

Impediments to Effective Altruism:
Charity as a Taste-Based Decision

Jonathan Z. Berman

Alixandra Barasch

Emma E. Levine

Deborah A. Small

Working Paper

Please do not distribute without authors' permission

AUTHORS' NOTE

Jonathan Berman (jberman@london.edu) is an Assistant Professor of Marketing at the London Business School; Alixandra Barasch (abarasch@stern.nyu.edu) is an Assistant Professor of Marketing at the NYU Stern School of Business; Emma E. Levine (emma.levine@chicagobooth.edu) is an Assistant Professor of Behavior Sciences and William S. Fishman Faculty Scholar at the University of Chicago Booth School of Business; Deborah Small (deborahs@wharton.upenn.edu) is the Laura and John J. Pomerantz Professor of Marketing and Psychology at the University of Pennsylvania.

ABSTRACT

Charity could do the most good if every dollar donated went to causes that produced the greatest welfare gains. In line with this proposition, recent years have seen the rise of the “Effective Altruism” movement, which seeks to provide individuals with information about the effectiveness of charities in hopes that they will donate to organizations that maximize social welfare. This paper investigates charitable decision-making when effectiveness information is made available. We find that individuals view charitable giving as a relatively taste-based decision, often prioritizing their personal preferences at the expense of maximizing effectiveness. Specifically, we show that (a) people believe it is more appropriate to donate based on an emotional connection rather than to maximize welfare; (b) people are less likely to engage in value-maximization when choosing charities than investments; (c) decreasing the variety in a choice set inhibits preference matching and increases selection of an effective option; (d) increasing accountability increases choice share of effective charities. Implications for the effective altruism movement are discussed.

Charity could do the most good if donations went to the causes produced the greatest welfare gains. The “Effective Altruism” movement seeks to make this proposition a reality. Effective altruism draws upon consequentialist ethics to argue that instead of consulting one’s feelings, individuals should apply expected value maximization when deciding where to donate their money (Singer 2009, 2015; MacCaskill 2015). In this sense, effective altruists argue that individuals should make charity decisions like investors make financial decisions. Just as an investor’s aim is to maximize the financial return of each dollar invested, so should a donor aim to maximize the *social* utility return of each dollar donated.

In recent years, the effective altruism movement has given rise to organizations such as GiveWell.org and Giving What We Can, which seek out and provide information to the public about the charities they deem to be highly effective and do the most good for each dollar donated. However, it is unclear whether the provision of this information will have a demonstrable impact on behavior.

We investigate how individuals make donation decisions when information about a charity’s effectiveness is available and comparable across options. Throughout the paper, we utilize the term “effective” to refer to the extent to which a dollar allocated to a charity achieves utility gains for others. Thus, the more effective a charity is, the more that donating to the charity maximizes expected social welfare.

Models of Charitable Decision Making

Normative models of giving argue that given scarce resources, individuals should allocate funds that maximize the outcomes of their donations (Baron 1993, 2005; Baron & Szymanska, 2011; Ord, 2012; Singer, 1979, 2009, 2015). Consider the choice between donating to the Make-

A-Wish foundation, which grants wishes to children suffering from life-threatening diseases, and donating to Against Malaria, which provides insecticide nets to prevent malaria in the developing world. In 2015, the Make-A-Wish foundation reported that they spend, on average, \$7,500 per wish granted. In contrast, it is estimated that for every \$3,500 donated, Against Malaria can save a child's life. Certainly if the goal of charity is to maximize the utility of others, it is much more cost-effective to donate to Against Malaria than to Make-A-Wish. Yet, in 2015, Make-A-Wish raised over \$200 million in donations in the United States alone (Make-A-Wish, 2016), while Against Malaria raised less than a tenth of that amount worldwide (Against Malaria, 2016).

Descriptive theories of charitable decision-making help explain why individuals do not maximize the value of their charitable donations. In particular, models of “distorted altruism” argue that people intend to maximize utility for others, but they make systematic mistakes (Loewenstein & Small, 2007; Slovic, 2007). Evidence in support of distorted altruism comes from studies that vary whether options are presented in joint versus separate evaluation (Hsee, 1996). In separate evaluation, individuals frequently fail to maximize utility for others: people donate more to single individuals, or “identifiable victims” than large groups of “statistical victims” (Kogut & Ritov 2005); pay more attention to the ratio of victims rather than the absolute number of victims being helped (Featherstonhaugh, 1997); care about changes in states (e.g., someone recently lost their home) rather than end states (e.g., someone has been homeless for an extended period; Small 2010); and over-weight overhead expenses relative to outcomes achieved (Baron & Szymanska, 2011; Caviola, et al. 2014; Gneezy, Keenan, Gneezy, 2014).

However, in joint evaluation, when metrics of value are made directly comparable, these effects are attenuated or reversed (Caviola, et al. 2014; Korgut & Ritov 2005; Small 2010),

suggesting that people seek to maximize value, but without clear information to make comparisons salient, they instead rely on their feelings to guide choice.

One limitation with these existing studies on distorted altruism is that they typically have participants evaluate donation options from a common underlying cause or condition (e.g., humanitarian aid in Africa). In the real world, individuals must choose among a variety of charities that support various causes (e.g., research and treatment for a wide range of medical diseases, disaster relief, support for civil liberties, etc.). Yet it remains unclear how individuals make charitable decisions across different types of causes when effectiveness information is available.

Charity as a Taste-Based Decision

The present research investigates how individuals make decisions about charitable options when effectiveness information is present and comparable, making it easy for individuals to maximize social utility if they desire to do so. We find that even when effectiveness information is comparable, a substantial portion of people do not choose to maximize social welfare. That is, despite knowing that Charity A is more effective than Charity B, a substantial portion of people still choose to donate to Charity B. Our results are consistent with the notion that many view charity as a taste-based decision, one in which their subjective preferences and intuitions are valid and appropriate in the decision-making process. In other words, when donating to charity, people view it as normative to select a cause that they personally prefer, even if the cause does not maximize social welfare.

We conducted an initial survey with participants from a Northeastern University's behavioral lab ($N = 126$) to investigate whether people believe that charitable giving decisions

should be a matter of subjective tastes or should be determined by objective metrics. Participants were asked to evaluate how they believe people should make decisions across six decision-making domains (charities, investments, cell phones, restaurants, art, and medical treatment). Specifically, they evaluated the extent to which they agreed with the statements: “It is important that the _____ I choose reflects my personal tastes or values”; “It is more important to rely on objective measures rather than personal feelings when choosing _____”; and “Objective measures are the best way to choose _____” on scales ranging from 1 = “Strongly disagree” to 7 = “Strongly agree.”¹

Table 1 displays the results for each question. Participants indicated that it is more important to rely on one’s tastes and personal feelings, rather than objective measures, when choosing a charity compared to choosing a cell phone, an investment, or a medical treatment. The importance of relying on personal tastes when making a charitable decision was roughly on par with choosing a restaurant. Only choosing a piece of art was evaluated to be more a matter of taste than choosing a charity.

This survey provides initial evidence that people think of charity as a relatively taste-based decision. However, it is possible that these evaluations reflect individuals’ beliefs for cases for which effectiveness is unavailable, and they might react differently when effectiveness information is present. The studies we describe in this paper directly examine whether judgments and choices conform to these meta-preferences when effectiveness is available, and further examine when individuals are more or less likely to utilize effectiveness information when choosing where to donate.

¹ See supporting materials for results of additional questions asked in the survey.

² We additionally ran a two-way ANOVA with decision type, search choice, and their interaction predicting rank of selected option. Results show a main effect of decision type, $F(1, 397) = 7.79, p = .006$, a main effect of search choice, $F(1, 397) = 69.19, p < .001$, and no interaction $F(1, 397) = 1.40, p = .24$. Regardless of how participants search, they choose a lower ranked option in the charity condition.

³ Across studies, we found varying results on this measure. For instance, in Study 2b, there was no difference

	Art	Charities	Restaurants	Cell Phones	Investments	Medical Treatments
It is important that the _____ I choose reflects my personal tastes or values.	6.01 (1.27)	5.90 (1.17)	5.52** (1.20)	5.22*** (1.40)	4.70*** (1.60)	4.63*** (1.65)
It is more important to rely on objective measures rather than personal feelings when choosing _____.	<i>2.48***</i> <i>(1.47)</i>	3.55 (1.50)	3.33 (1.43)	4.46*** (1.56)	5.20*** (1.49)	4.98*** (1.59)
Objective measures are the best way to choose _____.	<i>2.48***</i> <i>(1.28)</i>	3.79 (1.50)	3.95 (1.52)	5.10*** (1.26)	5.21*** (1.34)	5.26*** (1.35)

Notes: Means and standard deviations for responses to each statement. Values in **bold** indicate that given decision is rated as being significantly less taste-based than charitable decisions. *Italics* indicate a given decision is rated as being significantly more taste-based than charitable decisions. Astericks correspond to the significance level when comparing to the charity condition. ** $p < .01$, *** $p < .001$

Table 1: Personal tastes and feelings are considered more valid when making charitable decisions than when choosing a cell phone, investment, or medical treatment.

We report four main findings. First, we evaluate people's beliefs about the normativity of choosing an ineffective charity and find that people believe that individuals should donate to charities they feel emotionally connected to, even if the charity does not maximize social welfare (Study 1).

Second, we analyze people's tendency to rely on their subjective preferences in their own behavior. We directly compare how individuals make decisions for charities to how they make investment decisions, and find that people are significantly more likely to value-maximize when choosing an investment than they are when choosing a charity (Study 2; Supplemental S1 & S2). When selecting an investment, people are less concerned with satisfying their personal tastes and more concerned with maximizing outcomes.

Third, we examine how choice sets impact charitable decisions and find that people are more likely to donate to the most effective charity when the choice set contains options belonging to a similar underlying cause than when the choice set contains options that are more differentiated (Studies 3a & 3b). When all options support the same cause, there is no conflict between tastes and effectiveness, but as the underlying causes in a choice set become more differentiated, people are more likely to find an option that they subjectively favor. We demonstrate these results using real charities and accurate information regarding effectiveness as well as in a hypothetical scenario where effectiveness information can be experimentally manipulated.

Finally, we find that increasing accountability to others reduces reliance on personal tastes when choosing charities: people are less likely to license themselves (Study 4) and others (Study 5) to select an ineffective option when a decision-maker is made responsible for the wellbeing of others.

In all studies, our sample size was determined in advance. No conditions or participants were dropped from any analysis performed. All measures assessed that are not reported in the manuscript can be found in the supporting materials section. All data is available online at (https://osf.io/6myfj/?view_only=9b9ce0e42ca1493081fe2e5a5d070ef7)

STUDY 1: PEOPLE BELIEVE IT IS NORMATIVE TO PRIORITIZE PERSONAL FEELINGS OVER WELFARE GAINS

This study examines whether people believe it is normatively appropriate to reject the value-maximizing charity option and chose a charity based on their personal tastes. Participants were presented with a scenario involving a woman, Mary, faced with a decision to donate either to a local charity or to a charity in a third-world country. Mary was also described as having an emotional connection towards either near or distant causes. Consistent with research showing that emotions are valued in charitable decision-making (Barasch et al., 2014), we expected participants would indicate that Mary should donate to the charity to which she felt an emotional connection, regardless of whether they felt that this charity was most effective.

Method

We pre-registered all aspects of this study on AsPredicted.org (<http://aspredicted.org/blind.php/?x=jd9cv4>). We aimed to recruit 400 participants from Mechanical Turk and ended up with a final sample of 404 participants (Mean age = 34.8; 49.6% male, 49.4% female, 1.0% did not disclose). Participants read a scenario involving Mary, who was looking to donate money to one of two hypothetical charities. One charity, “Hunger Care in Africa,” was described as providing micronutrients for children in Africa, and for each \$100 donated, this charity could feed five children for a month. A second charity, “Jump Start Your

Community,” was described as training and educating homeless in Mary’s community, and for every \$100 donated, this charity could provide two hours of job training to one homeless person.

Participants were then randomized to one of two Emotional Connection conditions. In the Local Causes condition, participants read that Mary feels an emotional connection towards helping those in local communities, whereas in the Distant Causes condition, participants read that Mary feels an emotional connection towards helping those in third-world countries. See Appendix A for scenario wording.

We then presented participants with two different questions on separate pages in random order. One question assessed where they felt Mary should donate (“Which charity should Mary donate to?”), while the other question asked which charity was the most effective at maximizing social welfare (“Suppose that Mary wanted to donate to the charity that did the greatest good for the greatest number of people. Which one should she donate to?”). Both questions were measured on seven-point scales ranging from 1 = “Definitely Jump Start Your Community” to 7 = “Definitely Hunger Care in Africa”.

Results

This study design allows us to test our hypotheses either by evaluating the first question posed to participants in a between-subjects analysis, or by evaluating both questions asked to participants in a within-subjects analysis. In the pre-registration, we specified that we would utilize a between-subjects analysis to test our focal hypothesis as we were unsure whether any order effects would emerge in the within-subjects design. We report the between-subjects analysis below. Nonetheless, we replicate these results in a within-subjects analysis, regardless of question order, as reported in the supporting materials.

A two-way ANOVA found a significant Emotional Connection \times Response Question interaction in the predicted direction, $F(1, 399) = 8.80, p = .003, \eta^2 = .02$. When participants read that Mary felt an emotional connection towards local communities, they indicated that she should donate to “Jump Start Your Community” ($M = 3.00, SD = 1.99$), yet when asked which charity did the most good, participants favored “Hunger Care For Africa” ($M = 4.55, SD = 2.32$), $t(202) = -5.12, p < .001, d = -0.72$. However, when participants read that Mary felt an emotional connection towards distant communities, this divergence between what she should do and what would do the most good disappeared: participants indicated that she should donate to “Hunger Care For Africa” ($M = 5.26, SD = 2.05$) and evaluated that “Hunger Care For Africa” was more effective ($M = 5.59, SD = 1.87$), $t(197) = -1.19, p = .24, d = -0.17$. Figure 1 displays these results.

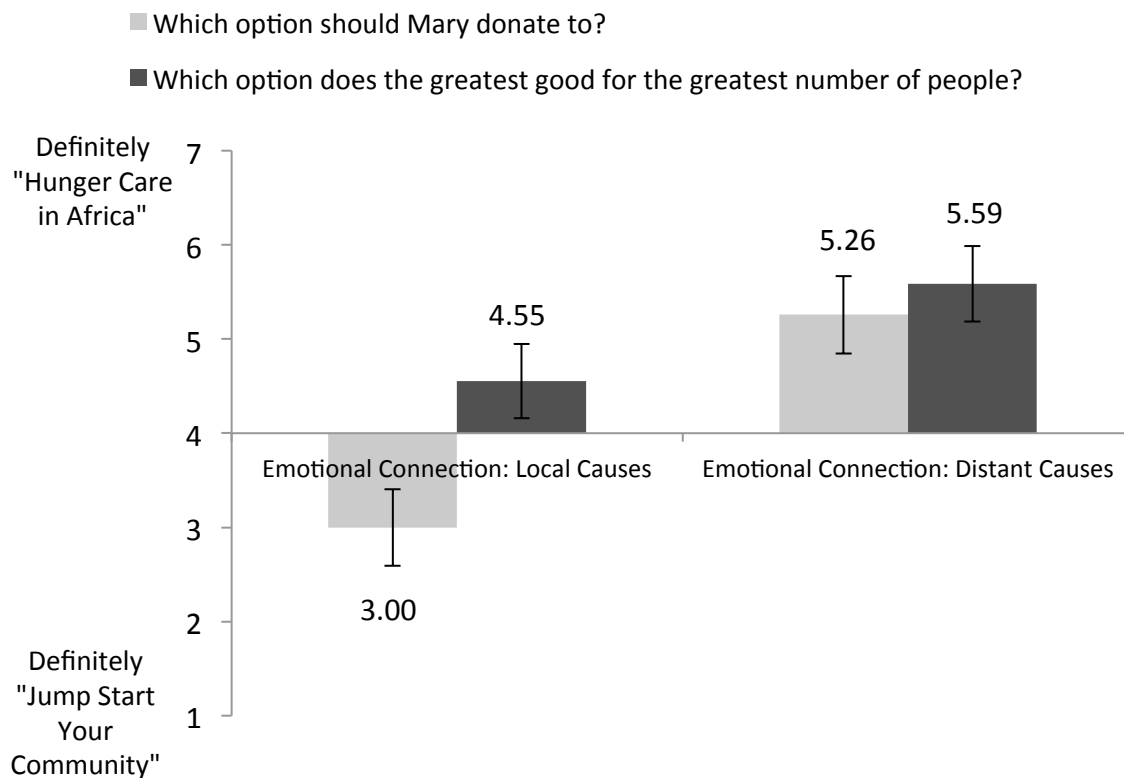


Figure 1. Participants believe that Mary should donate to the cause for which she feels an emotional connection, even if they do not believe that it maximizes social welfare. Error bars represent 95% confidence intervals.

Discussion

Study 1 shows that people believe that others should donate to charities that they emotionally favor even when the charity is not maximally effective. In the next set of studies, we move beyond how individuals evaluate others, and investigate the extent to which they seek to maximize social welfare in their own decisions.

STUDIES 2: PEOPLE VALUE MAXIMIZE FOR INVESTMENTS MORE THAN CHARITIES

In Study 2 we compare how individuals make charitable decisions relative to how they make investment decisions. Proponents of effective altruism argue that in order to do the most good, individuals should maximize expected welfare just as investors should seek to maximize financial returns. We expect that when faced with investment decisions, individuals will seek to maximize the expected value of their money, and will strongly weigh this information when making a choice. However, for charitable decisions, we expect that individuals will care less about maximizing expected value of their donation, and will more heavily weigh their personal tastes when making a choice. To test this, we examine the processes that individuals use to make charity versus investment decisions, and whether individuals choose sort options within a choice set according to how effective they are or what causes they represent. We then examine the subsequent choices participants make.

Method

We recruited 401 participants from Mechanical Turk (Mean age = 30.5; 31.4% female, 68.6% male). Participants were randomly assigned to imagine that they decided either to donate \$250 to a charity (Charity condition) or to invest \$250 into a stock (Investment condition) through a website recommended by a friend.

Participants then read (Investment condition in brackets): “The website is renowned for being extremely accurate in evaluating which charities [stocks] are the most effective and do the most good for each dollar donated [provide the best return on each dollar invested]. The website provides a 0 to 100 rating: the higher the number, the better the rating. In order to be listed on the website, each charity [stock] must have a rating of at least 70.” Participants were also told that the website categorizes options into six different categories (Alternative Energy, Cancer, Education, Food, Housing, and International). These categories were chosen because they could plausibly represent groupings of either charities or investments.

Participants were then given the choice to sort the options either by effectiveness or by category type. This served as our primary dependent variable for this study. Sorting by effectiveness suggests that an individual is primarily concerned with selecting one of the most effective charities (investments) whereas sorting by category suggests that an individual is primarily concerned with selecting a charity (investment) that reflects a preferred cause.

They then were presented with the list of options, shown in accordance with their sorting choice: a screen displayed six interactive headings that ordered the options either by effectiveness rating (100-96; 95-91; 90-86; 85-81; 80-76; 75-70) or by category type (Alternative Energy; Cancer; Education; Food; Housing; International). Participants could click on each heading in order to reveal the options within each heading (see Appendix B).

We created three options for each of the six different category types, for a total of 18 options. The options were named such that they could represent either charities or investments (see Appendix B for full list). Only the name, category, and effectiveness rating of each option was displayed (e.g., “Arcadia Grocers | Category = Food | Rating = 85”). Each option was randomly assigned an effectiveness rating between 70 to 100. We further structured the randomization so three options were assigned a rating between 100 and 96; three options were assigned a rating between 95 and 91; three options were assigned a rating between 90 and 86; and so on down to 70. This was done so that regardless of how participants chose to sort the options, there would always be three options under each heading.

After selecting their preferred option, participants evaluated a series of statements intended to probe their decision process. For each statement, participants rated the extent to which they agreed or disagreed with the statements on a seven-point scales ranging from 1 = “Strongly Disagree” to 7 = “Strongly Agree.”

We expected that the reason why individuals are more likely to select a less effective charitable option is that participants are more likely to weigh their subjective preferences when making charitable decisions. To examine this, participants evaluated four statements, “It is important for me to choose charities [investments] that I personally care about”; “It is important for me to choose charities [investments] that reflect my personal views and beliefs”; “Objective measures are the best way to select a charity [investment]” (reverse coded) and “It is more important to rely on objective measures rather than personal feelings when choosing a charity [investment]” (reverse coded), and were averaged to create a four-item Subjective Preferences measure ($\alpha = .73$).

We also included questions intended to probe three additional factors that may have affected how participants made their decision. First, it is possible that participants attended less to effectiveness ratings for charity because they wished to have agency when making charitable decisions (Andreoni, 1990; Harbaugh et al., 2007), or may have reacted negatively if they felt like they were being directed where to donate their money (Brehm, 1966). To examine this, participants evaluated two statements, “It is important that I actively select charities [investments] rather than letting someone else decide for me” and “I prefer that other people tell me where to donate [invest] my money than deciding for myself”(reverse coded), which were averaged to create a Preference for Agency measure ($r = .54$).

Second, it is possible that participants felt that they had more subjective knowledge about charitable decisions than investment decisions, and thus felt qualified to rely less on expert ratings when making charitable decisions. To examine this, participants evaluated two statements, “I feel that I am quite knowledgeable about charities [stocks]” and “I typically know a great deal about choosing the best charity to donate my money [stock to invest my money]”, which were averaged to create a Subjective Knowledge measure ($r = .81$).

Finally, it is possible that participants simply did not believe that evaluating charities by effectiveness is possible, and as a result, they discounted effectiveness information when objective evaluation was difficult to assess (Inbar, Cone & Gilovich, 2010). To examine this, participants evaluated the following statement: “It is possible for charities [investments] to be objectively rated according to their effectiveness.”

Results

Consistent with our hypothesis, significantly fewer participants chose to sort by effectiveness rating in the Charity condition (67.8%) than in the Investment condition (83.4%), $\chi^2(1) = 13.20, p < .001; \phi = .18$.

We supplement this analysis by examining the effectiveness of the options chosen. We first evaluated the percent of participants choosing to value-maximize by selecting the option with the highest effectiveness rating. Significantly fewer participants chose the highest rated option in the Charity condition (32.2%) than in the Investment condition (50.3%), $\chi^2(1) = 13.52, p < .001; \phi = .18$.

We further examine the rank of the selected option relative to the other options in the choice set, whereby the most effective option was assigned a rank of 1 and the least effective option was assigned a rank of 18. An ordinal regression with decision type (Charity vs. Investment) as the IV and rank of selected option as the DV, showed that the decision type affected the rank of the option selected, $\beta = .73, SE = .18, \chi^2(1) = 15.93, p < .001, OR = 2.08$. The average rank of the chosen option was lower in the Charity condition (Mean Rank = 4.11, $SD = 4.00$) than in the Investment condition (Mean Rank = 2.78, $SD = 2.83$)². Figure 2 displays these results broken down by search choice.

Additional Measures: Participants indicated that they weighed their subjective preferences more heavily in the Charity condition ($M = 4.51, SD = 0.87$) than the Investment condition ($M = 3.54, SD = 1.03$), $F(1, 399) = 104.42, p < .001, d = 1.02$. Participants also preferred to have more agency when making charitable decisions ($M = 5.66, SD = 1.23$) than when making investment decisions ($M = 4.81, SD = 1.29$), $F(1, 399) = 46.13, p < .001, d = .68$.

² We additionally ran a two-way ANOVA with decision type, search choice, and their interaction predicting rank of selected option. Results show a main effect of decision type, $F(1, 397) = 7.79, p = .006$, a main effect of search choice, $F(1, 397) = 69.19, p < .001$, and no interaction $F(1, 397) = 1.40, p = .24$. Regardless of how participants search, they choose a lower ranked option in the charity condition.

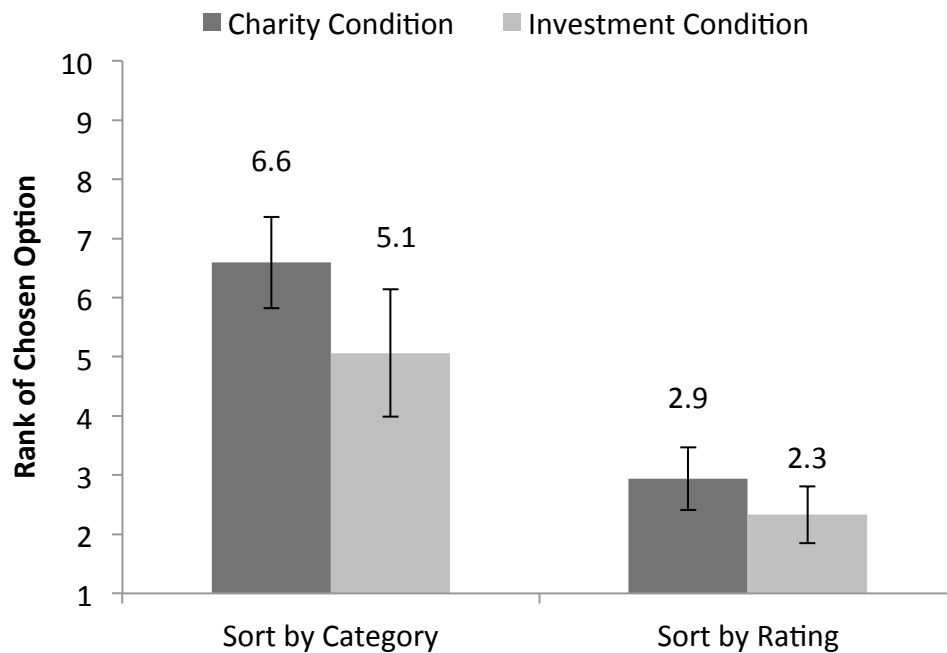


Figure 2. Mean rank of the selected option by condition and sort choice. Lower numbers (i.e., ranks closer to 1) indicate more effective options. Error bars represent 95% confidence intervals.

There was no statistically significant difference between subjective knowledge for charitable decisions ($M = 3.87$, $SD = 1.51$) and investment decisions ($M = 3.62$, $SD = 1.59$), $F(1, 399) = 2.68$, $p = .10$, $d = .16$. Finally, participants believed that it is more possible to objectively rate the effectiveness of charities ($M = 5.39$, $SD = 1.23$) than the effectiveness of investments ($M = 5.07$, $SD = 1.19$), $F(1, 399) = 7.27$, $p = .007$, $d = .27$.³ This suggests that people are not utilizing effectiveness ratings less for charitable decisions, because they believe charities are less objectively evaluable (Inbar, Cone, & Gilovich, 2010).

Mediation Analysis: We conducted a multiple mediation analysis using the bootstrap method with 5,000 samples to test what psychological factors drive participants' sort choice decision (SPSS Process Macro, Model 4; Hayes, 2013). The model included decision type (0 =

³ Across studies, we found varying results on this measure. For instance, in Study 2b, there was no difference between charities and investments on this measure. However, in the pilot survey, participants reported that it was more possible to rate investment decisions based on effectiveness than charities (see supporting materials).

investment decision, 1 = charity decision) as the IV, and sort choice as a binary DV (0 = search by effectiveness rating, 1 = search by category type). We included the Subjective Preferences, Preference for Agency, Subjective Knowledge, and Possible to Rate measures as independent mediator variables. Results show a significant indirect effect only on the Subjective Preferences measure (Indirect Effect = 0.68; 95% C.I. = [0.34, 1.11]). Relative to investment decisions, when making charity decisions, individuals increase the weight placed on subjective preferences ($a = 0.97$, $SE = .10$, $p < .001$), and the more weight placed on subjective preferences, the more likely participants were to sort by category type ($b = 0.69$, $SE = .16$, $p < .001$). Once controlling for the mediator, the effect of decision type on search choice, from $c = 0.87$, $SE = .24$, $p < .001$ to $c' = 0.25$, $SE = .29$, $p = .38$, suggesting full mediation. We did not find significant mediation for the Preference For Agency (Indirect Effect = 0.04; 95% C.I. = [-0.16, 0.25]), Subjective Knowledge (Indirect Effect = 0.05; 95% C.I. = [-0.01, 0.17]), or the Possible to Rate (Indirect Effect = -0.03; 95% C.I. = [-0.14, 0.03]) measures, as indicated by indirect effects that include zero.

Discussion

In Study 2 we find that participants are less likely to sort information by effectiveness ratings for charities than they are for investments. Moreover, participants were less likely to value-maximize for investments, regardless of search choice.

To further test the robustness of these results, we ran an additional study that directly assigned participants to the search by category condition of Study 2. We again found that fewer participants chose the highest rated charity (22.0%) than the highest rated investment (51.3%), $\chi^2(1, N = 224) = 20.58$, $p < .001$; $\phi = .30$ (Supporting Materials, Study S1). It is also possible that participants are more likely to satisfice for charities than investments, and these results would not

replicate in small choice sets. Yet when we ran a study that presented participants with choice sets containing only three options, we again find that significantly fewer still chose the highest rated charity option (46%) than the highest rated investment (64%), $\chi^2(1, N = 201) = 6.91, p = .009, \phi = .19$ (Supporting Materials, Study S2).

In the next set of studies, we investigate features of the choice context that influence participants' likelihood of selecting an option that maximizes social welfare.

STUDIES 3A & 3B: DECREASING DIFFERENTIATION INCREASES VALUE MAXIMIZATION

Studies 3a & 3b examine how the composition of a choice set affects participants' likelihood of selecting a charity that maximizes effectiveness. We expect that when individuals are faced with highly differentiated options, their tastes are more likely to align with one cause more than the others, and will select such an option at the expense of maximizing effectiveness. Study 3a examines this hypothesis using real donation behavior and presents participants with a set of charities, one of which has been endorsed by an effective altruism organization. Studies 3b replicates these findings in a more controlled setting and utilizes QALY (Quality Adjusted Life-Years), a standardized measure of an intervention's effectiveness to evaluate options.

Study 3a Method

We recruited 405 participants from Mechanical Turk (mean age = 36.8; gender = 47.7% male, 51.9% female, .5% did not disclose). Participants first completed a short data entry task that involved translating numbers into words (e.g., "632" into "six hundred thirty-two"). After completing the task, participants learned that as a thank you, we would donate \$0.20 to a charity

on their behalf, and that they would have their choice of donating this bonus to one of three charities. Participants also read that we had consulted organizations (specifically, GiveWell.org and TheLifeYouCanSave.org) to identify charities that are deemed to be highly effective, and do the greatest amount of good for each dollar donated.

Participants then viewed a table consisting of the three charities. The name of the charity, the type of charity, and a short description of each charity was displayed in the table, along with a column indicating whether the charity had been identified to be highly effective.

Those in the Single-Category Condition chose between three charities all belonging to the same underlying cause (either three blindness, three children's health, or three clean water charities). Those in the Mixed-Category Condition chose between three charities all belonging to different underlying causes (one blindness, one children's health, and one clean water charity). Regardless of condition, only one of the three options was identified as being highly effective. The presentation order of the options was fully randomized (See Appendix C for a full list of charities and a sample of the stimuli shown). Participants then made their selection. All the information provided to participants was accurate, and donations were made to charities in accordance with each participant's choice.

Study 3a Results

Consistent with our hypothesis, significantly fewer participants selected the option endorsed by effective altruism organizations in the Mixed Category condition (41.3%) than in the Single Category condition (54.3%), $\chi^2(1) = 6.87, p = .009, \phi = .13$.

Study 3b Method

We recruited 304 participants from Mechanical Turk (mean age = 30.4; gender = 66.4% male, 32.6% female, 1% did not disclose). Participants imagined they decided to donate \$250 to a charity through a website recommended by a friend. They were told that the website reports QALY statistics for each charity, defined as follows: “QALY stands for quality-adjusted life-year, and represents the number of years of life that would be added to someone in need by a charity's intervention. Specifically, the website lists QALY per ten thousand dollars donated. Thus, a QALY = 10 means that for every ten thousand dollars donated, an individual in need will gain ten years of healthy living to their life.”

We created three different categories of charities (Food & Water; Health & Medicine; International) and three options per category, for a total of nine options (see Appendix D). In the Single Category condition, participants were presented with three options belonging to a single category type. In the Mixed Category condition, participants were presented with three options, such that one option was randomly selected from each of the three category types. Each option was given a randomly determined QALY rating between 10 and 20. Thus, we randomly assigned options to effectiveness ratings. Only the name, the category, and the QALY of each option was displayed (e.g., “Medical Goods Initiative | Category = Health & Medicine | QALY = 18”). The presentation order of the options was also randomized.

Study 3b Results

We examined which options participants chose across the two conditions. Significantly fewer participants selected the top-ranked option (i.e., the option with the highest QALY rating) in the Mixed Category condition (69.3%) than in the Single Category condition (84.1%), $\chi^2(1) = 9.33, p = .002, \phi = .18$. In contrast, significantly more participants selected the middle-ranked

option in the Mixed Category condition (22.9%) than in the Single Category condition (9.9%), $\chi^2(1) = 9.29, p = .002; \phi = .17$. There was no significant difference in selection of the option with the lowest-ranked option (Mixed Category = 7.8% vs. Single Category = 6.0%), $\chi^2(1) = 0.42, p = .52; \phi = .04$.

We supplement this analysis by computing the rank of the option selected in terms of its QALY relative to the QALY rating of the other options. To assess whether the choice set composition affected which option participants selected, results were analyzed with an ordinal regression, with choice-set condition as the IV and QALY rank of the chosen option as the DV. Results show that those in the Single Category condition chose higher-ranked options than those in the Mixed Category condition, $\beta = 0.81, SE = 0.28, \chi^2(1) = 8.31, p = .004, OR = 2.24$.

Discussion

Together, these studies show that when the options within a choice set are highly differentiated, an individual's tastes are more likely to align with one cause more than the others, and will select that cause at the expense of maximizing effectiveness. We further document these results with both real (Study 3a) and hypothetical (Study 3b) donation decisions.

STUDY 4: INCREASING RESPONSIBILITY TOWARDS THE AFFECTED INCREASES VALUE MAXIMIZATION

In Study 4, we investigate circumstances in which people may be more likely to utilize effectiveness ratings when allocating resources to help others. One reason why individuals may feel licensed to select a charitable option that they personally prefer is that they do not face personal consequences for their actions, and are thus less sensitive to outcomes that affect others

compared to the self (Imas, 2014). For instance, earlier, we found that people are more likely to utilize outcome information for investments than charities. With investment decisions, individuals make less money if they choose an ineffective investment. In contrast, individuals do not directly experience the consequences when donating to charity, and thus feel permitted to select an option that does not maximize expected social welfare.

However, when individuals assume a role of responsibility to a larger group, they may face personal consequences for their decisions. For example, when acting on behalf of a larger organization, individuals may be expected to make a justifiable decision that appeals to all stakeholders, not just the tastes of the decision-maker.

In this study, we examine if people are more likely to maximize the social welfare when they assume a role of responsibility towards the affected parties. To test this, we experimentally manipulate the role that participants adopted while making a resource allocation decision. We additionally manipulate the presence or absence of effectiveness information. We do not expect role to affect choices when effectiveness information is absent. However, when effectiveness information is present, we expect that participants would be more likely to utilize effectiveness information when they assume a role of responsibility for the organization than if they are an independent donor.

Method

We recruited 419 participants (Mean age = 27.4; Gender: 36.5% male, 61.8% female, 1.7% did not disclose) from a behavioral lab subject pool in the United Kingdom.

We conducted a 2 (Role: Donor vs. Medical Center President) \times 2 (Effectiveness Information: Present vs. Absent) between-subjects design. Participants assigned to the Donor

condition imagined that they lived near a local medical research center. Participants assigned to the President condition imagined that they were the President of a local medical research center. They were then told (President Condition in brackets): “Because the stock market has done well, you have [the research center has] more cash than usual. You would like to help by making a donation to medical research [As President, it is part of your job to allocate these funds to conduct additional medical research.].” They were then told that they could select one research group in the medical center to allocate the cash.

Participants in the Effectiveness Ratings-Present condition then proceeded to read about an independent research firm that specializes in analyzing patient benefits for the medical research community. The firm evaluated research departments on a scale from 0 to 100, whereby the higher the number, the more effective additional funds would be for helping patients. Participants were then given their choice between allocating funds to one of three departments, with effectiveness ratings provided for each department (Arthritis [Rating = 92]; Heart Disease [Rating = 86]; Cancer [Rating = 74]). These ratings were purposefully selected so that the most intuitively appealing choice was rated as the least effective (Cancer), and the least appealing choice was rated as the most effective (Arthritis), thereby creating a tension between the option that would be most personally gratifying and the option that would be value-maximizing.

Participants in the Effectiveness Ratings-Absent condition did not read about the independent research firm, and were simply presented with the same three options, with no mention of effectiveness ratings.

Results

We conducted an ordinal regression with Role (Donor vs. President), Effectiveness Ratings (Present vs. Absent), and their interaction in a model to predict the rank of chosen option. Results revealed a significant Role \times Effectiveness Ratings interaction, $\chi^2(1) = 3.95$, $p = .047$. To probe this interaction, we ran separate ordinal regressions for the Effectiveness Ratings-Absent and Effectiveness Ratings-Present conditions. When the effectiveness ratings were absent, the assigned role did not significantly affect the option that participants chose, $\beta = .05$, $SE = .27$, $\chi^2(1) = 0.03$, $p = .87$, $OR = 1.05$. However, when the ratings were present, participants in the President condition were significantly more likely to select a higher-rated options than those in the Donor condition, $\beta = .79$, $SE = .26$, $\chi^2(1) = 9.70$, $p = .002$, $OR = 2.20$. Figure 4 displays these results.




Figure 4: Percent of participants choosing to fund a department as a function of their role (Donor vs. President) and whether effectiveness information was provided. The effectiveness ratings of options were only known to those in the Effectiveness Ratings-Present condition. Effectiveness information has a bigger impact on the choices that people make when placed in the role of a President than when placed in the role of a donor.

Discussion

Study 4 shows that individuals are more likely to donate to an ineffective charity when they make a personal decision. However, when they assume a role in which they are responsible towards the affected parties, they are more likely to utilize effectiveness ratings. One may wonder if participants in this study preferred to donate to aid deadly diseases (cancer, heart disease) as opposed to a non-deadly one (arthritis). However, this cannot explain why they were less likely to utilize effectiveness ratings in the donor condition compared to the President condition. In the next study, we extend these results to examine how people judge others who either utilize or disregard effectiveness information when making decisions.

STUDY 5: IS SELECTING A LOW-EFFECTIVENESS OPTION CONSIDERED A BAD DECISION?

In study 5, we examine how individuals judge the decisions of others who select either a relatively effective or ineffective option. The purpose of this study is two-fold. First, we test whether external judges recognize the validity of effectiveness ratings as a guide to charitable decisions. If so, then individuals who choose a highly effective option will be seen as making a higher quality and more altruistic decision than those who select effective less effective option. That is, we expect that people will acknowledge that selecting a highly effective option is indicative of making a good decision.

Second, we examine if individuals who make independent donations are evaluated differently than those who are responsible to others. If so, then the main effect described above should be moderated by accountability. In particular, we expect that those who are directly accountable for the well-being of others will be viewed particularly harshly for choosing to

allocate funds to an ineffective option. In contrast, for personal contributions, we expect that individuals will be given more leeway to make a decision that they subjectively prefer.

Method

We recruited 227 participants (Mean age = 23.5; Gender: 33.5% male, 65.6% female, 0.9% did not disclose) from a behavioral lab pool at a Northeastern university. We conducted a 2 (Role: Donor vs. Medical Center President) \times 2 (Choice: Most Effective Option vs. Least Effective Option) between-subjects design.

Participants read a scenario about Mr. Peterson, an individual who was allocating money to a local medical research center. In the Donor condition, participants read that due to recent stock market gains, Mr. Peterson was looking to donate money to a department within the medical center. In the President condition, participants read that Mr. Peterson was the President of the medical research center, and that due to recent stock market gains, he was looking to allocate surplus funds to a department within the medical center.

In all conditions, participants then read about an independent research firm that evaluates the effectiveness of departments in the same fashion as the previous study. Participants were presented with a list of three departments (Elderly Care, Heart Disease, and Arthritis), the rating for each department, and which department Mr. Peterson ultimately chose to fund. Mr. Peterson either chose the top-rated option or the bottom rated option. We varied whether elderly care or arthritis was the top-rated option to account for the possibility that people believe funding a specific department is particularly important for a President of a medical center. This manipulation had no effect on any of our DVs, thus we collapse across this factor when

presenting our analyses. Heart disease was held constant as the middle option and was never chosen by Mr. Peterson.

Participants then evaluated the extent to which they thought the decision was responsible, appropriate, and thoughtful on a scale ranging from 1 = “Not at all” to 7 = “Very much so”.

These were averaged to create a three-item Perceived Decision Quality measure ($\alpha = .92$).

Participants also evaluated the extent to which they thought the decision was ethical, selfish (reverse coded), and good on the same seven-point scale. These were averaged to create a three-item Perceived Altruism measure ($\alpha = .72$).

Results

Perceived Decision Quality: A two-way ANOVA revealed a significant Role \times Choice interaction, $F(1,223) = 6.64, p = .01, \eta^2 = .03$. In both the President and the Donor conditions, participants felt that Mr. Peterson’s decision was of lower quality when he selected the least effective option over the most effective option. However, the drop in decision quality was significantly greater when he was President ($M_{\text{MostEffectiveOption}} = 5.65, SD = 1.07$ vs. $M_{\text{LeastEffectiveOption}} = 3.89, SD = 1.35$), $t(114) = 7.64, p < .001, d = 1.46$, than when he was a donor ($M_{\text{MostEffectiveOption}} = 5.56, SD = 1.09$ vs. $M_{\text{LeastEffectiveOption}} = 4.83, SD = 1.30$), $t(109) = 4.09, p < .001, d = .61$. Figure 5 displays these results.

Perceived Altruism: A two-way ANOVA revealed a significant Role \times Choice interaction, $F(1,223) = 10.23, p = .002, \eta^2 = .04$. In both the President and the Donor conditions, participants felt that Mr. Peterson’s decision was less altruistic when he selected the least effective option over the most effective option. However, this difference was greater when Mr. Peterson was President ($M_{\text{MostEffectiveOption}} = 5.75, SD = 0.86$ vs. $M_{\text{LeastEffectiveOption}} = 4.30, SD =$

1.09), $t(114) = 7.85, p < .001, d = 1.49$, than when he was a donor ($M_{\text{MostEffectiveOption}} = 5.78, SD = 0.91$ vs. $M_{\text{LeastEffectiveOption}} = 5.19, SD = 1.00$), $t(109) = 3.30, p = .001, d = .62$.

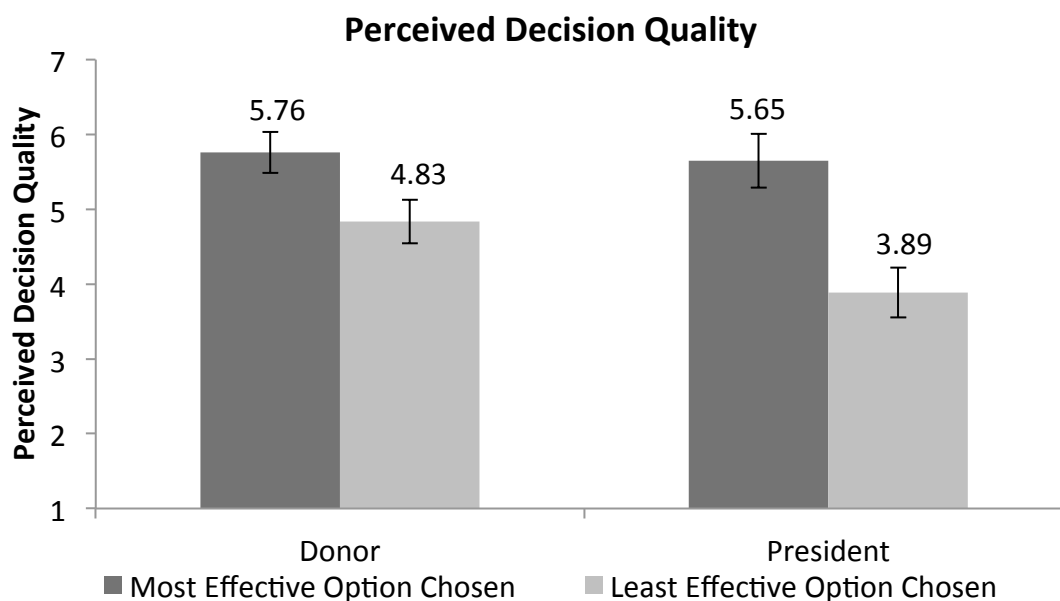


Figure 5: Participants made more negative evaluations of a decision to allocate funds to the least effective option when evaluating the President of a medical center than when evaluating a donor to the medical center. Error bars represent 95% confidence intervals.

Discussion

This study shows that people believe others have made a worse decision when they choose an option that is relatively ineffective than one that is relatively effective. Regardless of Mr. Peterson's role, he was deemed to make a poorer decision when he donated to an ineffective charity. We further find that people judge those who choose a less effective option more harshly when those individuals are accountable for their actions. In other words, people give voluntary donors more leeway to help in the way that they personally prefer rather than to help the charity that will do the most good.

GENERAL DISCUSSION

The studies reported here show that individuals generally view utilizing personal tastes to be normatively appropriate when choosing among charities, even when there exists a transparently more effective option. When effectiveness information is available, people often discount the information and select an option that they personally prefer. We find that individuals are less likely to weigh effectiveness information for charities compared to investments, and that increasing the diversity of the options reduces likelihood of selecting an option that maximizes effectiveness. We additionally find that individuals are more likely to use effectiveness information when they are in a position of responsibility to those affected, and that those who are in a position of responsibility are also judged more harshly for selecting an option that does not maximize effectiveness.

Evidence for “Warm Glow” Giving

These findings are consistent with models of “warm glow” giving, which argue that individuals gain utility from performing generous acts, and are less sensitive to the utility created by those acts (Andreoni, 1988, 1989, 1990). As a result, donations are frequently allocated to the causes that are emotionally gratifying, as decision makers seek to maximize positive feelings resulting from doing good deeds at a minimal cost to the self. Indeed, the results from Study 1 suggest that people view it to be appropriate to donate based on one’s emotional feelings above and beyond considerations of what will do the most good.

Donating to charities can be emotionally gratifying for a number of reasons. Individuals may derive pleasure from donating to causes that provide material and reputational benefits (e.g., supporting public goods they benefit from; donating to obtain a building’s naming rights;

Lacetera & Macis, 2010), that satisfy emotional states (e.g., relieving distress from being exposed to an emotionally evocative cause; Cialdini, et al., 1997), or in order to uphold a value or principle other than minimizing harm (e.g., out of a sense of duty or loyalty towards helping in-group members; Graham, Haidt & Nosek, 2009). Each of these factors can cause individuals to deemphasize effectiveness when making charitable decision.

Our studies further suggest that warm glow giving may be particularly prevalent because donors are not held personally responsible for the impact of their giving. One proposed solution, therefore, is to make donations more public and transparent in hopes to create a “culture of giving” that rewards those who donate high amounts to effective causes (Kraft-Todd et al., 2015; Singer, 2009). In Study 5, we found that donors who chose to support an ineffective charity were seen as making a less altruistic decision than those who chose to support an effective charity, suggesting that making charitable donations and effectiveness ratings public could provide reputational incentives for people to donate towards highly effective causes. However, one caveat to this finding is that we did not supply a reason for why the individual chose a poorly rated option. It is possible that if the donor said he chose the charity in order to support research for a disease that harmed a loved one, individuals may view those who selected a low-effectiveness option to be just as altruistic—or more so—as those who chose a highly-rated one (but see Lin-Healy & Small, 2012). Study 1 suggests that simply stating that an individual has an emotional connection with a cause may be enough to justify donating to a less effective option.

Moreover, a culture of giving requires that individuals be open and transparent about their donations. However, donations tends to be private, and individuals who voluntarily publicize their own good deeds are often seen as motivated by reputational concerns, which can

backfire (Berman et al., 2015). While creating a culture of giving has the potential to improve charitable outcomes, barriers exist to making this a reality.

Implications for “Effective Altruism”

Overall, our findings suggest that since people view their subjective preferences to be valid when donating to charity, the impact of providing effectiveness ratings to donors is likely to be muted. However, this does not mean that the effective altruism movement will be itself “ineffective.” Rather, we find that some people do utilize effectiveness ratings when making a decision—across our studies, the number of participants who select the value-maximizing option was significantly greater than chance.

From a utilitarian standpoint, effective altruism can be considered a success as long as the welfare benefits produced from shifting preferences exceed the costs. Central to this calculation is how the presence of effectiveness information impacts donation rates. While effectiveness information may encourage reluctant donors to give by decreasing uncertainty regarding the consequences of a donation, it is possible that this information would reduce giving if it severs people’s emotional connection to a cause and thereby inhibits them from opening their wallet to begin with (Small, Loewenstein & Slovic, 2007).

A recent study by Karlan and Wood (in press) suggests that responses to effectiveness information is heterogeneous, with some responding positively and others responding negatively. They found that when effectiveness information was added, those who donated small amounts in the past reduced their donations, whereas those who donated large amounts in the past increased their donations. The authors argue that small donors are “warm glow givers”, whereby the additional information inhibited an emotional connection towards the cause, reducing their

tendency to donate. Large donors, in contrast, increased their donations when faced with favorable information regarding the effectiveness of a donation. When combined, these two effects canceled each other out, and there was no overall change in donation rates or amounts.⁴

Future research should investigate moderators that determine how effectiveness information impacts donation rates. For instance, donors may become discouraged upon finding out that their past donations were relatively ineffective, and disengage from donating in the future. To align their beliefs with their actions, they may additionally reject the validity of effectiveness information if past contributions were revealed to be ineffective (Bem 1972; Festinger, 1962). If so, then the movement may be better targeted to the young, who are less likely to have a history of donating to charity, and thereby may be more receptive to hearing the arguments put forth by the movement.

Some individuals may also disagree with the premises that the effective altruism movement stands on. For instance, some effective altruists endorse that one should treat all human life equally (Singer 2009, 2015). However, a consequentialist calculation may dictate that saving the life of a skilled doctor—who would in turn save the lives of others—may produce greater welfare gains than saving the lives of two accountants. Individuals may also believe that certain welfare comparisons are inherently incommensurable, such as determining the relative value of reducing animal suffering versus human suffering, reducing vision loss versus hearing loss, or helping a life for sure in the present versus the uncertain possibility of helping more lives in the future. One challenge effective altruists face is to convince people that, while quantifying these tradeoffs may be imperfect and require some amount of tolerance for error, doing so can

⁴ To investigate this further, we conducted four additional studies that examine how effectiveness information impacts donation rates or exerted effort in support of a charity. In two studies we found no significant effect of the presence of effectiveness information on donation rates. In the third, we found that effectiveness information reduced overall effort, but we failed to replicate this effect in a fourth study. Further investigation is warranted to identify when effectiveness information will impact donation rates and amounts.

engender more positive outcomes than relying on intuition alone. Indeed, recent writings of effective altruists argue that at the extremes, the cost-benefit analysis of helping those in the developing world are orders of magnitude greater than helping those in the developed world, making these concerns negligible (McCaskill, 2015; Singer, 2015).

In conclusion, the present research suggests that providing effectiveness information may not go far to convince people of their value. For effective altruists, success therefore requires an understanding of how best to change behavior. Our results present two possible paths: introducing better incentives to encourage effective giving (e.g., increasing accountability) and altering the decision context to nudge people to select more effective options. To this latter point, as of the writing of this article, on the website for TheLifeYouCanSave.Org, individuals can search for effective charities first by selecting what cause they would like to help before looking at the options within each type of cause. The site has also cleverly assigned a few charities to multiple categories. For instance, those who click on “humanitarian aid”, “children” or “women and girls” will all be informed about Against Malaria which can fit in any of these categories. This way, individuals can match their preferences with a category and still be directed towards donating to highly effective charities. Future research can identify other avenues that encourage individuals to rely less on their tastes and intuitions and more on comparing outcomes when making charitable decisions.

REFERENCES

- Against Malaria (2016) *Donation Statistics*. Accessed from:
<https://www.againstmalaria.com/DonationStatistics.aspx>
- Andreoni, J. (1988). Privately provided public goods in a large economy: The limits of altruism. *Journal of Public Economics*, 35(1), 57-73.
- Andreoni, J. (1989). Giving with impure altruism: Applications to charity and Ricardian equivalence. *The Journal of Political Economy*, 1447-1458.
- Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal*, 100(401), 464-477.
- Barasch, A., Levine, E. E., Berman, J. Z., & Small, D. A. (2014). Selfish or selfless? On the signal value of emotion in altruistic behavior. *Journal of Personality and Social Psychology*, 107(3), 393-413.
- Baron, J. (1993). *Morality and Rational Choice* (Vol. 18). Springer Science & Business Media.
- Baron, J. (2005). *Thinking and Deciding*. Cambridge, MA: Cambridge University Press.
- Baron, J., & Szymanska, E. (2011). Heuristics and biases in charity, In D. M. Oppenheimer & C. Y. Olivola (Eds.), *The Science of Giving: Experimental Approaches to the Study of Charity*. (pp. 215-235). New York, NY: Psychology Press; US.
- Bem, D. J. (1972). Self-perception theory. *Advances in Experimental Social Psychology*, 6, 1-62.
- Berman, J. Z., Levine, E. E., Barasch, A., & Small, D. A. (2015). The braggart's dilemma: On the social rewards and penalties of advertising prosocial behavior. *Journal of Marketing Research*, 52(1), 90-104.
- Brehm, J. W. (1966). *A Theory of Psychological Reactance*. New York: Academic Press.

- Caviola, L., Faulmüller, N., Everett, J. A., Savulescu, J., & Kahane, G. (2014). The evaluability bias in charitable giving: Saving administration costs or saving lives? *Judgment and Decision Making*, 9(4), 303-315.
- Cialdini, R. B., Brown, S. L., Lewis, B. P., Luce, C., & Neuberg, S. L. (1997). Reinterpreting the empathy-altruism relationship: When one into one equals oneness. *Journal of Personality and Social Psychology*, 73, 481–494.
- Festinger, L. (1962). *A Theory of Cognitive Dissonance (Vol. 2)*. Stanford University Press.
- Fetherstonhaugh, D., Slovic, P., Johnson, S., & Friedrich, J. (1997). Insensitivity to the value of human life: A study of psychophysical numbing. *Journal of Risk and Uncertainty*, 14(3), 283-300.
- Gneezy, U., Keenan, E. A., & Gneezy, A. (2014). Avoiding overhead aversion in charity. *Science*, 346(6209), 632-635.
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96(5), 1029-1046.
- Harbaugh, W. T., Mayr, U., & Burghart, D. R. (2007). Neural responses to taxation and voluntary giving reveal motives for charitable donations. *Science*, 316(5831), 1622-1625.
- Hayes, A. F. (2013). *Introduction to Mediation, Moderation, and Conditional Process Analysis: A Regression-Based Approach*. Guilford Press.
- Hsee, C. K. (1996). The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives. *Organizational Behavior and Human Decision Processes*, 67, 247–257.
- Imas, A. (2014). Working for the “warm glow”: On the benefits and limits of prosocial incentives. *Journal of Public Economics*, 114, 14-18.

- Inbar, Y., Cone, J., & Gilovich, T. (2010). People's intuitions about intuitive insight and intuitive choice. *Journal of Personality and Social Psychology*, 99(2), 232-247.
- Karlan, D., & Wood, D. H. (in press). The effect of effectiveness: Donor response to aid effectiveness in a direct mail fundraising experiment. *Journal of Behavioral and Experimental Economics*
- Kogut, T., & Ritov, I. (2005). The singularity effect of identified victims in separate and joint evaluations. *Organizational Behavior and Human Decision Processes*, 97(2), 106-116.
- Kraft-Todd, G., Yoeli, E., Bhanot, S., & Rand, D. (2015). Promoting cooperation in the field. *Current Opinion in Behavioral Sciences*, 3, 96-101.
- Lacetera, N., & Macis, M. (2010). Social image concerns and prosocial behavior: Field evidence from a nonlinear incentive scheme. *Journal of Economic Behavior & Organization*, 76(2), 225-237.
- Lin-Healy, F., & Small, D. A. (2012). Cheapened altruism: Discounting personally affected prosocial actors. *Organizational Behavior and Human Decision Processes*, 117(2), 269-274.
- Loewenstein, G., & Small, D. A. (2007). The Scarecrow and the Tin Man: The vicissitudes of human sympathy and caring. *Review of General Psychology*, 11(2), 112-126.
- MacAskill, W. (2015). *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference*. London: Guardian Faber Publishing.
- Make-A-Wish (2016). *FY2015 Make-A-Wish Annual Report*. Downloaded from:
<http://wish.org/about-us/making-a-difference/managing-funds>
- Ord, T. (2012). The moral imperative toward cost- effectiveness in global health. *Center for Global Development*.

Singer, P. (1979). *Practical Ethics*. Cambridge: Cambridge University Press.

Singer, P. (2009). *The Life You Can Save: Acting Now to End World Poverty*. New York: Random House.

Singer, P. (2015). *The Most Good You Can Do: How Effective Altruism is Changing Ideas About Living Ethically*. New Haven: Yale University Press.

Slovic, P. (2007). “If I look at the mass I will never act”: Psychic numbing and genocide. *Judgment and Decision Making*, 2(2), 79–95.

Small, D. A. (2010). Reference-dependent sympathy. *Organizational Behavior and Human Decision Processes*, 112(2), 151-160.

Small, D. A., Loewenstein, G., & Slovic, P. (2007). Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organizational Behavior and Human Decision Processes*, 102(2), 143-153.

Appendix A: Study 1 stimuli

Distant condition in brackets

Recently, Mary received a bonus from her company and is looking to donate a portion of this bonus to charity. Mary has received brochures for two charities and is considering which one to donate to.

The first charity, "Jump Start Your Community", trains and educates local homeless so they can find work. Jump Start invests most of its dollars to hiring educators and trainers. For every \$100 donated, they can provide 2 hours of job training to one homeless person.

The second charity, "Hunger Care in Africa", provides micronutrients for hungry children in Africa. For each \$100 donated, they can provide enough food to feed 5 children for a month.

Mary has always felt a strong emotional pull towards helping those in her local community, and cares deeply about helping those close to home. [Mary has always felt a strong emotional pull towards helping those in third world countries, and cares deeply about helping those around the world.]

Appendix B: Studies 2 stimuli

Alternative Energy

- EcoLogic Biofuel
- Spectral Solar
- Vector Wind Energy

Cancer

- Certec BioMedical
- The Cancer Research Institute
- United Cancer Services [Inc.]

Education

- Knowledge Tech.
- Learning Tree
- Teacher Network Services [Inc.]

Food:

- Arcadia Grocers
- Harvest Distribution Services [Inc.]
- The Fresh Water Group [Company]

Housing:

- Franklin Residences
- Main Street Housing
- Urban Development Agency [Company]

International:

- Action International
- Global Resurgence Agency [Company]
- Navara Worldwide

Note: Investment condition in brackets

Sample Screenshot from the “Search by Category” Interactive Table

Alternative Energy
Cancer
Education
<ul style="list-style-type: none"><input type="radio"/> Knowledge Tech Category = Education Rating = 96<input type="radio"/> Teacher Network Services Category = Education Rating = 90<input type="radio"/> Learning Tree Category = Education Rating = 84
Food
Housing
International

Select Item

Participants click on a heading to reveal the options in that heading. In this example, clicking on the Education heading reveals the three options. Within each heading, options are ordered by effectiveness rating.

Appendix C - Study 3a stimuli

Possible Charity Options

Blindness Charities

- Guide Dogs for the Blind: Trains Seeing Eye dogs to help the blind
- American Foundation for the Blind: Provides services to help the blind
- *Fred Hollows Foundation: Provides low cost cataract surgery to prevent blindness

Children's Health Charities

- St. Jude's Children's Hospitals: A leading children's hospital focusing on treatment for children with life-threatening disease
- Make-A-Wish Foundation: Grants wishes to children suffering from life-threatening diseases
- *Schistosomiasis Control Initiative: Provides children in Africa with medicine to combat parasitic diseases

Water Charities

- The Water Project: Provides water infrastructure to villages in Africa
- Charity: Water: Provides clean water to people around in third-world countries
- *Evidence Action: Provides low-cost clean water to people in Africa

*Indicates charity has been designed to be highly effective by GiveWell.org or TheLifeYouCanSave.org

Sample Screenshot from the Mixed-Category Condition

Below are the list of charities you can choose from:

Name	Type	Description	Rated as Highly Effective?
Guide Dogs for the Blind	Blindness	Trains seeing eye dogs to help the blind	
The Water Project	Water	Provides water infrastructure to villages in Africa	
Schistosomiasis Control Initiative	Children's Health	Provides children in Africa with medicine to combat parasitic diseases	✓

Appendix D - Study 3b stimuli**Food & Water Charities:**

- The Fresh Water Agency
- Harvest Distribution
- Arcadia Nutrition

Health & Medicine Charities:

- Navara Health
- Medical Goods Initiative
- Heart Aid

International Charities:

- Action International
- International Aid Fund
- Global Revival