher, 2007), with the exception of 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. Busara staff (Kenyan citizens) implemented the experimental protocol in Swahili, the national language. 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 is common. This experiment included 570 participants, spread across four treatments and one control group.¹ Sixty-three percent of the participants 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

¹ 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 (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.

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

The experimental design features a behavioral game intended to mimic some key aspects of real-world decision processes, particularly as often deployed through participatory budgeting. In particular, the research design involves groups of people deciding how to spend resources in order to maximize a group outcome. In the lab setting, the resources being allocated are the individuals' time and effort, and the group outcome takes the form of a pot of money, which allows for objective comparison across results.

Participants were first randomly assigned to a session, which was then randomly assigned to a treatment group. On the day of the experiment, participants engaged in a collective decision-making process within teams of five to determine which of several real-effort tasks they would participate in to earn money toward a shared 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. In the control group, the decision was made via random assignment.

The experiment took place in several stages, summarized below. The full experimental protocol documents are attached as Appendix K & L. The experimental design described below was registered as a Pre-Analysis Plan in the EGAP (Evidence in Governance and Politics) repository.²

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

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. Below I describe the four treatment groups, which were determined by two overlapping treatment variations: (i) public vs private voting procedures and (ii) majority rules voting vs deliberative consensus decision-making procedures.

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 team effort task 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. Team Effort Task

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 shillings³ 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 toward a collectively shared outcome (an aggregate team pot).

³ The exchange rate is approximately 100 Shillings = 1USD.

In the final dataset, 11 individuals completed no tasks at all during the team effort task, 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 real 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.

Analytic Methods

My main analyses use individual performance on the effort task (pro-social behavior as it contributes to the collective outcomes) 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 team with which individuals shared their collective earnings).⁴ After testing for differences between the coefficients of the public and private vote groups, my final model collapses the 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 pro-sociality.⁵

⁴ 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.

⁵ 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.

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 toward collective outcomes, 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.

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 1: Summary of Outcome Variables Used in Analyses

Primary Outcome Variables	
Contributions	Individual performance on the team effort task, which takes the form of contributions to a team fund. 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 a 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 1 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.

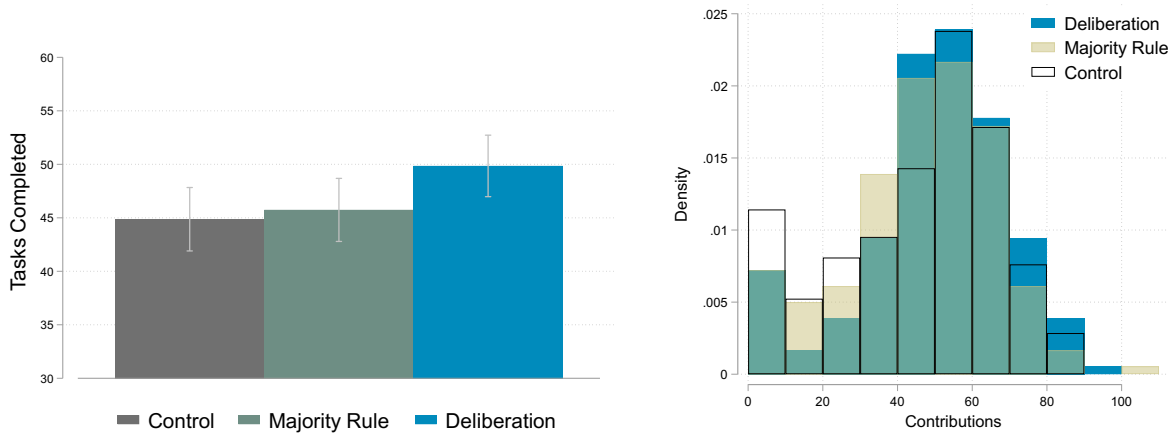
Finally, where the prior analyses indicate that it is appropriate, I use causal mediation analysis to test whether there is a mediated effect of the treatment through the hypothesized mediating variables.⁶ For this part of the analysis, I used the stata package “mediation” (Hicks & Tingley 2011), which 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 pro-social behavior. I then 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 behavior, by which I mean the effort an individual exerted (as opposed to greater ability achieved 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 or not. This means that in contrast to the main regression models, the

⁶ 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.

comparison group for the causal mediation analysis includes the majority rule treatments as well as the control group. Thus, these could be considered particularly conservative tests of mediated effects.

Results

Figure 1: Individual Output on Team Effort Task by Treatment Group
(Left: Mean Output, Right: Distribution of Output)

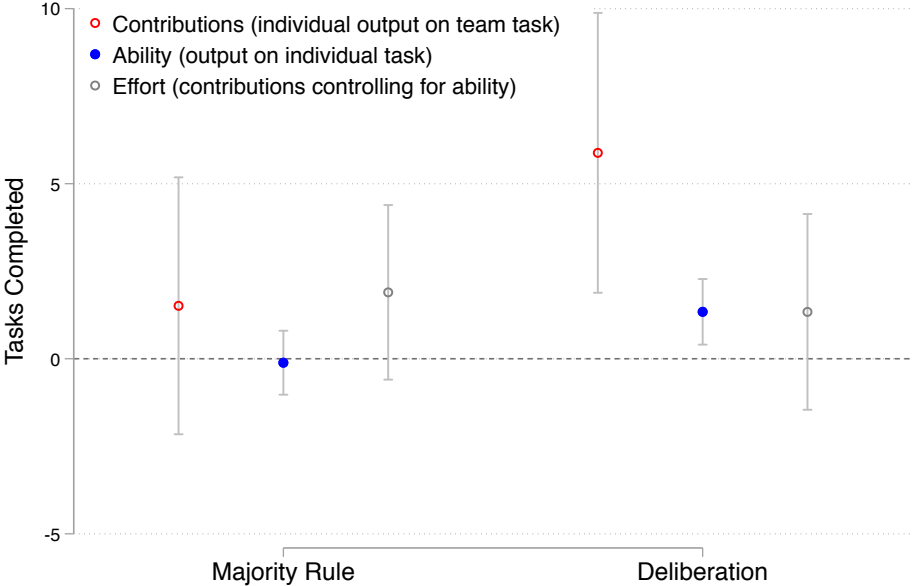


Deliberation improves collective outcomes (higher earnings for the team)

Contributions are defined here as the number of tasks completed by an individual on the team-level incentivized effort task. Individual performance on the team effort task is 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 the entire team within the context of this controlled setting. This is apparent even from basic descriptive statistics showing mean contributions across the treatment groups (See Figure 1, left-hand side), and is corroborated by regression analysis (See Appendix C). On average, individuals in the deliberation treatment completed around fifty tasks within the ten minutes allotted. This amounts to an additional five tasks completed correctly on average (corresponding to an additional 25 Shillings in earnings for their team) as compared with the control group, whereas the majority rule treatment groups do not differ significantly from the control group. The visual

distribution of outcomes (Figure 1, right-hand side) reveals that the effect is partially driven by a higher frequency of very low output in the control and majority rule groups. 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 effort investments on behalf of the team and thus greater earnings for members of the team).

Figure 2: Treatment Effects - Decision Quality vs. Behavior Change



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

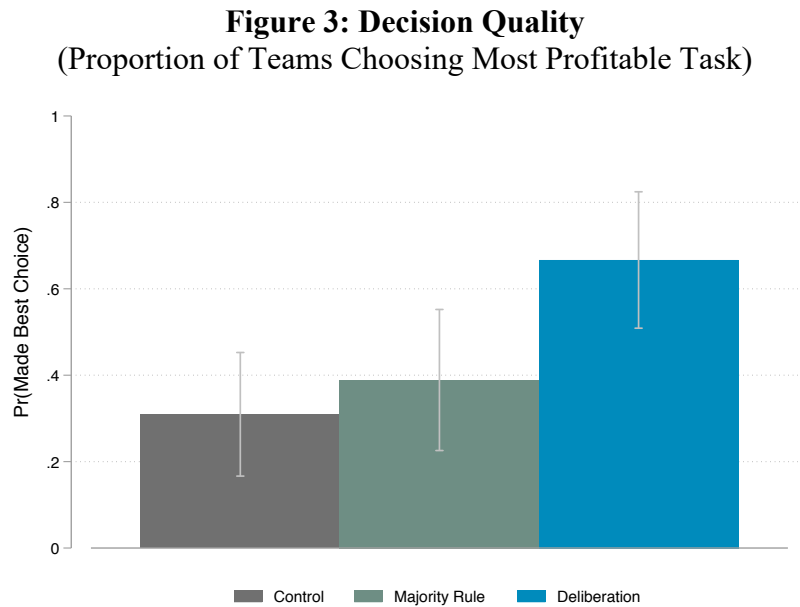
While individuals in the deliberation group perform better on the team effort task (higher contributions), they also have higher task-specific ability (performance on the individually incentivized version of the same task chosen by the team). When I control for task-specific ability in the main model (See Appendix E), the statistical significance of the effect on contributions disappears (although the coefficient is still positive). Figure 2 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.

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. “Best decision” here refers to the socially optimal decision - that for which they would have earned the most money based on their collective ability on each task. (Collective ability is judged based on performance on the individually compensated practice rounds for each task.) 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 3). 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.

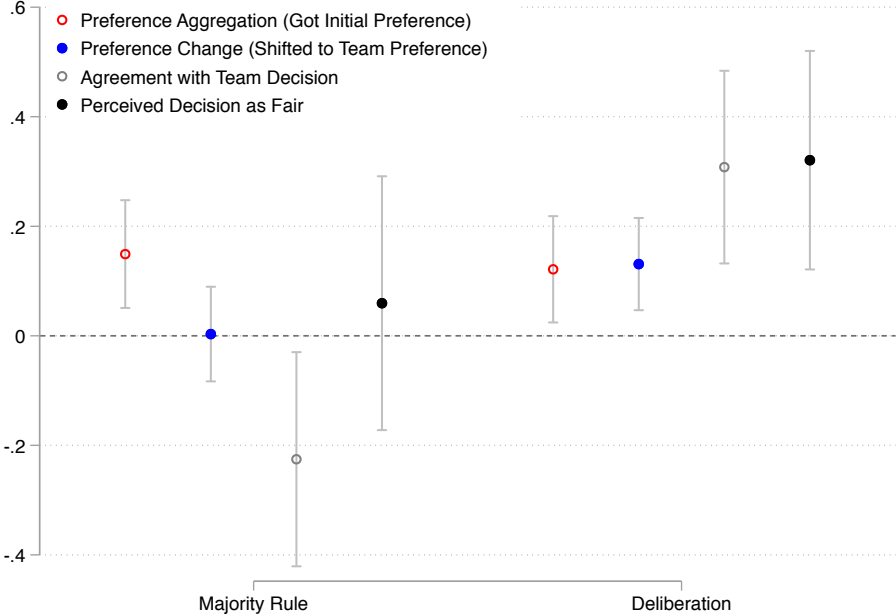


My data strongly supports the conclusion that the deliberation treatment led teams to more frequently choose the most locally appropriate task for their team – the task that was most profitable for their own mix of abilities. This suggests that, while individuals in the deliberation treatments do generate more contributions to the team fund, 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 4 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.)

Figure 4: Treatment Effects – Intermediate Outcomes



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.

Despite slightly more people in the majority rule group getting the outcome they initially preferred, individuals in the majority rule treatment actually exhibited lower levels of agreement with the decision that was made – even compared with the control group, where the task was assigned using a random number generator (See Appendix F). Individuals in the deliberation treatment, on the other hand, were more likely to say that they agreed with the decision that was made (an effect equivalent to about 0.3 points on a likert scale relative to the control group), and they were more likely to perceive the process as having been fair. In the majority rule treatments, the answer to the fairness question was not significantly different from the control group. 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 5 compares the match between the individual and team preference across the treatments. The left-hand column shows, for each treatment group, the proportion of individuals whose initial preference was that chosen by their team, while the right-hand column 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 switch their preference to 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.

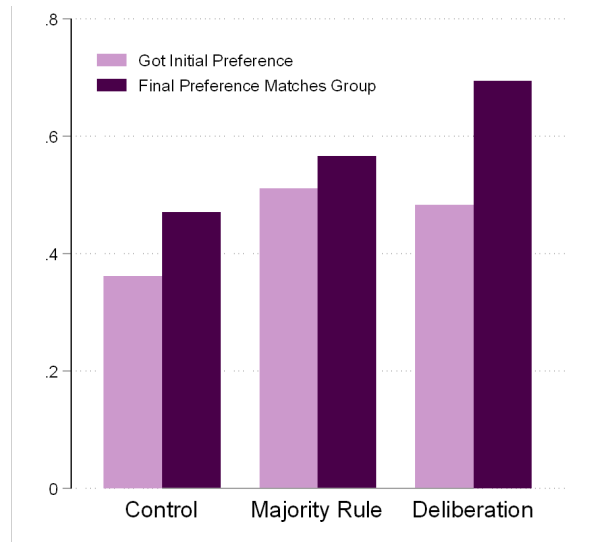


Figure 5: Preference Alignment & Change

Deliberation may indirectly increase effort through preference change

I also estimated the effect of each intermediate variable on the final outcomes: contributions and effort (Appendix G & H). 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 is associated with 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 performed better during the team task, after accounting for the ability they displayed on the individually incentivized version of that same task. Preference change is the only intermediate variable that is a statistically significant predictor of effort.

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 improve collective outcomes 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 improvements in information available.

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 (Hicks & Tingley 2011) 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 achieving collective outcomes if they have been persuaded by reason-based argumentation to change their minds about the best way to achieve those outcomes. 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

Some further exploratory analyses can help us begin to tease out the causal mechanisms underpinning the effects on both decision quality and behavior change. 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 likely to make the choice that would maximize collective contributions. Table 2 disaggregates teams in each treatment group by how well their initial individual preferences (based on the pre-treatment survey question) mapped onto the best (most profitable) choice for their team. Some teams already had a clear majority preference for the socially optimal choice (3 or more individuals preferred it). In other teams, the best choice was tied for first place (with 2 individuals listing it as their preference, equal to a second task). Still other teams had a clear majority preference that did not include the best choice. For all three categories of initial

preferences, the deliberation teams have a higher success rate in terms of arriving at the socially optimal choice through group decision making.

Table 2: Decision Success by Initial Preference Category

Initial Preferences for Best Choice Task	Top Preference (3 or more prefer it)		Tied for Top Choice (2 prefer it)		Not Preferred (>2 prefer another)	
	Maj.	Delib.	Maj.	Delib.	Maj.	Delib.
Total Teams	7	12	11	9	18	15
# Made Best Choice	5	9	5	6	4	9
Success Rate	0.71	0.75	0.45	0.67	0.22	0.60

Where the deliberation treatment especially outshines majority rules voting procedures is where there is an initial preference for a suboptimal decision. Of 18 such teams in the majority rules treatment, only 4 manage to arrive at an optimal solution.⁷ Of 15 such teams in the deliberation treatment, more than half arrive at the socially optimal decision despite initially having a majority preference for a different task.

Since preference aggregation alone cannot explain the outcomes, improvements in decision quality could be driven by some combination of (i) correction of information asymmetries, (ii) a change to the decision criteria being applied, and (iii) constrained self-interest.

It seems clear that information asymmetries exist, since individuals have some information about their own ability with respect to each task, but cannot know the ability of their teammates absent an explicit discussion with them. Participants are not very good at guessing where they stand relative to others in the

⁷ These are special cases in which, for example, there was a tie between the two non-optimal tasks based on initial preferences, and then one or more individuals voted differently from their top preference during the actual voting procedure.

room. After each of the practice rounds, individuals were asked to guess their rank on each task for a chance to earn a small cash bonus if they guessed correctly +/- 1 place out of 15. Fewer than 30% of participants managed to earn that bonus for any given task. With only 15 people in the room to compete against during any given session, a random guess should have yielded a 20% chance of being correct. This suggests that deliberation could have improved decision outcomes by allowing for information about others' abilities to be shared. This is supported by the fact that most of the "reasons" given during the reason-based argumentation related to how easy the task was. If this represented new information that was revealed during deliberation, then it could be one driver of improved decision-making.

It is also possible that, even those participants who had accurate information about their relative level of ability on each task 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 (altered decision criteria or correction of cognitive biases) without necessarily changing their underlying personal preference for a given task. While many of the arguments revolved around how easy the different task were, not all of them did – some referred to how enjoyable or even "challenging" their preferred tasks were. Furthermore, what participants perceive to be the "easiest" task may not be the task that is in fact the most profitable. This may be another example of information asymmetries, but it also could be that participants interpret "easy" differently than "most profitable." They may instead be referring to how much cognitive effort they feel while doing the task, even if they do not actually achieve the most points from that task. About 58% of the participants list as their top preference the same task that they rate as being "easiest" in the pre-treatment survey. However, only about 50% of participants rate as "easiest" the task that they actually performed best on (earned most money from) during the practice rounds.

With regards to constrained self-interest and the so-called “forceless force of the better argument”, I asked our facilitators to observe the deliberative discussions and code them according to whether the arguments primarily appealed to the good of the group, mostly referred to personal preferences, or if there was an equal mix of individual vs collective reason-giving. Most groups used a mix of the two, but in the 9 cases where deliberation teams overrode the initial majority preference to arrive at the socially optimal decision, more than half primarily used collectivist reasoning.

Since we only see effects on behavior through preference change, and not through perceptions of fairness, the evidence seems to be more supportive of procedural utility as an agent for pro-social behavioral change rather than procedural justice (though procedural justice may of course still play an important role in legitimacy more generally). Further supporting the procedural utility hypothesis is the fact that of those whose general preference shifted toward the task their team selected also are more likely to rate that task as the “most enjoyable” in the endline survey. Of those who shifted preferences, 88.46% of them rate that team task as the most enjoyable. (For comparison, 70.8% of them rated the task as “easiest” in the post-survey.) Of those preference-shifting participants who rate the task as most enjoyable in the post-treatment survey, only 22.6% percent of them already believed the task to be the most enjoyable in the pre-survey.

Conclusion

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 allocation of resources toward achieving collectively valued outcomes. Second, deliberation has an indirect effect through preference change that leads some individuals to invest more effort into achieving those socially optimal outcomes.

The effects of deliberation on socially optimal decision-making do not occur through preference aggregation alone, lending some evidence to the notion that deliberation is more transformative than other forms of aggregating preferences. At a bare minimum, deliberation allows for the exchange of useful information that individuals do not have access to in isolation. Beyond that, individuals genuinely seem to be persuaded by the arguments of their peers and come to change their views about the experience of performing the tasks themselves.

Regarding the other pathway to improved collective outcomes, behavior change, the main analysis did not provide 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.

Furthermore, some specific features of the study context may have limited my ability to detect effects with respect to effort – meaning that this effect could be larger or less heterogeneous in a different setting. First, although I offer alternative activities (in the form of an activity sheet with a short story and puzzles), 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. Only about 2% of participants contribute nothing at all to the team effort task and the majority of these also failed to complete any tasks during the individually compensated version rounds.

Second, the study participants are, on average, 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 a measure of complementary effort – a standard voluntary contribution mechanism public good game – more than 70% of participants contribute at least half of their endowment. 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 achieving collective outcomes in this context. However, this may actually be true of many resource-constrained communities in developing countries, where social capital is high but 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 decisions.

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 concerning the achievement of shared collective goals. Recent research suggests that government policy

does have substantial potential to influence the quality of deliberation (Sanyal & Rao 2018), and so an emphasis on reason-based argumentation could yield concrete improvements in outcomes. However, 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 those 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 closer dialogue with democratic theory, as well as 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.

Acknowledgements

I 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 Chaning Jang. For input into the research design, analytical methods and theoretical focus, I acknowledge Jay McCann, Molly Scudder, Mike Touchton, Brian Wampler, and Richard Zeckhauser. For helpful comments on previous drafts of this paper, I thank participants from conferences of APSA, APPAM, IASC, the NYU-CESS Conference on Experimental Political Science, the Workshop on the Ostrom Workshop, the Earth Systems Governance Conference 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 research was supported by the Harvard Law School Project on Negotiations' Graduate Student Grants Program.

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
- Bäckstrand, Karin. (Ed.). (2010). *Environmental politics and deliberative democracy: Examining the promise of new modes of governance*. Edward Elgar Publishing.
- 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, Hélène 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
- Dryzek, J. S., & Pickering, J. (2017). Deliberation as a catalyst for reflexive environmental governance. *Ecological Economics*, 131, 353-360. doi: <http://dx.doi.org/10.1016/j.ecolecon.2016.09.011>
- 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.
- Ganuza, 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. (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. (1975). *Legitimation crisis* (Vol. 519). Beacon Press.
- Habermas, Jürgen. 1984. *Theory of Communicative Action*. Boston, MA: Beacon
- Habermas, Jürgen. (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.
- Lerner, J., & Secondo, D. (2012). By the people, for the people: Participatory budgeting from the bottom up in North America. *Journal of Public Deliberation*, 8(2), 2.
- 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.

Appendix A – Descriptive Statistics

	<u>FULL SAMPLE</u>		<u>CONTROL</u>		<u>T1:MajSec</u>		<u>T2:MajPub</u>		<u>T3:DelSec</u>		<u>T4:DelPub</u>	
	Mean	sd	Mean	sd	Mean	sd	Mean	sd	Mean	sd	Mean	sd
Female	.63		.63		.63		.67		.57		.67	
Age	33.83	11.38	33.29	11.29	33.08	11.72	34.34	11.97	35.06	11.75	34.13	10.32
Education	10.25	2.16	10.29	2.08	10.06	2.18	10.29	2.18	10.30	2.16	10.26	2.32
Ever Married	.65		.63		.67		.66		.63		.66	
No. of Children	2.24	1.91	2.25	1.94	2.19	1.79	2.28	1.87	2.30	2.06	2.16	1.88
Coll. Action	.74		.70		.80 ⁺		.76		.71		.78	
<i><u>Ability (Performance on Individual Task)</u></i>												
Letter Counting	5.01	2.85	5.37	2.94	4.84	2.76	4.68 ⁺	2.88	4.67 ⁺	2.61	5.04	2.91
Sliders	6.31	4.13	6.51	3.95	6.81	4.59	5.80	3.83	6.23	3.98	5.90	4.47
Stroop	6.85	5.06	6.82	5.35	6.67	4.88	6.59	5.06	7.09	4.66	7.16	5.00
<i><u>Preferred Task (Individual)</u></i>												
Letter Counting	.46		.46		.53		.44		.41		.46	
Sliders	.20		.21		.17		.27		.17		.19	
Stroop	.34		.33		.30		.29		.42		.36	
<i><u>Selected Task (Team)</u></i>												
Letter Counting	.43		.36		.61 ^{***}		.67 ^{***}		.22 [*]		.39	
Sliders	.15		.26		.11 ^{**}		.06 ^{***}		.11 ^{**}		.06 ^{***}	
Stroop	.42		.38		.28 ⁺		.28 ⁺		.67 ^{***}		.56 ^{**}	
<i><u>Contributions (Performance on Team Task)</u></i>												
Raw	46.72	20.72	44.87	21.88	47.70	20.39	43.78	19.80	48.74 ⁺	20.38	50.96 [*]	18.89
Standardized	.00	1.00	-.06	1.05	.10	1.06	-.08	1.00	.01	.92	.12	.88
Observations	570		210		90		90		90		90	

Significance stars indicate results of ttest comparing each treatment group to the control

⁺ $p < 0.1$, ^{*} $p < 0.05$, ^{**} $p < 0.01$, ^{***} $p < 0.001$

Appendix B – Statistical Power Calculations

This study included 570 experimental subjects, and randomization occurred at the level of the individual. Where unit of randomization is the individual, the normalized minimum detectable effect size (MDES) is given by the following formula (Glennerster & Takavarasha 2013, p.269):

$$MDES = \left(t_{(1-k)} + t_{\frac{\alpha}{2}} \right) * \frac{\sqrt{(1 - R^2)}}{\sqrt{P(1 - P)} \sqrt{N}}$$

In this equation, k is the probability of producing a false negative, 1-k is equivalent to statistical power, α is the probability of producing a false positive, N is the sample size and P is the proportion of the sample that belongs to the treatment group. Throughout my calculations, I assume standard values for statistical power equal to 0.80 and alpha of 0.05, which reduces the first term in the equation above to ~2.80.

In this experiment with 90 participants in each of 4 treatment groups and 210 individuals in the control group, the sample for any pairwise comparison between the control group and any single treatment group is 300 individuals. If the overlapping treatments can be pooled according to a single treatment characteristic (i.e. deliberation), then the sample is 390. The models include additional covariates that are known to influence pro-social behavior and performance on effort tasks, and so R^2 may increase statistical power.

This formula yields an MDES as high as 0.35σ in the case of a pairwise comparison between the control and a single treatment group with $R^2=0$ and as low as 0.20σ in the case of pooled treatment groups with $R^2=0.5$. The table below shows MDES conditional on model type and R^2 .

	<i>single treatment vs. control</i>			<i>2 pooled treatments vs. control</i>		
<i>R²</i>	0.00	0.25	0.50	0.00	0.25	0.50
<i>N</i>	300	300	300	390	390	390
<i>P</i>	0.3	0.3	0.3	0.46	0.46	0.46
<i>Power</i>	0.80	0.80	0.80	0.80	0.80	0.80
<i>alpha</i>	0.05	0.05	0.05	0.05	0.05	0.05
<i>MDES</i>	0.35σ	0.31σ	0.25σ	0.28σ	0.25σ	0.20σ

Following Cohen (1988), effect sizes below 0.3 are considered small, so this sample size allows us to identify most moderate effects, if they exist, with a probability of 80%. I may not be able to rule out the existence of small effects, but given the additional cost of deliberative decision-making processes in practice, we may conclude that only moderately sized effects would be sufficient to single-handedly justify their use – with the caveat that there may still be other, normative reasons why deliberation is worthwhile.

Appendix C – Main Analyses: Individual Performance on Team Task (linear regression)

	All Treatment Groups			Collapsed Treatments	
	(1)	(2)		(3)	(4)
Majority Rule + Secret Ballot	2.120 (2.623)	2.714 (2.440)	Majority Rule	0.835 (2.068)	1.512 (1.852)
Majority Rule + Show of Hands	-0.453 (2.363)	0.377 (1.959)			
Deliberation + Secret Ballot	3.165 (2.305)	4.207 ⁺ (2.290)	Deliberation	4.946* (2.242)	5.881** (2.017)
Deliberation + Show of Hands	6.725* (3.143)	7.557** (2.633)			
SizeSession	-1.214* (0.464)	-1.103* (0.435)	SizeSession	-1.188** (0.416)	-1.080** (0.395)
Controls?	No	Yes	Controls?	No	Yes
Constant	62.45*** (7.167)	69.10*** (8.346)		62.06*** (6.498)	68.61*** (7.935)
R^2	0.025	0.261		0.022	0.258
Observations	570	570		570	570

Standard errors in parentheses
⁺ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix D – Main Analyses: Individual Performance on Team Task (poisson models)

	All Treatment Groups			Collapsed Treatments	
	(1)	(2)		(3)	(4)
Majority Rule + Secret Ballot	0.0457 (0.0563)	0.0564 (0.0526)	Majority Rule	0.0180 (0.0456)	0.0300 (0.0410)
Majority Rule + Show of Hands	-0.0114 (0.0533)	0.00326 (0.0443)			
Deliberation + Secret Ballot	0.0673 (0.0492)	0.0917 ⁺ (0.0491)	Deliberation	0.104* (0.0473)	0.126** (0.0428)
Deliberation + Show of Hands	0.140* (0.0635)	0.159** (0.0530)			
SizeSession	-0.0240** (0.00859)	-0.0207* (0.00814)	SizeSession	-0.0234** (0.00758)	-0.0202** (0.00729)
Controls?	No	Yes	Controls?	No	Yes
Constant	4.150*** (0.134)	4.259*** (0.176)		4.141*** (0.120)	4.251*** (0.169)
Observations	570	570		570	570

Standard errors in parentheses
⁺ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix E – Decision Quality vs. Behavior Change (Effort)

	(1) Contributions	(2) Ability	(3) Effort	(4) Decision Quality
Majority Rule	1.512 (1.852)	-0.113 (0.461)	1.896 (1.258)	0.207 (0.296)
Deliberation	5.881** (2.017)	1.340** (0.473)	1.338 (1.411)	0.955** (0.300)
Ability			3.391*** (0.161)	
Controls?	Yes	Yes	Yes	Yes
Constant	68.61*** (7.935)	6.809*** (1.773)	45.52*** (5.566)	-1.715 (1.057)
Model Type	ols (appendix C, model 4)	ols	ols	probit (by team)
Observations	570	570	570	114

Standard errors in parentheses
⁺ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix F – Effect of Treatments on Intermediate Outcomes

	(1) Preference Alignment	(2) Preference Change	(3) Perceived Fairness	(4) Acceptance of Outcomes
Majority Rule	0.379** (0.130)	0.00816 (0.167)	0.0749 (0.120)	-0.219* (0.0966)
Deliberation	0.333** (0.127)	0.416** (0.139)	0.313** (0.0965)	0.297** (0.0895)
Final Preference Matches Team				0.399*** (0.0924)
Controls?	Yes	Yes	Yes	Yes
Constant	-1.858*** (0.446)	-0.288 (0.615)	4.944*** (0.308)	3.019*** (0.455)
Model Type	probit	probit	ols	ols
Observations	570	570	570	570

Standard errors in parentheses
⁺ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix G – Effect of Intermediate Outcomes on Contributions

	(1)	(2)	(3)	(4)
	Contributions	Contributions	Contributions	Contributions
Preference Alignment	6.520*** (1.651)			
Preference Change		0.0232 (2.026)		
Perceived Fairness			1.068 (0.799)	
Acceptance of Outcomes				1.645* (0.702)
Final Preference Matches Team				5.093** (1.884)
Controls?	Yes	Yes	Yes	Yes
Constant	71.42*** (8.018)	70.57*** (8.003)	65.18*** (8.700)	63.58*** (7.945)
Observations	570	570	570	570

Standard errors in parentheses
⁺ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: All the models in this table take the same form as Appendix C, Model 4, other than the addition of each intermediate outcome.

Appendix H – Effect of Intermediate Outcomes on Effort

	(1) Effort	(2) Effort	(3) Effort	(4) Effort
Ability	3.464*** (0.160)	3.436*** (0.157)	3.393*** (0.157)	3.351*** (0.156)
Preference Alignment	-1.637 (1.135)			
Preference Change		3.249* (1.436)		
Perceived Fairness			0.433 (0.708)	
Acceptance of Outcomes				0.874 (0.529)
Final Preference Matches Team				1.084 (1.353)
Controls?	Yes	Yes	Yes	Yes
Constant	45.57*** (5.354)	44.64*** (5.291)	44.10*** (6.491)	43.43*** (5.401)
Observations	570	570	570	570

Standard errors in parentheses
⁺ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: All the models in this table take the same form as Appendix E, Model 3, other than the addition of each intermediate outcome.