

LBS Research Online

J Z Berman, A Barasch, E E Levine and D A Small

Impediments to effective altruism: the role of subjective preferences in charitable giving

Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/958/>

Berman, J Z, Barasch, A, Levine, E E and Small, D A

(2018)

Impediments to effective altruism: the role of subjective preferences in charitable giving.

Psychological Science, 29 (5). pp. 834-844. ISSN 0956-7976

DOI: <https://doi.org/10.1177/0956797617747648>

SAGE Publications (UK and US)

<http://journals.sagepub.com/doi/10.1177/0956797617...>

Users may download and/or print one copy of any article(s) in LBS Research Online for purposes of research and/or private study. Further distribution of the material, or use for any commercial gain, is not permitted.

This is the accepted version of the following article:

Jonathan Z Berman, Alixandra Barasch, Emma E Levine, and Deborah A Small, Impediments to Effective Altruism: The Role of Subjective Preferences in Charitable Giving, Psychological Science . Copyright © 2018 (Copyright Holder). Reprinted by permission of SAGE Publications.

<https://doi.org/10.1177/0956797617747648>

Impediments to Effective Altruism:
The Role of Subjective Preferences in Charitable Giving

Jonathan Z. Berman

Alixandra Barasch

Emma E. Levine

Deborah A. Small

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. Corresponding Author: Jonathan Berman
(jberman@london.edu).

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, the “Effective Altruism” movement seeks to provide individuals with information regarding the effectiveness of charities in hopes that they will donate to organizations that maximize the social return of their donation. This paper investigates the extent to which presenting effectiveness information leads people to choose more effective charities. We find that even when effectiveness information is made easily comparable across options, it has a limited impact on choice. Specifically, people frequently choose less effective charity options when those options represent more subjectively preferred causes. In contrast to making a personal donation decision, outcome metrics are used to a much greater extent when choosing financial investments and when allocating aid resources as an agent of an organization. Implications for Effective Altruism are discussed.

Charity could do the most good if donations went to the causes that produced the greatest welfare gains. The “Effective Altruism” (EA) movement seeks to make this proposition a reality. EA 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). Thus, just as an investor aims to maximize the financial return of each dollar invested, donors should aim to maximize the social welfare returns of each dollar donated.

In recent years, EA has given rise to organizations, such as GiveWell.org, that provide information to the public about charities’ effectiveness. However, it is unclear whether the provision of this information will have a demonstrable impact on choice.

While normative models of altruism argue that individuals should allocate their limited resources to funds that maximize welfare (Baron, 1993, 2005; Ord, 2012; Singer, 1979, 2009, 2015), descriptive accounts of giving show that individuals often do not maximize the outcomes of their donations. People tend to donate more to single identifiable victims than to large groups of victims (Kogut & Ritov, 2005a, 2005b; Small, Loewenstein & Slovic, 2007); pay more attention to the ratio of victims rather than the absolute number of victims (Fetherstonhaugh et al., 1997); and over-weight overhead expenses relative to outcomes achieved (Baron & Szymanska, 2011; Caviola, et al., 2014; Gneezy, Keenan & Gneezy, 2014). Whereas those findings persist in between-subjects designs (i.e., separate evaluation), when outcome metrics are made directly comparable (i.e., joint evaluation, Hsee, 1996), individuals dramatically increase their selection of the welfare-maximizing option (e.g., Caviola et al., 2014; Kogut & Ritov, 2005a). Thus, people appear to be “distorted altruists”—they care about welfare maximization,

but without clear information to make comparisons, they rely on their feelings to guide choice (Loewenstein & Small, 2007; Slovic, 2007).

This past research suggests that people will welfare maximize so long as the decision environment affords comparisons of effectiveness across options. However, one limitation of existing studies is that participants typically evaluate donation options from a single underlying cause, such as choosing between saving one child or eight children suffering from the same disease (e.g., Kogut & Ritov, 2005a). In the real world, individuals must choose amongst charities that support various types of causes, such as choosing between saving children from illness or reducing hunger. Yet, we know little about how people utilize effectiveness information when causes differ.

The present research investigates how individuals make donation decisions when effectiveness information is provided across multiple different causes, making it easy for people to maximize welfare if they desire to do so. We find that even when effectiveness information is comparable, a substantial portion choose not to welfare-maximize. Our results suggest that individuals view charity as a relatively subjective decision (cf. Goodwin & Darley, 2008; Spiller & Belogolova, 2017), one in which people often feel justified to discount welfare maximization concerns in order to choose in accordance with their personal preferences.

If people view charity solely in terms of welfare maximization, then the idiosyncratic preferences of decision-makers should be irrelevant to their choice, and people will strive to choose the option that does the most good. However, if people view charity as a relatively subjective decision, then the personal preferences of decision-makers will be considered valid inputs to choice, in spite of their consequences for welfare maximization. In other words, believing that charity is a subjective decision licenses individuals to donate in personally

gratifying ways at the cost of maximizing welfare. Thus, evaluating how people construe charitable giving is of utmost importance to understanding the impact of effectiveness information on choice.

The following studies investigate beliefs regarding the subjectivity of charitable giving and its consequences for social welfare maximization. Study 1 shows that, relative to many personal decisions, individuals view charity as being relatively subjective. Study 2 shows that individuals prioritize emotional connection over welfare maximization when comparing charity options. Study 3 examines choice processes and finds that people are more likely to override welfare maximization when choosing a charity than when choosing a financial investment. The final two studies find that individuals are less likely to license themselves (Study 4) and others (Study 5) to select an ineffective option when a decision-maker assumes a position of responsibility.

STUDY 1: PERCEIVED SUBJECTIVITY OF CHARITY

We conducted an initial survey to investigate perceived subjectivity of a variety of common decisions, including charity decisions.

Method

We recruited 126 participants (55.6% female; 44.4% male; Mean age = 25.0) from a Northeastern University's behavioral lab pool. Participants were asked to evaluate how they believe they should make decisions across six decision domains (choosing a charity, investment, cell phone, restaurant, a piece of art, and a 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 = “not at all” to 7 = “very much so.” We reverse-coded responses from the latter two statement and then averaged all items to create a three-item measure of perceived subjectivity for each decision domain.

In this and all subsequent studies, sample sizes were determined in advance. For the lab studies, sample size was determined by the number of participants who appeared for the assigned lab session. For online studies, we recruited a minimum of 100 participants per cell. No participants or conditions were dropped from any analysis performed. All measures not reported herein can be found in the supporting materials. Data files are stored at: <https://osf.io/6myfj/>.

Results

We first compared the three-item perceived subjectivity measure for each decision domain to the charity condition. Participants indicated that it is more important to rely on one’s subjective preferences, rather than objective metrics, when choosing a charity ($M = 4.85$, $SD = 1.02$) compared to choosing a cell phone ($M = 3.88$, $SD = 1.03$; $t(125) = 9.03$, $p < .001$, $d = .81$ (95% $CI_d = [.56, 1.05]$), an investment ($M = 3.74$, $SD = 0.64$; $t(125) = 10.57$, $p < .001$, $d = .97$ (95% $CI_d = [.64, 1.30]$), or a medical treatment ($M = 3.42$, $SD = 0.54$; $t(125) = 13.34$, $p < .001$, $d = 1.23$ (95% $CI_d = [.81, 1.66]$). The importance of relying on subjective preferences for a charity decision was roughly on par with choosing a restaurant ($M = 4.74$, $SD = 0.96$; $t(125) = 1.02$, $p = .31$, $d = .10$ (95% $CI_d = [-.11, .30]$). Only choosing a piece of art was evaluated to be more a matter of subjective preference than choosing a charity ($M = 5.68$, $SD = 1.03$; $t(125) = -7.49$, $p < .001$, $d = -.67$ (95% $CI_d = [-.91, -.42]$). Table 1 displays the results broken down by each item.

	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 the given decision is rated as being significantly more objective than charitable decisions, whereas *italics* indicate a given decision is rated as being significantly more subjective than charitable decisions. Astericks correspond to the significance level when comparing to the charity condition. ** p < .01, *** p < .001

Table 1: Subjective preferences are considered more valid when making charity decisions than when choosing a cell phone, investment, or medical treatment.

Discussion

This survey shows that people think of charity as a relatively personal, or subjective decision, and not necessarily one that should be made using objective measures. Even so, it is possible that people will behave differently in a decision environment when comparative effectiveness information is transparent. The following studies examine whether judgments and choice conform to these meta-preferences.

STUDY 2: PERSONAL FEELINGS VERSUS WELFARE GAINS

Study 2 examines whether people believe it is normatively appropriate to discount welfare-maximizing options in order to choose a more subjectively-preferred charity. Whereas past work suggests people aim to welfare maximize, even if they make mistakes when doing so (Loewenstein & Small, 2007; Slovic, 2007), this research has not examined people's beliefs when both causes and welfare gains differ. We expect that people readily identify that some charities do more good than others. However, we don't expect individuals will endorse donating to the more effective charity. Consistent with research showing that one's personal feelings are

normatively valued in prosocial behavior (Barasch et al., 2014), we expect that people will endorse donating to causes that people feel connected to rather than causes that do the most good.

Method

We pre-registered all aspects of this study on AsPredicted.org (<https://aspredicted.org/pe8j2.pdf>). 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 supporting materials for full 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 with distant charities, they said that she should donate to “Hunger Care For Africa” ($M = 5.26, SD = 2.05$) and also evaluated it as being more effective ($M = 5.59, SD = 1.87$), $t(197) = -1.19, p = .24, d = -0.17$ (95% $CI_d = [-.44, .11]$). However, when she felt connected to local communities, they indicated that she should donate to “Jump Start Your Community” ($M = 3.00, SD = 1.99$), despite indicating that “Hunger Care For Africa” was more effective ($M = 4.55, SD = 2.32$), $t(202) = -5.12, p < .001, d = -0.72$ (95% $CI_d = [-1.00, -.43]$). Figure 1 displays these results.

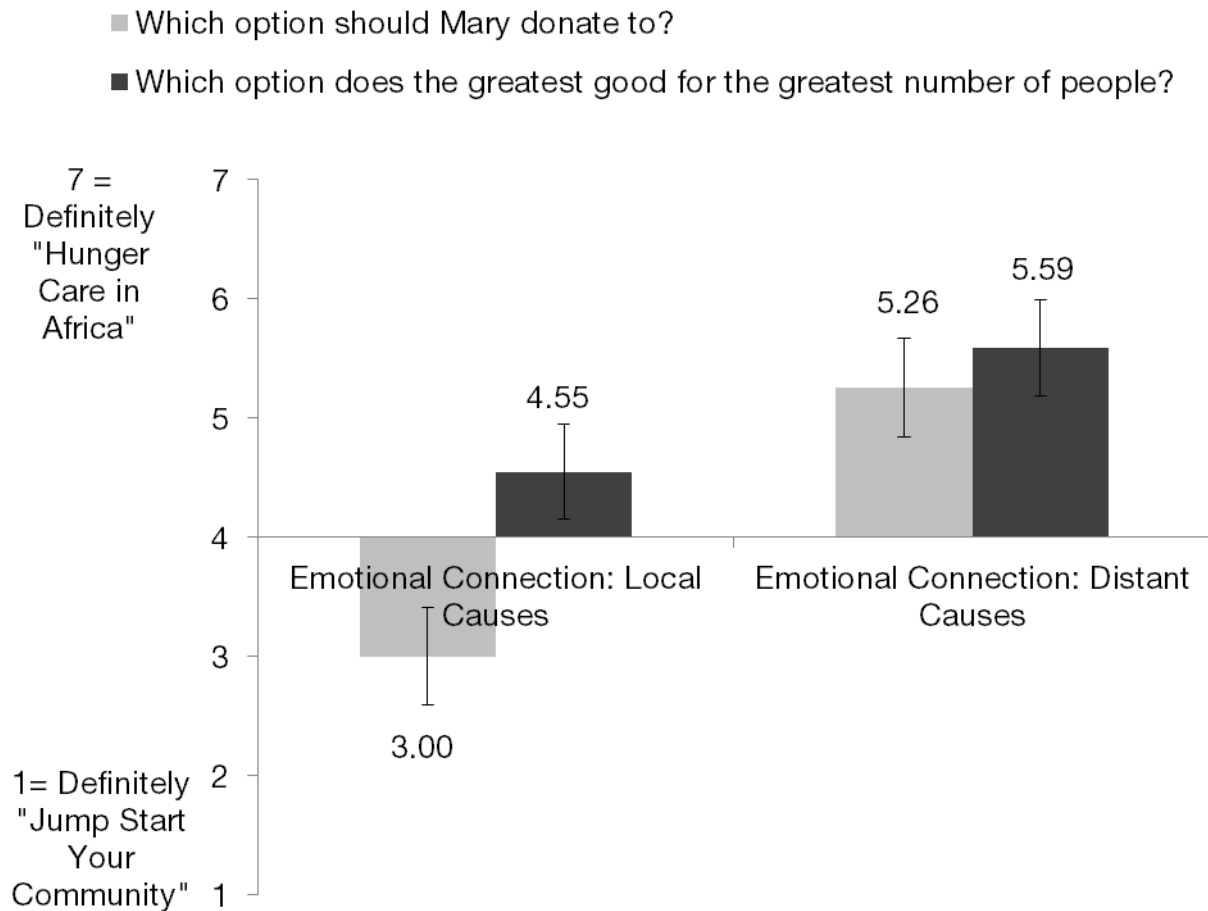


Figure 1. Participants believe that Mary should donate to the cause for which she feels an emotional connection, even if it does not maximize social welfare. Error bars represent 95% confidence intervals.

STUDY 3: CHARITY VERSUS INVESTMENT CHOICE

Effective altruists argue that in order to do the most good, individuals should maximize expected welfare just as investors should seek to maximize financial returns. However, in Study 1, we found that individuals more strongly endorse utilizing objective information in financial decisions compared to charitable decisions. This study directly contrasts choices between

charities and financial investments, and further evaluates potential explanations for how people choose amongst options.

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 option type.

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.

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 supporting materials for the 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. We expected that people will be more likely to weigh their personal preferences when making charitable decisions. To examine this, participants evaluated four statements and rated their agreement of each on scales ranging from 1 = “strongly disagree” to 7 = “strongly agree.”: “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) and Sort Choice (Category Type vs. Effectiveness Ratings) predicting the rank of the chosen option, showed that participants chose a less effective Charity (Mean Rank = 4.11, $SD = 4.00$) than Investment (Mean Rank = 2.78, $SD = 2.83$), $\beta = .57, SE = .19, \chi^2(1) = 9.36, p = .002, OR = 1.77$ (95% $CI_{OR} = [1.22, 2.57]$), and that participants who searched by category type chose a less effective option (Mean Rank = 6.07, $SD = 4.58$) than those who searched by effectiveness ratings (Mean Rank = 2.60, $SD = 2.57$), $\beta = 1.71, SE = .22, \chi^2(1) = 59.39, p < .001, OR = 5.52$ (95% $CI_{OR} = [3.59, 8.51]$). Figure 2 displays these results.




Figure 2. Mean rank of the selected option by condition and sort choice. Lower numbers (ranks closer to one) indicate more effective options. Error bars represent 95% CIs.

Additional Measures: Participants indicated that they weighed their subjective preferences more heavily in the Charity condition ($M = 4.51$, $SD = 0.87$) than in the Investment condition ($M = 3.54$, $SD = 1.03$), $F(1, 399) = 104.42$, $p < .001$, $d = 1.02$ (95% $CI_d = [.81, 1.22]$). Participants also preferred to have more agency when making charitable decisions ($M = 5.66$, $SD = 1.23$) than when making investment decisions ($M = 4.81$, $SD = 1.29$), $F(1, 399) = 46.13$, $p < .001$, $d = .68$ (95% $CI_d = [.47, .88]$). There was no statistically significant difference between subjective knowledge for charitable decisions ($M = 3.87$, $SD = 1.51$) and investment decisions ($M = 3.62$, $SD = 1.59$), $F(1, 399) = 2.68$, $p = .10$, $d = .16$ (95% $CI_d = [-.04, .36]$). Finally, participants believed that it is more possible to objectively rate the charities ($M = 5.39$, $SD = 1.23$) than investments ($M = 5.07$, $SD = 1.19$), $F(1, 399) = 7.27$, $p = .007$, $d = .27$ (95% $CI_d =$

[.07, .46]). This suggests that the observed effects are not due to a belief that charities are less objectively evaluable.

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 = 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.

Robustness Check Studies: To further test the robustness of these results, we ran an additional study that directly assigned participants to the sort by category presentation format ($N = 224$; Supporting Materials, Study S1). We again found that fewer participants chose the highest rated charity (22.0%) than the highest rated investment (51.3%), $\chi^2(1, N = 224) = 20.58$, $p < .001$; $\phi = .30$. It is also possible that participants are more likely to satisfice when faced with

large choice sets for charities than investments, and these results would not replicate in small choice sets when it is easy to compare all options. Yet when we ran a study that presented participants ($N = 201$, Supporting Materials, Study S2) with choice sets containing only three options, we still find that significantly fewer chose the highest rated charity option (46%) than the highest rated investment (64%), $\chi^2(1, N = 201) = 6.91, p = .009, \phi = .19$.

We also examined how sensitive these findings are to the range of effectiveness ratings in a choice set. In an additional study ($N = 403$, Supporting Materials, Study S3), participants were presented with either five charity or investment options. We then randomly assigned options to have effectiveness ratings that were either between 75 and 100 (Narrow Range condition) or between 0 and 100 (Wide Range condition). A binary logistic regression found a main effect of decision type, such that participants were more likely to choose the highest rated investment than the highest rated charity, $\beta = .65, SE = .21, \chi^2(1) = 9.21, p = .002, OR = 1.92$ 95% CI = [1.27, 2.89], and a main effect of range, such that participants were more likely to maximize on effectiveness for wide ranges rather than narrow ranges, $\beta = .44, SE = .21, \chi^2(1) = 4.20, p = .04, OR = 1.55$, 95% CI = [1.02, 2.35]. A separate regression confirmed that there was no significant interaction ($p = .42$). Thus, even when choice sets contain large differences in effectiveness ratings, people still remain more hesitant to maximize social welfare than financial returns. Additionally, we find some evidence that the more substantial the difference in effectiveness ratings of options, the more likely people will weigh effectiveness information in choice.

STUDY 4: DECISION-MAKING ROLE AND WELFARE MAXIMIZATION

In Study 4, we investigate when individuals are more likely to welfare-maximize when helping others. One reason why individuals may discount effectiveness information in choice is

that they do not experience the consequences for choosing an ineffective option, and are thus rendered insensitive to the impact of actions on others (Imas, 2014). However, we expect that when individuals assume a role of responsibility, they may feel obligated to act in accordance of the welfare of the entire group, and will discount their personal preferences in order to do so.

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 welfare-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, $\beta = .05, SE = .27, \chi^2(1) = 0.03, p = .87, OR = 1.05$ (95% $CI_{OR} = [0.62, 1.78]$). 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, $\beta = .79, SE = .26, \chi^2(1) = 9.70, p = .002, OR = 2.20$ (95% $CI_{OR} = [1.32, 3.67]$). While participants utilized effectiveness ratings when given this information, significantly more did so when placed in a position of responsibility. Figure 3 displays these results.

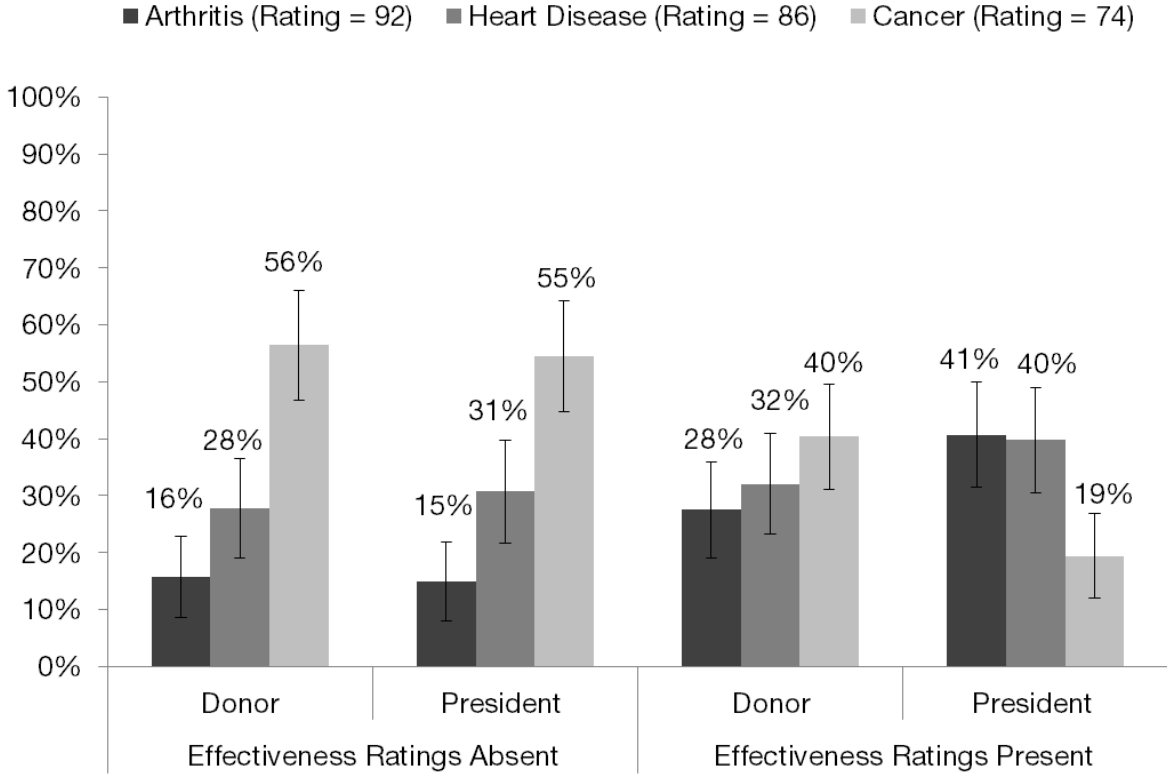


Figure 3: Percent of participants choosing to fund a department as a function of their role (Donor vs. President) and whether effectiveness information was provided. Ratings were only known to those in the Effectiveness Ratings Present condition. Error bars indicate 95% CIs.

STUDY 5: JUDGMENTS OF DECISION QUALITY

This study examines how individuals judge the decisions of others who select either a relatively effective or ineffective option. We expect that people will acknowledge that selecting a highly effective option is indicative of making a higher quality decision than a less effective option. However, we further examine if this judgment depends on whether the decision-maker assumes a role of responsibility. We expect that those in a role of responsibility will be viewed particularly harshly for choosing to allocate funds to an ineffective option. In contrast, we expect more leeway will be given to individual donors who select an option they personally 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. $M_{\text{LeastEffectiveOption}} = 3.89, SD = 1.35$), $t(114) = 7.64, p < .001, d = 1.46$ (95% $CI_d = [.21, 1.86]$), than when he was a donor ($M_{\text{MostEffectiveOption}} = 5.56, SD = 1.09$ vs. $M_{\text{LeastEffectiveOption}} = 4.83, SD = 1.30$), $t(109) = 4.09, p < .001, d = .61$ (95% $CI_d = [.22, 1.00]$).

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. $M_{\text{LeastEffectiveOption}} = 4.30, SD = 1.09$), $t(114) = 7.85, p < .001, d = 1.49$ (95% $CI_d = [1.06, 1.90]$), than when he was a donor ($M_{\text{MostEffectiveOption}} = 5.78, SD = 0.91$ vs. $M_{\text{LeastEffectiveOption}} = 5.19, SD = 1.00$), $t(109) = 3.30, p = .001, d = .62$ (95% $CI_d = [.22, 1.00]$).

GENERAL DISCUSSION

This paper shows that many view relying on subjective preferences to be normatively appropriate when choosing where to donate, even when there are transparently more effective options available to donors. Thus, these results provide an important caveat to the “distorted altruist” view of charitable giving, which implies that failures to maximize welfare are due to systematic mistakes that are corrected in the right information environment—when people can compare the impact of different charities (see Small, 2010). In contrast, the current research shows that the benefits of comparing charity effectiveness are limited when causes vary by type, as people often believe it is more important to choose an option they emotionally connect with, rather than an option that does the most good.

In this sense, these findings are more consistent with the theory of “warm-glow” giving which argues that individuals gain utility from committing instances of generous acts, but are insensitive to the benefits created by the acts (Andreoni, 1990). Our results suggest that one reason individuals are insensitive to the magnitude of benefits is due to how they construe charity; believing it is a relatively subjective decision, and not one that should be made solely by consulting the numbers.

One limitation of the current research is that we rely primarily on hypothetical scenarios (except, see supporting materials, Study S4). However, there is strong reason to believe that these results will hold in real world behavior. For one, the stimuli we utilize are affect-poor, which tend to induce calculative mindsets when assessing value (Hsee & Rottenstreich, 2004). As stimuli become more vivid—as is often the case in the field—people may rely on their feelings and intuitions to a greater extent, favoring more emotionally evocative options at the expense of welfare maximization (Loewenstein & Small, 2007). It is also possible that with

earned money, people feel even more licensed to choose a cause they subjectively prefer (cf. Cherry, Frykblom & Shogren, 2002).

While our findings suggest that there will be little impact of providing effectiveness ratings to donors, this does not mean that EA will fail. Across studies, the number selecting a welfare-maximizing option was significantly greater than chance. However, we expect that the extent to which people utilize effectiveness information will be contingent on the diversity of options faced: the more diverse the choice set, the more likely individuals are to find and choose an option they personally prefer. In two studies reported in the supporting materials (Studies S4 & S5), participants were significantly more likely to welfare-maximize when choice sets contained similar (vs. differentiated) causes. When choosing between causes, people appear more willing to discount welfare considerations and choose in line with their personal preferences, but once a cause has been identified, welfare considerations become increasingly paramount. Additional factors may further affect people's tendency to utilize effectiveness information when comparing options. For instance, drawing explicit attention to the opportunity costs associated with donating to less effective charities may encourage individuals to make more welfare-maximizing choices.

More broadly, EA can be considered a success as long as the welfare gains produced exceed the costs. Central to this calculation is how information impacts donation rates. While effectiveness information may encourage reluctant donors to give by decreasing uncertainty regarding the consequences of a donation, it may also reduce giving if it severs emotional connections to causes, inhibiting action (Karlan & Wood, 2017; Small, et al., 2007). Individuals may additionally become discouraged if past donations are revealed to be ineffective, which could cause them to disengage from donating in the future.

Some individuals may additionally disagree with the premise that EA stands on, which presumes that donors should treat all human life equally (Singer, 2009, 2015). However, a consequentialist calculation could dictate that saving the life of a skilled doctor would produce greater welfare gains than saving the lives of two accountants. Moreover, certain welfare comparisons may be considered incommensurable, such as determining the relative value of reducing animal suffering versus human suffering or helping a life for sure versus the uncertain possibility of helping more. 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 engender more positive outcomes than permitting feelings to drive choices. Some 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 many of these concerns negligible (McCaskill, 2015; Singer, 2015).

In conclusion, our results suggest that people view charity decisions as being relatively subjective, which inhibits the impact of effectiveness information on welfare maximization. Thus, to persuade people to make donation decisions that maximize social welfare, providing information alone may not be sufficient. Rather, it may require altering how individuals view their role as a donor altogether.

REFERENCES

- 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.
- 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.
- Cherry, T. L., Frykblom, P., & Shogren, J. F. (2002). Hardnose the dictator. *The American Economic Review*, 92(4), 1218-1221.
- 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.

- Goodwin, G.P., & Darley, J.M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106, 1339-1366.
- 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.
- Hsee, C. K., & Rottenstreich, Y. (2004). Music, pandas, and muggers: on the affective psychology of value. *Journal of Experimental Psychology: General*, 133(1), 23-30.
- 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. (2017). The effect of effectiveness: Donor response to aid effectiveness in a direct mail fundraising experiment. *Journal of Behavioral and Experimental Economics*, 66, 1-8.
- Kogut, T., & Ritov, I. (2005a). The singularity effect of identified victims in separate and joint evaluations. *Organizational Behavior and Human Decision Processes*, 97(2), 106-116.
- Kogut, T., & Ritov, I. (2005b). The “identified victim” effect: An identified group, or just a single individual? *Journal of Behavioral Decision Making*, 18(3), 157-167.

- 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.
- 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, 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.
- Spiller, S. A., & Belogolova, L. (2017). On consumer beliefs about quality and taste. *Journal of Consumer Research*, 43 (6), 970-991.