Can charitable appeals identify and exploit belief heterogeneity?

Drouvelis, Michalis; Marx, Benjamin M.

DOI:
10.1016/j.jebo.2022.04.019

Can Charitable Appeals Identify and Exploit Belief Heterogeneity?*

Michalis Drouvelis[†‡] and Benjamin M. Marx[§]

December 2021

Abstract

Charitable fundraisers frequently announce giving by others, and research shows that this can increase donations. However, this mechanism may not put information about peers to the most efficient use if it is costly to inform individuals who are indifferent to peer actions or causes some individuals to give less. We investigate whether a simple mechanism without incentives can predict heterogeneity in charitable responses to peer decisions. We elicit beliefs about donations in a baseline solicitation, and in subsequent solicitations we randomly assign information about others' donations. We find that elicited beliefs are often logically inconsistent and that many subjects fail to update beliefs when treated. Elicited beliefs did not predict heterogeneous treatment effects on beliefs or donations, making the strategy unlikely to succeed unless individuals are engaged and the information is salient.

Keywords: charitable, donation, norm, social preferences, peer effects, experiment. JEL: D01, D64, A13.

*We are grateful to the Birmingham-Illinois Partnership for Discovery, Engagement, and Education for research support. We thank Masaki Aoyagi, Antonio Cabrales, Gary Charness, Tongzhe Li, Mark Ottoni-Wilhelm, Michael Price, Aldo Rustichini, Kimberly Scharf, and seminar participants at the 2020 Science of Philanthropy Initiative conference for helpful comments. Sakshi Bhardwaj provided excellent research assistance. The project was approved by the UIUC Institutional Review Board under protocol #20967. Our experiment was pre-registered in the AEA RCT Registry, and its unique identifying number is AEARCTR-0006301.

[†]Department of Economics, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK. Email: m.drouvelis@bham.ac.uk.

[‡]CESifo, Munich, Germany.

[§]Department of Economics, University of Illinois, 214 David Kinley Hall, 1407 W. Gregory, Urbana, Illinois 61801, MC-707, US. Email: benmarx@illinois.edu.

1 Introduction

Charitable giving has garnered attention from researchers seeking to uncover the nature of social preferences or to address the under-provision expected due to a donation's positive externality. One influential group of these studies has built a body of evidence that individuals increase their donations when informed of high levels of giving by others (Hermalin, 1998; List and Lucking-Reiley, 2002; Vesterlund, 2003; Potters et al., 2007; Alpizar et al., 2008; Martin and Randal, 2008; Shang and Croson, 2009, Huck and Rasul, 2011; Bekkers, 2012; Karlan and List, 2012; Smith et al., 2013, Huck et al., 2015; Kessler, 2017; Klinowski, 2020; van Teunenbroek and Bekkers, 2020).[1] Such informational treatments could influence either beliefs about the giving of others or desire to give in amounts that are similar to what one believes others are giving. Either way, if donors differ in their beliefs about others' giving, then the average effect of information could mask considerable heterogeneity. Indeed, incorporating information about others' giving into a charitable solicitation could be wasteful or even counterproductive if potential donors were overly optimistic about others' giving prior to receiving the information.

Whether a fundraiser can improve on the average treatment effect of providing such information depends on whether heterogeneity can be predicted. Charities regularly engage in and can pay considerable amounts for data about donors in what is sometimes called "prospect research." Such research typically examines markers of wealth and past donation patterns to identify donors who might have the capacity and willingness to make generous donations. We consider an alternative form of donor research that gauges individuals' beliefs about others' donations. If beliefs can be elicited at minimal cost to the fundraiser and the prospect, and if donations depend on these beliefs, then information treatments can be targeted

[1]This pattern arises for both large and small stakes, as experiments involving donations of $4 or less have also found that such giving responds positively to the average donated by others (Drouvelis and Marx, 2020) and the percentage of others donating (Frey and Meier, 2004; Heldt, 2005).

towards those who would have positive treatment effects.

Our study tests whether collecting a small amount of information about beliefs in one solicitation can allow a fundraiser to predict and take advantage of donor belief heterogeneity. To do so, we employed a short survey following one solicitation and then randomly assigned information in later solicitations. Each solicitation invited subjects to donate some of their earnings from completing a survey for another study. In the baseline solicitation, we elicited subjects' beliefs about amounts donated by others and about the amount they themselves would donate if they knew that the average donation of others was \$0.50 or \$1. We elicited these beliefs without providing incentives for accuracy, thereby testing whether fundraisers can reasonably collect predictive information in their campaigns. We elicited beliefs about three statistics of baseline donations: the mean, the share of subjects making any donation, and the share donating \$1 or more. We found dispersion in beliefs about each statistic, with a number of subjects overestimating and underestimating each.

We then conducted experiments to test for heterogeneity in treatment effects of information about others' baseline donations. As we described in our pre-registration, the quantity of interest was not the average treatment effect but whether the information collected in the baseline solicitation predicts treatment effect heterogeneity. Our three treatments each informed subjects about one of the statistics of baseline donations, any which could have the strongest effect on donations depending on such factors as the extent to which each statistic affected beliefs, the number of subjects wishing to conform with the statistic, and the number of subjects who would consider a donation amount consistent with the statistic. We elicited beliefs again after the experimental solicitation to check whether the treatments shifted beliefs from their baseline value towards the value shown in the treatment. After assessing this, we test whether donations change in the same direction as beliefs about others' behavior.

Results of our first experiment were limited. Any treatment effects did not differ significantly between those with high versus low beliefs at baseline and did not depend on whether the subject reported beliefs about their own donations that indicated they would give more. Estimates were not significant for changes in beliefs and therefore did not significantly affect donations.

In a second experiment, we made the share-donating treatment more salient. Speculating that its effects on beliefs had been minimal because subjects missed the information, we amended the treatment by placing the information near the beginning of the solicitation in bold text. This version of the treatment also did not significantly shift beliefs towards the value provided.

We conclude that it would be possible but likely difficult for charities to identify and exploit belief heterogeneity. More than half of our subjects report beliefs that are logically inconsistent and therefore not reliable in predicting how they will respond to treatment. Among those treated with information, two thirds or more appear to not see the information, even after efforts to increase its salience. Such lack of attention by potential donors would limit a fundraiser's ability to use belief heterogeneity to increase donations. However, the sizable donation responses to changes in beliefs about others' donations suggest that fundraisers could profitably exploit belief heterogeneity among donors who have a stronger connection with the charity and hence might be more attentive and accurate in the belief elicitation stage.

In examining whether individuals attempt to follow the donation patterns of others, our study was similar to the literature on "conditional cooperation" in public goods games. In a review of research with public goods games, Chaudhuri (2011) discusses a bevy of studies finding that half or more of their subjects are conditional cooperators whose contributions are increasing in the contributions of others, typically by slightly less than dollar-for-dollar (e.g. Keser and Van Winden, 2000; Fischbacher et al., 2001; Kurzban and Houser, 2005; Gächter, 2007; Kocher et al., 2008; Fischbacher and Gachter, 2010; Kessler et al., 2020; Croson, 2007; Kumru and Vesterlund, 2010). In the public goods game, conditional cooperation may be due to reciprocity, but Bardsley and Sausgruber (2005) also find an effect of informing subjects about the amounts given by a group other than their own, thereby identifying conformity that may also occur in giving to charity.

There are several reasons that information about others' donations might affect one's own giving decision. The donations of others might signal charity quality (Glazer and Konrad, 1996; Vesterlund, 2003; Andreoni, 2006; Potters et al., 2007; Bracha et al., 2011), but we have minimized this factor in our experiment by allowing donors to choose from a list of charities and not informing them about which charities others chose. Individuals may donate amounts similar to others to improve their image or social status (Andreoni and Bernheim, 2009; Ariely et al., 2009; Kumru and Vesterlund, 2010) or because they feel social pressure (DellaVigna et al., 2012; Andreoni et al., 2017), particularly when peers are soliciting the donations (Meer, 2011; Meer and Rosen, 2011; Castillo et al., 2014, 2017). In our experiment, we give subjects no reason to think that anyone other than the researchers will observe their donation.

van Teunenbroek et al. (2020) conduct a systematic review of 35 studies and propose factors that may determine whether social information affects donations. The factors that they propose included perceptions of social norms. Social norms are rules for appropriate behavior (Elster, 1989; Ostrom, 2000; Krupka and Croson, 2016). They should be commonly recognized, but incentive-compatible elicitation has shown heterogeneity in beliefs about what others consider appropriate (Krupka and Weber, 2013; Drouvelis et al., 2019). We considered adherence to an uncertain social norm to be the most likely reason for a treatment effect in our setting.

Finally, we note that our study reflects recent interest in targeting treatments among heterogeneous donors. Cagala et al. (2015) study how to target a gift-exchange treatment, while Adena and Huck (2019) propose matching gifts with thresholds that are personalized to each potential donor. Both of these papers employ heterogeneity in past donations, whereas our paper studies heterogeneity according to elicited beliefs about one's own and others' donations.

The paper proceeds as follows. Section 2 describes the design and results of the baseline solicitation and elicitation. Section 3 presents the first experiment, and Section 4 presents the second. Section 5 concludes.

2 Baseline Solicitation

The baseline solicitation and survey captured variation in donations and reported beliefs. This section describes how we elicited these and provides descriptive analysis that motivated the choices in the experimental design.

2.1 Design and Procedures

We conducted our study using Prolific (www.prolific.co), a web-based platform for recruiting participants, and oTree (Chen et al., 2016). We carried out the baseline survey in mid-August 2020. We restricted our sample to be of US nationality, and 80 percent of our sample is nationally representative in terms of age, sex, and race.[2]

Subjects were recruited to complete a survey on judicial politics (Garcia, 2020). Prolific users who met our selection criteria saw a link for this survey and the description "This study on court-case decisions pays $4 (£3.05) and will take approx-

[2]With some uncertainty about the number of subjects who would provide complete responses, we initially collected a nationally representative sample of 309 subjects before discovering that the Prolific feature for adding subjects does not allow nationally representative samples to be increased by less than 300 subjects at a time.

imately 20 minutes." In this way, subjects earned income through a task unrelated to donating.[3] The baseline solicitation and those in each experiment followed variants of the same survey task. The consent screen explained that "If you complete the survey, you will be paid $4, which is the sum of the $1 shown on Prolific and $3 to be paid as a bonus." Base payments on Prolific cannot be donated in a way that researchers can observe, but by paying $3 of the $4 as a bonus, we were able to give subjects the option of donating any portion of the $3.

The solicitation occurred after subjects had completed the survey, and it was unannounced as in checkout campaigns and other charitable solicitations. The text read "Thanks, you're almost done! When you finish, you will be paid a $3 bonus. Would you like to donate some of your bonus to one of these charities fighting the COVID-19 pandemic? These are some of Charity Navigator's highest-rated charities." Subjects were then able to select from the following list of charities that we collected from a web page in which the charity evaluator Charity Navigator listed highly rated charities fighting the effects of the coronavirus pandemic.[4] We provided subjects with a list of charities partly in the belief that this would increase donation rates and thus statistical power. The choice to provide a list of charities and treat subjects with information about giving to the list rather than to an individual charity was also meant to isolate the normative content of the information about others' giving, essentially taking the opposite of the approach of Ottoni-Wilhelm et al. (2017), who created individualized charitable causes for each subject in order to isolate crowd-out to identify altruism. Below this list were two lines of text: "How

[3]Previous research has demonstrated that donation responses may vary depending on whether income is earned or unearned (Cherry et al., 2002; Cherry et al., 2005; Harrison, 2007; Kroll et al., 2007; Erkal et al., 2011; Reinstein and Riener, 2012; Carlsson et al., 2013; Tonin and Vlassopoulos, 2017; Drouvelis et al., 2019; Drouvelis and Marx, 2020).

[4]The charities (and total donated to each across all sessions) were Save the Children ($184.36), CDC Foundation ($83.45), National Foundation for Infectious Diseases ($57.50), Matthew 25: Ministries ($52.00), United Way Worldwide ($45.11), and Give Directly Inc. ($42.50). Because the choice of charity was not a topic of inquiry, we chose not to randomize the order so that the screen was identical for subjects.

much of the $3 bonus would you like to donate? _____ You will receive a receipt for this tax-deductible donation.”

The solicitation was followed by the belief-elicitation screen. Text at the top stated “Please answer these final questions to help us learn a bit more about donations in Prolific surveys.” The questions were as follows

- How much of the $3 bonus do you think other participants donate on average? $_.___
- How much of the $3 bonus do you think you would donate if
  - You knew that others gave $0.50 on average? $_.___
  - You knew that others gave $1.00 on average? $_.___
- Out of 100 people, how many do you think donate? ___
- Out of 100 people, how many do you think give $1 or more? ___

This screen of questions always followed the solicitation so that all subjects knew what solicitation the questions were regarding, and they were appeared as listed here in the order we thought most intuitive. A key component of our design is that we purposefully did not incentivize responses to these questions. Our choice reflected a desire to mimic a realistic tool for a fundraiser to use, recognizing that it would be unusual to incentivize accurate prediction within a fundraising interaction in the field. Responses to these questions in our study should therefore contain inaccuracies or biases that are similar to those in data that a charitable organization would collect.

In total, we collected data from 382 individuals. Of these, 51 percent identified as female and 65 percent as white. One benefit of using Prolific was the dispersion of subject ages: 40 percent were ages 18 to 34, 33 percent ages 35 to 54, 26 percent ages 55 to 74, and 2 percent 75 or over.

2.2 Results of Baseline Solicitation

Figure 1 shows the distribution of donations made. The maximum amount that subjects could donate was \$3, which was the decision made by 9.9 percent of subjects. The modal donation was \$0, while the second most popular choice was \$1. On average, donations were equal to \$0.64 (s.d. = 0.96), which is line with results from similar donation experiments. Subjects often chose round numbers or increments of \$0.50, though there are some who chose smaller fractions.

Figure 1: Distribution of baseline donations

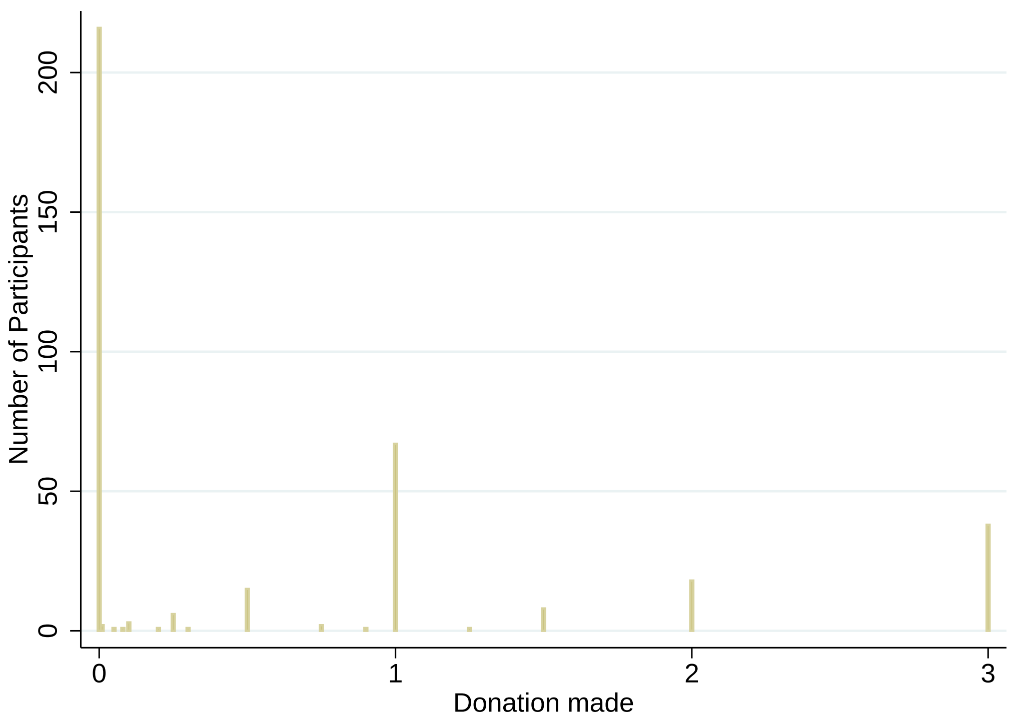

Notes: N = 382. Bin width = 0.1. Mean = \$0.64.

Figure 2 displays the distribution of beliefs about the mean of others' donations. We found that subjects expect others to donate a bit more than they actually do (average = \$0.89, s.d. = 0.75). Subjects' own donations were positively correlated with their beliefs about the average of others' donations ($\rho = 0.34$; $p < 0.001$). Causality here could run in either direction, with subjects wanting to give an amount similar to what they think others give, or with subjects stating beliefs that justify their donation choice. Our experiments attempted to examine causality by randomly assigning information about the giving of others.

Figure 2: Baseline beliefs about mean donation

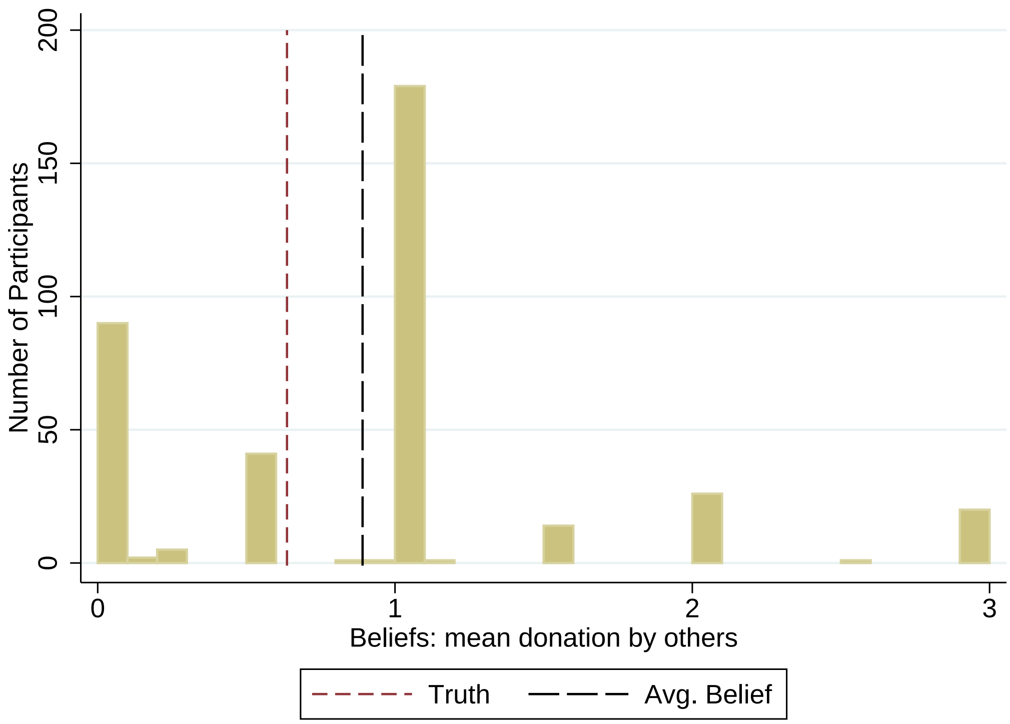

Notes: Bin width = 0.1.

It is worth noting the discrete nature of the responses in Figure 2. The most common belief was that other subjects give exactly $1 on average, and other multiples of $0.50 were common choices. "Heaping" at round numbers has been observed in many contexts (Barreca et al., 2011), and in our context there are few round numbers to choose from. Here, heaping could be due to low precision in subjects' beliefs.

Figure 3 displays the distribution of beliefs about the share of others donating particular amounts. We constructed these shares by dividing subjects' responses to the "Out of 100 people" questions by 100. The first panel of the figure shows subjects' beliefs about the number of individuals making a positive donation. Here the average belief (0.427) was very close to the truth (0.434) that we observed in the data ($p = 0.58$ for a test of equality). The most common response was a belief that 50% of others donate a positive amount. The second panel of Figure 3 shows beliefs about the share donating $1 or more. Here the average belief (0.27) somewhat underestimated ($p < 0.01$) the true share of subjects who give $1 or more (0.35).

Figure 3: Beliefs about share of others donating threshold amounts

Donate $>$ $0 Donate $\geq$ $1

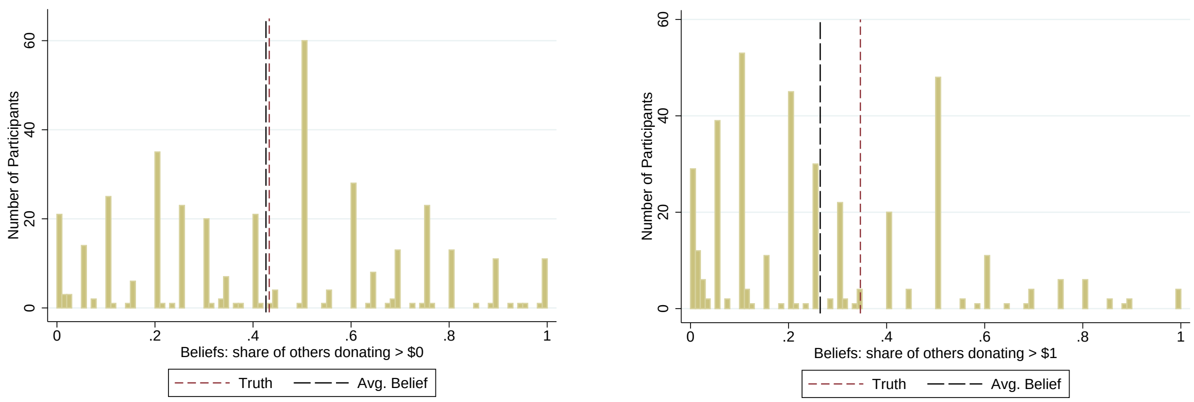

Notes: Bin width = 0.01.

Comparing the subjects' beliefs about each of the donation statistics, we found that the majority of subjects report logically inconsistent beliefs. Let b_{i0}^{avg}, b_{i0}^{0}, and b_{i0}^{1} denote, respectively, the reported baseline beliefs of subject i about the average donation, share donating, and share donating $1 or more. It is logically inconsistent for a subject to hold beliefs with any of the following properties: $b_{i0}^{1} > b_{i0}^{0}$, $b_{i0}^{avg} < 0.01b_{i0}^{0}$, $b_{i0}^{avg} > 3b_{i0}^{0}$, $b_{i0}^{avg} < 1b_{i0}^{1}$, or $b_{i0}^{avg} > 3b_{i0}^{1} + 0.99\,(b_{i0}^{0} - b_{i0}^{1})$. We found that 238 subjects (62 percent) reported beliefs that were not logically consistent. This fact provided suggestive initial evidence that unincentivized beliefs are imprecise and would be difficult to exploit to increase donations. We could have treated this as an attention check and dropped subjects reporting beliefs that are not consistent with each other, but it could be that any one of these subjects' beliefs is predictive, and the three treatments in our design each test for heterogeneity related to a particular belief. We therefore invited subjects with logically inconsistent beliefs to participate in the experiments, but we estimate each regression with and without these subjects included.

The other baseline questions asked subjects how much they would donate conditional on particular values of donations by others. We found that only 49 subjects (13 percent) report that they would donate more if others donate $1 than they would if others donate $0.50. These self-proclaimed conditional donors (similar to "conditional cooperators" in the public goods literature) reported distributions of beliefs that were similar to those who reported that their donations are not conditional on those of others: across the three belief measures, the largest difference in means between the two groups was 0.019, and the largest difference in absolute error was 0.028. The similarity in the two groups' belief distributions offered the opportunity to test whether self-proclaimed conditional donors do in fact respond more to information about the giving of others.

Baseline donations and beliefs about the average donation of others are highly correlated; a regression of the former on the latter gives a coefficient of 0.45 with standard error 0.06. This relationship need not be causal, however, as there may be unobserved characteristics that correlate with both, or it may be the case that one's donation has a reverse-causality effect on one's reported belief. Our experiment tests for a causal effect by randomizing information about the donations of others.

3 Experiment 1

3.1 Design of Experiment 1

Our first experiment took place in late August and early September. We invited all participants from the baseline solicitation. Subjects earned their endowment by completing a different survey on judicial politics for Garcia (2020) and were again asked to decide how much to donate from the \$3 fixed bonus they received. The list of selected charities was identical to the ones available at baseline, and we asked the same questions about beliefs following the donation decision in the experiment, allowing us to estimate treatment effects on beliefs.

Our experiment was designed to answer the following research questions:

R1. Do reported hypothetical conditional donations predict the effect of information about the average donation of others?

R2. Do reported beliefs about the donations of others predict the effect of information about:

R2a) the average donation?

R2b) the share of individuals donating?

R2c) the share of individuals donating \$1 or more?

Our design randomized information between subjects. In the control group (denoted by *CG*) subjects made their donation decisions without being provided any information about giving behavior in the baseline solicitation. In our treatments, we accompany the solicitation of subject i with information about past donations from the first wave. The "Average Donation" treatment (T^{avg}) stated that "Participants in our last Prolific survey donated an average of \$.64." The "Share Donating" treatment (T^{0}) stated that "43% of participants in our last Prolific survey made a donation." The "Share Donating \$1+" treatment (T^{1}) stated that "35% of participants in our last Prolific survey donated \$1 or more." These were the true statistics, and hence we do not employ deception. We provided this information in the middle of the solicitation, as shown in Appendix Figure A.1.

We distributed subjects between treatment arms so as to address question **R1** and question **R2**. Research question **R1** relates to heterogeneity in reported hypothetical conditional donations. We used the elicited hypotheticals to define $\triangle\, hypo_{i0}$

as the predicted treatment effect of learning the average of others' donations. We calculated this using a linear combination of the hypothetical donations conditional on others' donations of $0.50 and $1 to approximate the response to the information that the average was $0.64.[5] Only 87 subjects reported that their donation would change, i.e. that $\triangle\, hypo_{i0} \neq 0$. We therefore split all of these subjects between *CG* and T^{avg}, with the remaining subjects distributed across the four treatment arms.

To promote balance across treatment arms and enhance precision for our heterogeneity analyses, we utilized stratification in our randomization. We stratified based on the sign of four variables: $\triangle\, hypo_{i0}$, b_{i0}^{avg}, b_{i0}^{0}, and b_{i0}^{1}. Because only 13 subjects have $\triangle\, hypo_{i0} < 0$, we did not further subdivide them using the other variables. Because subjects with $\triangle\, hypo_{i0} > 0$ are assigned to either *CG* or T^{avg}, we did not subdivide them based on b_{i0}^{0}, or b_{i0}^{1}. We created a separate category for subjects with $b_{i0}^{avg} = 0$. Then, to the extent that we could do so without reducing the count within strata to single digits, we further subdivided strata into "high" and "low" groups based on whether their b_{i0}^{0} (which was the most accurate belief on average) was above or below the median for their stratum. We assigned subjects in each stratum as evenly as possible across treatment arms and then as needed assigned any additional subject to arms in the following order: Control, then "Average Donation," and then "Share Donating $1+."

Appendix Tables A.1 and A.2 describe the randomization. Appendix Table A.1 details the stratification and displays the number of subjects in each stratum. The only single-digit strata contain exactly four subjects, allowing one to be assigned to each treatment arm. Appendix Table A.2 shows tests of balance across treatment arms in baseline donations, beliefs, and demographics. None of the coefficients in either panel are statistically significant, nor are any of the *F* tests of joint signif-

[5]If we denote the baseline donation by d_{i0} and by $hypo_{i0}^{1}$ and $hypo_{i0}^{0.5}$ the reported donations in the hypothetical cases in which others donate an average of $1 or $0.50, respectively, then $\triangle\, hypo_{i0} = 2(.64 - .5)\left(hypo_{i0}^{1} - hypo_{i0}^{0.5}\right) + hypo_{i0}^{0.5} - d_{i0}^{0}$.

icance. All subjects from the baseline solicitation were invited to the experiment, and Appendix Table A.3 shows the balance across treatment arms and strata of those who chose to participate.

3.2 Estimation

Our goal was to estimate heterogeneous effects across subject types. This necessitates a more complicated estimating equation than what one would use to estimate the average treatment effect. We therefore pre-registered our design in the AEA Registry prior to the experiment. We estimate separate equations for the different types of elicited variation that might predict treatment effects.

To answer question **R1**, we restrict the sample to subjects assigned to either the control group (C) or the "Average Donation" treatment (T1). We estimate the following regression:

$$Y_i = \alpha + \beta_1 T_i^{avg} + \beta_2 pos\, hypo_{i0} + \beta_3 T_i^{avg} * pos\, hypo_{i0} + \epsilon_i \tag{1}$$

We consider two types of outcomes Y_i: the change in individual i's belief between the baseline solicitation and the experiment, and the change in the individual's donation. In equation (1), we estimate these outcomes as a function of the treatment dummy variable T_i^{avg} and an indicator $pos\, hypo_{i0} = 1\left\{\triangle\, hypo_{i0} > 0\right\}$ for subjects for whom we should see a positive treatment effect according to their hypothetical donations conditioning on the amount given by others. In this regression, the constant (α) gives the average value of the outcome among those who are in *CG* and *do not* report hypothetical donations greater than their baseline donation. Coefficient β_2 gives the difference between the average values of these subjects and those in *CG* who *do* report hypothetical donations greater than their baseline donation. Coefficient β_1 gives the effect of treatment among those who *do not* report hypothetical donations greater than their baseline donation, while β_3 gives the relative effect of treatment (in addition to β_1) among those who do. We test whether $\beta_3 = 0$ to assess whether subjects' beliefs about their own donations are predictive of treatment effects. We report heteroskedasticity-robust standard errors for all regressions.

To answer the three parts of question **R2**, we estimate three regressions. In each, we restrict the sample to the control group and a single. We then estimate the following regression:

$$Y_i = \alpha + \beta_1 T_i^X + \beta_2 b_{i0}^X + \beta_3 T_i^X * b_{i0}^X + \epsilon_i \tag{2}$$

In equation (2), we estimate outcome Y_i as a function of the treatment dummy variable T_i^X, an indicator for receiving treatment $X \in \{avg, 0, 1\}$, and the corresponding belief variable b_{i0}^X. The control variable b_{i0}^X captures any differential trends, such as if subjects who were optimistic on the day of the baseline solicitation systematically give less in the experiment, regardless of treatment. The coefficient of interest is β_3, which captures relative treatment effects. If elicited beliefs are predictive, then we would expect information about the true value of the statistic to have a positive (negative) effect on those whose baseline beliefs were below (above) the truth, i.e. we would expect $\beta_3 < 0$.

3.3 Results of Experiment 1

Figure 4 shows the density of beliefs, elicited after the experiment, about the average donation. Among subjects who were treated with the information that the previous average donation was \$0.64, 34 percent reported beliefs strictly between 0.6 and 0.7. No subjects in the control group report such values. The figure also shows that the treatment reduces the likelihood of other values throughout the distribution that subjects would have entered if not for treatment.

Figure 4: Beliefs after Experiment 1: average donation

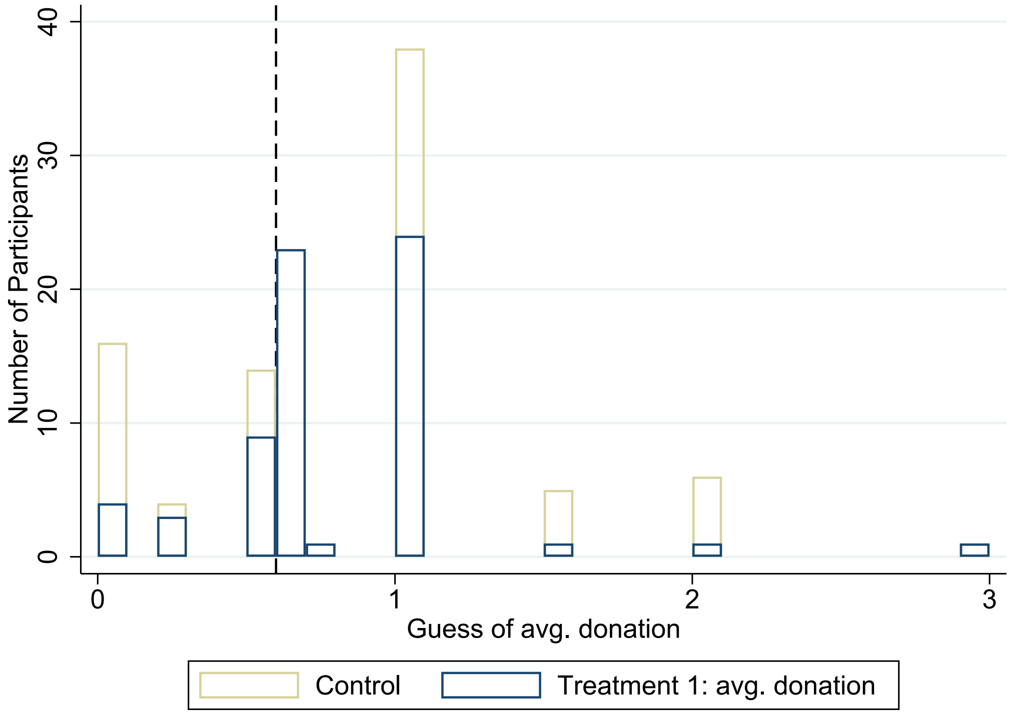

Notes: N = 151. Bin width = 0.1. Dashed vertical line indicates mean baseline donation, which was reported to the treatment group.

Figure 5 shows the essence of the estimation strategy. Panel A of the figure plots the change from baseline beliefs to post-experiment beliefs as a function of baseline beliefs and treatment. The x axis is binned to group subjects by ranges of beliefs about baseline donations. It was expected that the change in beliefs among those with a low value would be more positive for the treatment group than the control group, while the opposite would be true at high baseline levels, and this is what we observe in the figure. In Panel B of Figure 5, the outcome variable is the difference between a subject's donation before and after the experiment. For subjects whose baseline guess was that the average donation was less than $1.50, treatment shifts donations in the same direction in which it shifted beliefs.

Figure 5: Effects of average-donation treatment in Experiment 1

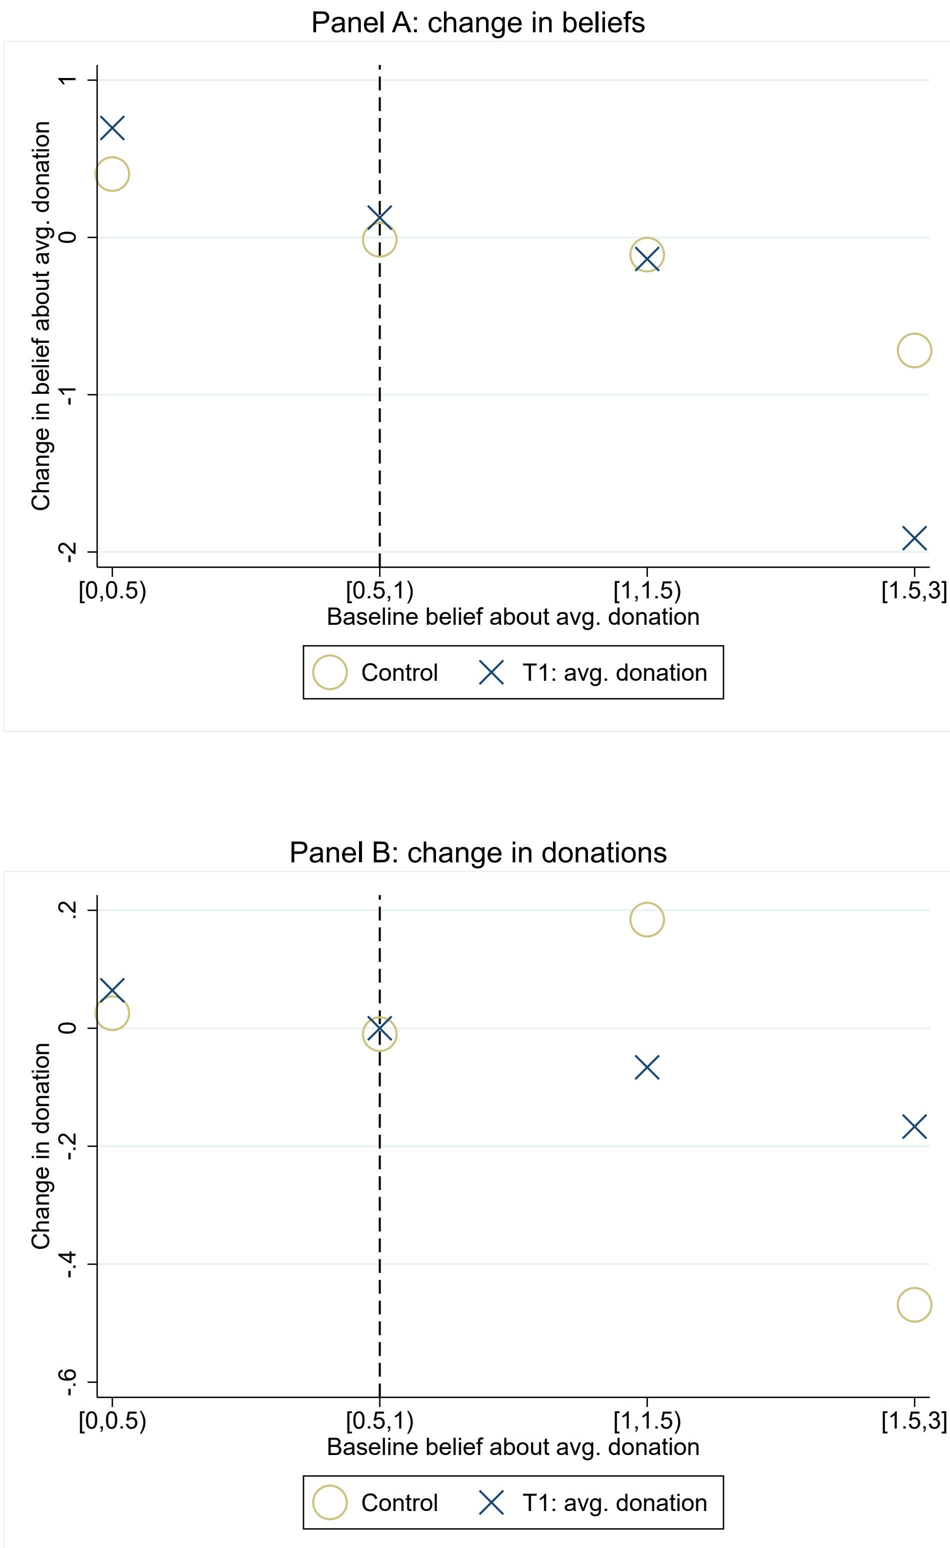

Notes: N = 151. Dashed vertical line indicates baseline average donation.

Two results in Figure 5 reveal challenges in collecting meaningful beliefs that will predict treatment effects on donations. First, while T^{avg} reduced beliefs among subjects who reported a belief greater than \$1.50 at baseline, it did not reduce these subjects' donations. We note that 18 of the 22 subjects with a baseline belief of \$1.50 or more had logically inconsistent beliefs, and we suspect that these subjects entered their beliefs about the mean conditional on donating rather than the unconditional mean, which would render any reduction in the reported belief meaningless. Second, there is a downward sloping pattern among *CG* subjects, with those reporting low (high) baseline beliefs regressing upwards (downwards) towards the mean. This suggests that many subjects did not firmly believe the numbers that they reported at baseline. We can quantify the mean reversion by restricting the sample to *CG* and regressing the change in the belief on the baseline belief, and this gives an estimate of -0.56 with standard error 0.08. However, if we further restrict to the group that report logically consistent beliefs, then the estimate shrinks to -0.17 with standard error 0.14, evidence that more of this group reported actual beliefs. In our regression analysis, we examine how results change when we drop subjects whose elicited beliefs are logically inconsistent.[6]

Appendix Figure A.2 shows that the share-donating treatments had less effect on the distributions of relevant beliefs. In both the treatment and control groups there is wide dispersion, with beliefs spanning the full range from 0 to 1. Treated subjects are more likely to report values near the value stated in the treatment but by a small amount. For T_i^0, the treatment group is only 4.3 percentage points more likely to appear in the bin containing the treatment value of 0.43 ($p = 0.085$). Thus, it appears that fewer than 5 percent of subjects absorbed the information provided through the treatment, and the effect is even smaller for T_i^1.

[6] It is possible that even our simple one-screen survey eliciting a handful of beliefs imposed an excessive cognitive load, but in that case we would expect to obtain the most reliable results from the first statistic elicited, and as we show in the next section, a more salient treatment was effective in influencing responses about the second-to-last statistic and even influenced donations.

Regression results appear in Table 1. In each panel of the table, we display the coefficient of interest from estimates of equation 1 (columns 1-2) or equation 2 (columns 3-8). We estimate each regression with the full experimental sample and then dropping subjects who reported beliefs that were not logically consistent with each other.

Panel A shows estimated effects of each treatment on the relevant belief. We first test (columns 1 and 2) whether treatment effects of T^{avg} on beliefs depend on whether a subjects' reported conditional donations imply a treatment effect on donations ($pos\,hypo_{i0}$ in equation 1). We did not expect effects on beliefs to differ by this variable, and neither estimate is significant at the 0.05 level. The remaining columns of Panel A of Table 1 test for heterogeneous effects by baseline beliefs about others' donations. Columns (3) through (8) show that none of the treatments significantly shifted beliefs towards the truth in either the full sample or subsample. The strongest result is that T^1 had the expected relatively negative effect (p=0.057) on beliefs that were higher at baseline.

Table 1: Results of Experiment 1

Panel A: beliefs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in belief about:	Avg. donation		Avg. donation		Share >0		Share 1+	
T^{avg} * 1{hypothetical>0}	-0.11	-0.44*						
	(0.32)	(0.23)						
T^{avg} * Baseline belief: avg.			-0.28	0.37				
			(0.17)	(0.27)				
T^{0} * Baseline belief: >0					0.04	-0.13		
					(0.19)	(0.33)		
T^{1} * Baseline belief: 1+							-0.21	-0.36*
							(0.16)	(0.18)
Observations	155	56	155	56	102	41	100	45
R-squared	0.00	0.05	0.53	0.20	0.25	0.27	0.26	0.29
Drop logically inconsistent		X		X		X		X

Panel B: donations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T^{avg} * 1{hypothetical>0}	0.32	0.06						
	(0.30)	(0.47)						
T^{avg} * Baseline belief: avg.			0.06	-0.45				
			(0.14)	(0.48)				
T^{0} * Baseline belief: >0					0.26	0.32		
					(0.50)	(0.55)		
T^{1} * Baseline belief: 1+							0.27	-0.35
							(0.80)	(0.81)
Observations	155	56	155	56	102	41	100	45
R-squared	0.04	0.06	0.01	0.03	0.09	0.05	0.05	0.01
Drop logically inconsistent		X		X		X		X

Notes: All regressions include the uninteracted heterogeneity variable and treatment variable in addition to their displayed interaction. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Panel B of Table 1 tests for predicted effects on donations. Given the lack of effect on beliefs, it is not surprising that no coefficients for donations are statistically significant. To assess whether they are economically significant, we consider the estimates for T^1 in the subsample reporting logically consistent baseline beliefs, the treatment and group for which we found the strongest evidence of an effect on beliefs. Column (8) implies that changing the perceived percentage of others donating at least $1 from 0% to 1% would increase a subject's own donation by $0.97 (=0.35/0.36). Although the point estimate has the expected sign, the estimate is imprecise due to the weakness of the first-stage effect on beliefs, with a confidence interval wider than the choice set. Thus, our treatments in Experiment 1 did not have a strong enough effect on beliefs for us to draw any conclusions about the effect of beliefs on donations.

Our regressions are based on the expectation of downward responses by those with high beliefs at baseline and upward responses by those with low beliefs at baseline. This is in contrast with several papers in the literature, which provide information about a large gift of value that would exceed most subjects' expected belief. Moreover, it is possible that responses are asymmetric, for example if motivated reasoning results in smaller responses to information about an amount greater than the baseline belief than to information about an amount less than the baseline belief. For these reasons we complement our main regressions with those in Appendix Table A.4 in which we simply estimate the average treatment effect when restricting attention to two subsamples, one with baseline beliefs that fall below the values communicated by the treatment, and one with baseline beliefs that fall above. As in Table 1, estimated effects on beliefs mostly have the expected sign but are not statistically significant, suggesting that our treatment did not have an asymmetric effect but rather was simply missed or ignored.

In summary, our treatments mostly failed to induce a first-stage effect on beliefs and thus could not affect donations. It was therefore important that we measured changes in beliefs because if we had assumed an effect on beliefs and only measured changes in donations, then we might have drawn the unwarranted conclusion that changing beliefs about the giving of others does not affect giving. Instead, we conclude from Experiment 1 that informational treatments are unlikely to work when they are not salient and when potential givers are not attentive.

4 Experiment 2

Our first experiment failed to reveal predictable heterogeneity in treatment effects because it mostly did not shift beliefs as expected. We find evidence that this is partly because many subjects reported beliefs at baseline that they did not hold firmly enough to report again when not treated in the experiment. Mean reversion is less of an issue among subjects who reported logically consistent beliefs at baseline, but even these subjects' beliefs do not respond strongly to treatments. This could be because they did not see the information or because they did not deem it relevant. Subjects could deem the information irrelevant in the sense that they believe the behavior of the past sample does not predict that of their contemporaries or in the sense that they do not care how much others give. We therefore ran a second experiment in which we focused on a single treatment and made it more salient. If doing so affected beliefs, then we would conclude that subjects had not seen the information in the first experiment. A further effect on donations would imply that subjects deemed the information relevant.

4.1 Design of Experiment 2

For maximum power, Experiment 2 employed a single treatment and control. We selected one of the treatments that had not significantly altered reported beliefs, namely the sentence "43% of participants in our last Prolific survey made a donation." Rather than appearing in between the charity options and the field for donation amount, in this experiment this sentence appeared on the line after the first line that said "Thanks, you're almost done! When you finish, you will be paid a $3 bonus." The sentence was in bold font and followed by the same content as in Experiment 1, starting with “Would you like to donate some of your bonus...” Appendix Figure A.1 provides a screen shot of the share-donating treatment in each experiment. We updated the AEA experiment registry before launching the second experiment in late December 2020.

Subjects were selected from the sample that had completed the baseline survey. Sample inclusion and stratification were determined by the baseline belief about the share of others donating. We excluded subjects whose reported belief was equal to zero and those who had entered a share from 0.4 to 0.59. The latter exclusion removed subjects whose baseline beliefs were already close to the truth. Subjects were then grouped by baseline-belief bins of $(0, 0.2)$, $[0.2, 0.4)$, $[0.6, 0.8)$, and $[0.8, 1]$. The potential subjects in each bin were randomly divided between treatment and control groups. Both versions of the survey were opened on December 21 and closed on December 31, at which time 54 members of the control group and 62 members of the treatment group had completed Experiment 2.[7]

Other procedures were similar to Experiment 1. The solicitation again followed a new political science survey. The list of selected charities was identical to the list in the baseline solicitation and Experiment 1. Among the questions that followed the solicitation, we removed those that asked how much the subject would donate based on amounts of donations by others, and we retained the questions about statistics of others' donations.

[7] To obtain a larger sample, we invited subjects whether or not they had participated in Experiment 1. Experiment 1 participants made up 73 percent of the control group and 83 percent of the treatment group, and a test of equality of these percentages gives a p-value of 0.19. Only two of these subjects had reported in Experiment 1 that 43 percent donate, and one of these subjects was assigned to the control group in Experiment 2, which would induce a bias against finding predictable treatment effects if this subject remembered that the share donating had been 43 percent, but the subject's estimate in Experiment 2 was 25 percent.

4.2 Results of Experiment 2

Figure 6 shows that the more salient information treatment increased the odds of subjects reporting a belief close to the value of 43 percent provided in the treatment. The treatment increased the probability of a subject reporting a belief $b_{i2} \in [0.4, 0.5)$ by 24 percentage points ($p < 0.01$). Even so, 69 percent of subjects continue to report beliefs outside of this bin, and there is no increase in the adjacent bins. Thus, it appears that subjects who respond to the information continue to be in the minority.

Figure 6: Beliefs after Experiment 2

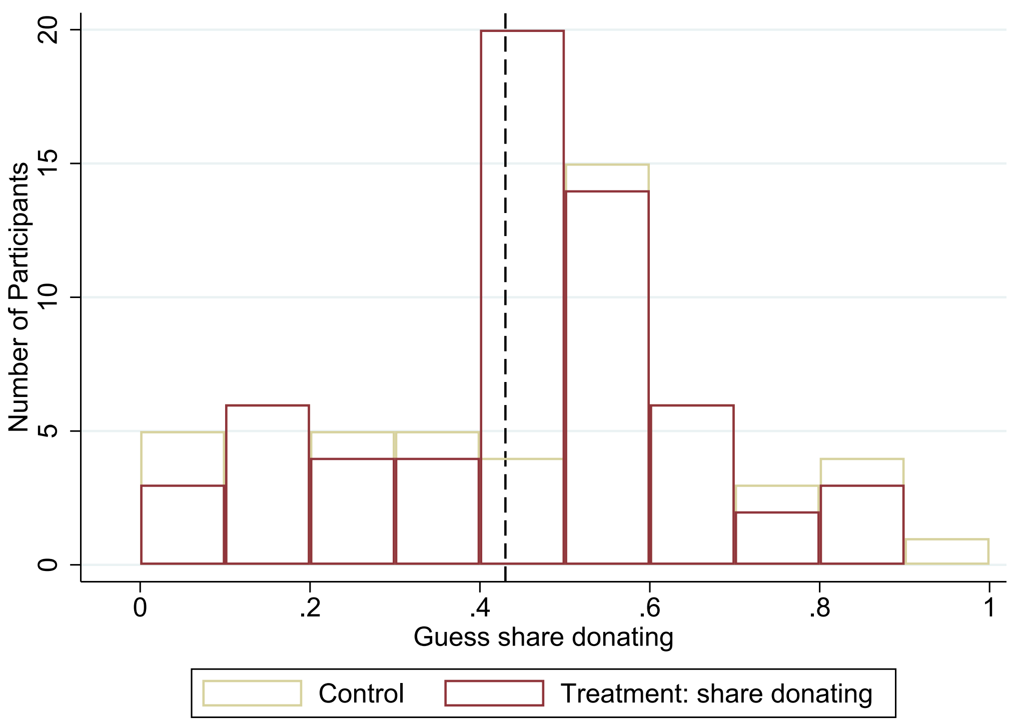

Notes: N = 115. Bin width = 0.1. Dashed vertical line indicates baseline share donating.

Figure 7 shows how beliefs and donations changed as a function of treatment and the baseline belief. In Panel A, we would have expected treatment to increase the perception of the share donating among those whose baseline belief was low and decrease it among those with a high baseline belief. This is true for most but not all bins of baseline beliefs, consistent with the observation in Figure 6 that the salient treatment had an effect that was stronger but still far from ubiquitous. It is not surprising, then, that in Panel B we do not see a treatment effect on donations that is systematically positive among those with low baseline beliefs and negative among those with high baseline beliefs.

Figure 7: Effects of treatment in experiment 2 by baseline belief

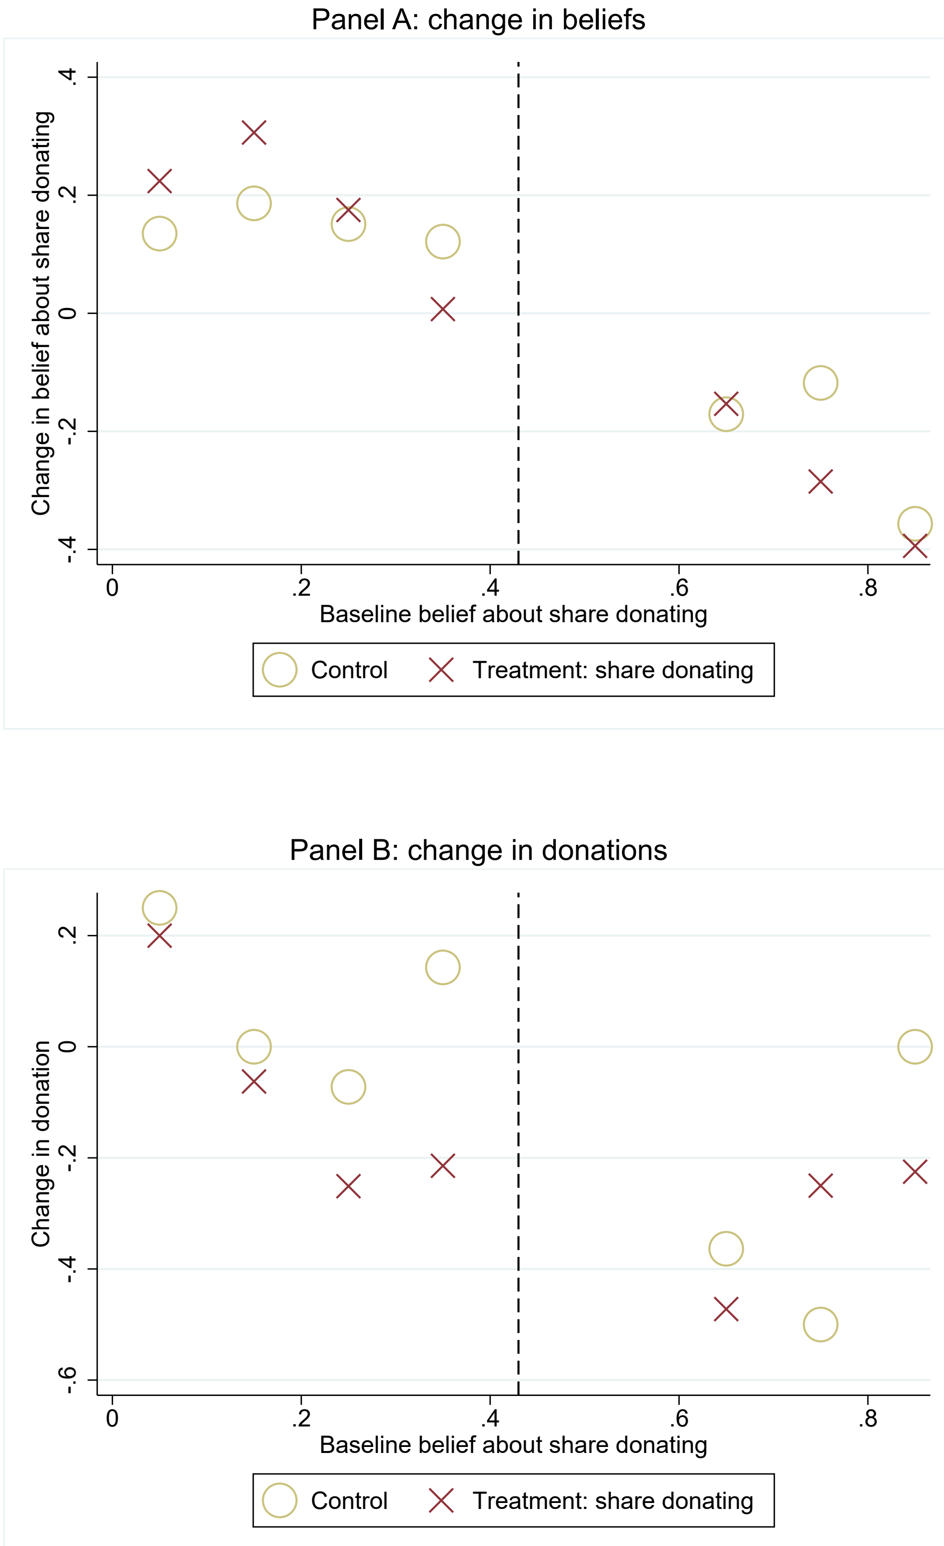

Notes: N = 116. Dashed vertical line indicates baseline share donating.

Regression results appear in Table 2. The outcome in columns (1) and (2) is the change in a subject's belief about the share of others donating (relative to that subject's belief in the baseline survey). The coefficient of interest has the expected negative sign, with treatment causing subjects with higher (lower) beliefs at baseline to revise these down (up) towards the value of 43 percent that was provided in the treatment. Indeed, both estimates are more negative than those in the first experiment. This suggests that increasing the salience of the treatment improved its effectiveness. However, the estimates are still not statistically significant. It is not surprising, then, that the estimated effects on donations in columns (3) and (4) are also not significant.

Table 2: Results of Experiment 2

	(1)	(2)	(3)	(4)
Change in:	Belief		Donation	
T^0 * Baseline belief: >0	-0.18	-0.50	-0.06	0.12
	(0.12)	(0.33)	(0.37)	(1.52)
Observations	116	46	116	46
R-squared	0.55	0.16	0.03	0.01
Drop logically inconsistent		X		X

Notes: All regressions include a dummy variable for treatment and control for the baseline level of the belief. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 2 shows that the more salient treatment was still not effective. Following Staiger and Stock (1997), researchers often use a first-stage F statistic of 10 as a rule-of-thumb for a weak instrumental variable. For the regressions in both column (1) and column (2) of the table, the F statistic is less than 2.5. When an instrumental variable is weak, the second-stage estimates are biased and may not be consistent. The weakness of our instrument therefore precludes us from drawing strong conclusions about the second-stage relationship between beliefs and donations. In fact, comparing columns (2) and (4) would suggest that donations move in the opposite direction from beliefs, though the wide confidence intervals would also include the possibility that donations move in the same direction as beliefs and with five times the magnitude. The weak instrument is therefore not an obvious improvement over ordinary least squares regression, which could be biased away from the causal effect of changes in beliefs for a number of reasons, such as if those who don't give then manipulate their beliefs downwards as justification. We still expect that donations are increasing in beliefs, but an unbiased estimate of the causal effect of beliefs on donations will require a different setting or design for which the randomized treatment has a significant effect on beliefs.

In light of the weak effects that we find in our experiments, we ask whether a treatment that was more successful in shifting beliefs would have been likely to make a substantial difference in donations. To do so, we return to our most precise estimates of effects on beliefs, shown in column 8 of Table 1. For each subject in Experiment 1, we can calculate the predicted effect of applying the 0.97 cents-per-percentage-point effect size to the difference between the subject's baseline belief and the actual share of 0.35 who donated at least $1 at baseline. Averaging these predicted effects across all subjects, we find that convincing all subjects that 35 percent of others donate at least $1 would increase the $0.54 average donation in Experiment 1 to $0.63. If the treatment were only applied to the subjects whose baseline belief was below the truth, the average donation would rise to $0.69, a 56 percent increase in the average treatment effect. However, this calculation assumes that the treatment effect on beliefs is perfect and that the effect on donations is the same for all subjects as what we estimate for subjects with consistent baseline beliefs. Our research shows that these assumptions are unrealistic and that charities would likely need to find better ways to identify and treat beliefs to exploit belief heterogeneity for more than moderate gain.

5 Conclusion

Our experiment tests whether beliefs about one's own giving and that of others, elicited without incentives, predict how a person's giving responds to information treatments. Our baseline survey found sufficient dispersion in beliefs to test whether they are informative of treatment effects. Our experiments showed that information about others' donations affected only a small number of subjects' beliefs. As a result, we find no significant effects on donations. The lack of a first-stage effect on beliefs appears to be due to an inattentive group of subjects, many of whom reported baseline beliefs that were not logically consistent. Some of the point estimates suggest that it may be possible for charities to profitably exploit belief heterogeneity, but it appears that such a strategy would only be viable in settings in which engaged donors report meaningful beliefs. Such a setting would be required to obtain the first-stage effect on beliefs that will be necessary to obtain consistent estimates of the effect of beliefs on donations.

Our mixed results highlight the question of when and where it would be most likely that a charity could elicit beliefs that would predict heterogeneous effects of information about others' donations. Our setting is a solicitation for a small donation following a transaction, similar to those of purchasers outside of a supermarket (Andreoni et al., 2017) or sellers on eBay (Elfenbein et al., 2012). Subjects are therefore likely less engaged than a charity's "warm list" of past donors would be and perhaps more like the "cold list" of potential new donors. Consistent with this, Bekkers (2012) finds that similar treatments are only effective when potential donors are asked to estimate others' generosity just before the solicitation. Future research might test, then, whether beliefs are more predictive when solicited from a charity's warm-list donors. Regardless of the setting, our study demonstrates the importance of conveying the information about others' donations in a clear and salient way.

References

Adena, Maja and Steffen Huck, "Personalized fundraising: A field experiment on threshold matching of donations," *Available at SSRN 3480472*, 2019.

Alpizar, Francisco, Fredrik Carlsson, and Olof Johansson-Stenman, "Anonymity, reciprocity, and conformity: Evidence from voluntary contributions to a national park in Costa Rica," *Journal of Public Economics*, 2008, *92* (5-6), 1047–1060.

Andreoni, James, "Leadership Giving in Charitable Fund-Raising," *Journal of Public Economic Theory*, 2006, *8* (1), 1–22.

_ **and Douglas B. Bernheim**, "Social Image and the 50–50 Norm: A Theoretical and Experimental Analysis of Audience Effects," *Econometrica*, 2009, *77*, 1607–1636.

_ **, Justin M. Rao, and Hannah Trachtman**, "Avoiding the Ask: A Field Experiment on Altruism, Empathy, and Charitable Giving," *Journal of Political Economy*, 2017, *125* (3), 625–653.

Ariely, Dan, Anat Bracha, and Stephan Meier, "Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially," *American Economic Review*, 2009, *99*, 544–555.

Bardsley, Nicholas and Rupert Sausgruber, "Conformity and reciprocity in public good provision," *Journal of Economic Psychology*, 2005, *26* (5), 664–681.

Barreca, Alan I., Melanie Guldi, Jason M. Lindo, and Glen R. Waddell, "Saving Babies? Revisiting the effect of very low birth weight classification," *The Quarterly Journal of Economics*, 2011, *123* (4), 2117–2123.

Bekkers, René, "Limits of social influence on giving: Who is affected when and why," in "workshop Social influences and charitable giving, Centre for Market and Public Organisation, University of Bristol" 2012.

Bracha, Anat, Michael Menietti, and Lise Vesterlund, “Seeds to succeed?: Sequential giving to public projects,” *Journal of Public Economics*, 2011, *95* (5), 416–427.

Cagala, Tobias, Ulrich Glogowsky, and Johannes Rincke, “Who to Target in Fundraising? A Field Experiment on Gift Exchange,” 2015.

Carlsson, Fredrik, Haoran He, and Peter Martinsson, “Easy come, easy go-The role of windfall money in lab and field experiments,” *Experimental Economics*, 2013.

Castillo, Marco, Ragan Petrie, and Clarence Wardell, “Fundraising Through Online Social Networks: a Field Experiment on Peer-to-peer Solicitation,” *Journal of Public Economics*, 2014, *114*, 29–35.

__ , __ , and __ , “Friends Asking Friends for Charity: The Importance of Gifts and Audience,” 2017.

Chaudhuri, Ananish, “Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature,” *Experimental economics*, 2011, *14* (1), 47–83.

Chen, D.L., M. Schonger, and C. Wickens, “oTree - An Open-Source Platform for Laboratory, Online and Field Experiments,” *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Cherry, Todd L, Peter Frykblom, and Jason F Shogren, “Hardnose the Dictator,” *American Economic Review*, 2002, *92* (4), 1218–1221.

__ , Stephan Kroll, and Jason F Shogren, “The impact of endowment heterogeneity and origin on public good contributions: evidence from the lab,” *Journal of Economic Behavior & Organization*, 2005, *57* (3), 357–365.

Croson, Rachel TA, “Theories of commitment, altruism and reciprocity: Evidence from linear public goods games,” *Economic Inquiry*, 2007, *45* (2), 199–216.

DellaVigna, Stefano, James A. List, and Ulrike Malmendier, "Testing for Altruism and Social Pressure in Charitable Giving," *The Quarterly Journal of Economics*, 2012, *127(1)*, 1–56.

Drouvelis, Michalis, Adam Isen, and Benjamin M. Marx, "The Bonus-Income Donation Norm," *CESifo Working Paper No. 7961*, 2019.

__ **and Benjamin M. Marx**, "Dimensions of Donation Preferences: the Structure of Peer and Income Effects," *Experimental Economics*, 2020.

Elfenbein, Daniel W, Raymond Fisman, and Brian McManus, "Charity as a Substitute for Reputation: Evidence from an Online Marketplace," *Review of Economic Studies*, 2012, *79* (4), 1441–1468.

Elster, Jon, "Social norms and economic theory," *Journal of economic perspectives*, 1989, *3* (4), 99–117.

Erkal, Nisvan, Lata Gangadharan, and Nikos Nikiforakis, "Relative Earnings and Giving in a Real-Effort Experiment," *The American Economic Review*, 2011, *101* (7), 3330–3348.

Fischbacher, Urs and Simon Gachter, "Social preferences, beliefs, and the dynamics of free riding in public goods experiments," *American economic review*, 2010, *100* (1), 541–56.

__ **, Simon Gächter, and Ernst Fehr**, "Are People Conditionally Cooperative? Evidence fro a Public Goods Experiment," *Economics Letters*, 2001, *71*, 397–404.

Frey, Bruno S and Stephan Meier, "Social comparisons and pro-social behavior: Testing" conditional cooperation" in a field experiment," *The American Economic Review*, 2004, *94* (5), 1717–1722.

Gächter, Simon, "Conditional Cooperation. Behavioral Regularities from the Lab and the Field and their Policy Implications," in Bruno S. Frey and Alois Stutzer, eds., *Economics* *and Psychology. A Promising New Cross-Disciplinary Field*, CESifo Seminar Series, The MIT Press, 2007.

Garcia, Luzmarina, "The Fairer Sex? Court Function and How Women Rule," *mimeo, www.illinois.edu*, 2020.

Glazer, Amihai and Kai A Konrad, "A signaling explanation for charity," *The American Economic Review*, 1996, *86* (4), 1019–1028.

Harrison, Glenn W, "House money effects in public good experiments: Comment," *Experimental Economics*, 2007, *10* (4), 429–437.

Heldt, Tobias, "Conditional cooperation in the field: cross-country skiers' behavior in Sweden," in "Economic Science Association, European Meeting, Alessandria, Italien, September 15-18, 2005" 2005.

Hermalin, Ben, "Toward an Economic Theory of Leadership: Leading by Example," *American Economic Review*, 1998, *88(5)*, 1188–1206.

Huck, Steffen and Imran Rasul, "Matched fundraising: Evidence from a natural field experiment," *Journal of Public Economics*, 2011, *95*, 351–362.

_ , _ , and Andrew Shephard, "Comparing charitable fundraising schemes: Evidence from a field experiment and a structural model," *American Economic Journal: Economic Policy*, 2015, *7* (2), 326–69.

Karlan, Dean and John A List, "How Can Bill and Melinda Gates Increase Other People's Donations to Fund Public Goods?," *NBER Working Paper No. 17954*, 2012.

Keser, Claudia and Frans Van Winden, "Conditional cooperation and voluntary contributions to public goods," *scandinavian Journal of Economics*, 2000, *102* (1), 23–39.

Kessler, Judd B, "Announcements of Support and Public Good Provision," *American Economic Review*, 2017, *107* (12), 3760–3787.

Kessler, Judd B., Corinne Low, and Monica Singhal, "Social Policy Instruments and the Compliance Environment," 2020.

Klinowski, David, "Reluctant donors and their reactions to social information," *Experimental Economics*, 2020, pp. 1–21.

Kocher, Martin G, Todd Cherry, Stephan Kroll, Robert J Netzer, and Matthias Sutter, "Conditional cooperation on three continents," *Economics letters*, 2008, *101* (3), 175–178.

Kroll, Stephan, Todd L. Cherry, and Jason F. Shogren, "The Impact of Endowment Heterogeneity and Origin on Contributions in Best-shot Public Good Games," *Experimental Economics*, 2007, *10* (4), 411–428.

Krupka, Erin L and Rachel TA Croson, "The differential impact of social norms cues on charitable contributions," *Journal of Economic Behavior & Organization*, 2016, *128*, 149–158.

Krupka, Erin L. and Roberto A. Weber, "Identifying Social Norms Using Coordination Games: Why Does Dicator Game Sharing Vary?," 2013, *11* (3), 495–524.

Kumru, Cagri S. and Lise Vesterlund, "The Effect of Status on Charitable Giving," *Journal of Public Economic Theory*, 2010, *12* (3), 709–735.

Kurzban, Robert and Daniel Houser, "Experiments investigating cooperative types in humans: A complement to evolutionary theory and simulations," *Proceedings of the National Academy of Sciences*, 2005, *102* (5), 1803–1807.

List, John A and David Lucking-Reiley, "The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign," *Journal of Political Economy*, 2002, *110* (1), 215–233.

Martin, Richard and John Randal, "How is donation behaviour affected by the donations of others?," *Journal of Economic Behavior & Organization*, 2008, *67* (1), 228–238.

Meer, Jonathan, "Brother, Can You Spare a Dime? Peer Pressure in Charitable Solicitation," *Journal of Public Economics*, 2011, *95* (7-8), 926–941.

__ **and Harvey Rosen**, "The ABCs of Charitable Solicitation," *Journal of Public Economics*, 2011, *95* (5-6), 363–371.

Ostrom, Elinor, "Collective action and the evolution of social norms," *Journal of economic perspectives*, 2000, *14* (3), 137–158.

Ottoni-Wilhelm, Mark, Lise Vesterlund, and Huan Xie, "Why Do People Give? Testing Pure and Impure Altruism," *American Economic Review*, 2017, *107* (11), 3617–3633.

Potters, Jan, Martin Sefton, and Lise Vesterlund, "Leading-by-example and signaling in voluntary contribution games: an experimental study," *Economic Theory*, 2007, *33* (1), 169–182.

Reinstein, David and Gerhard Riener, "Reputation and Influence in Charitable Giving: an Experiment," *Theory and Decision*, 2012, *72(2)*, 221–243.

Shang, Jen and Rachel Croson, "A Field Experiment in Charitable Contribution: the Impact of Social Information on the Voluntary Provision of Public Goods," *Economic Journal*, 2009, *119*, 1422–1439.

Smith, Sarah, Frank Windmeijer, and Edmund Wright, "Peer Effects in Charitable Giving: Evidence from the (Running) Field," *The Economic Journal*, 2013, *125*, 1053–1071.

Staiger, Douglas and James H. Stock, "Instrumental VariablesRegression with Weak Instruments," 1997, *65*, 557–586.

Tonin, Mirco and Michael Vlassopoulos, "Sharing one's fortune? An experimental study on earned income and giving," *Journal of Behavioral and Experimental Economics*, 2017, (66), 112–118.

van Teunenbroek, Claire and René Bekkers, "Follow the crowd: Social information and crowdfunding donations in a large field experiment," *Journal of Behavioral Public Administration*, 2020, *3* (1).

_, _, and Bianca Beersma, "Look to others before you leap: A systematic literature review of social information effects on donation amounts," *Nonprofit and Voluntary Sector Quarterly*, 2020, *49* (1), 53–73.

Vesterlund, Lise, "The informational value of sequential fundraising," *Journal of Public Economics*, 2003, *87* (3), 627–657.