Supplementary Information for

Donors Vastly Underestimate Differences in Charities'
Effectiveness

Supplementary Text on Method and Results of Study 4

The post-experimental questionnaire in this study gave the experts the opportunity to comment on their answers and asked them about their self-estimated level of expertise. We also asked about what organization they worked for (optional) but do not report this for privacy reasons. We conducted a correlation analysis between those estimates and the responses to the explicit comparison question (excluding two outliers) but found no significant correlations.

Pay-to-reveal-effectiveness study (not reported in the main manuscript)

In this study, we studied whether correcting this misconception would lead to more effective giving. One way in which correcting this misconception could increase effective giving is indirectly via an increased willingness to invest resources to find the most effective charities. In general, the larger the difference between the most effective charities and the average charity is, the more you will increase your impact if you identify, and donate to, the most effective charities. To the extent that people want to maximize their impact, this should make people more willing to invest resources in finding the most effective charities before choosing a charity to donate to. This means that underestimating the difference in effectiveness between charities could lead to underinvestment in finding the most effective charities, which in turn could lead to ineffective giving. Potentially, this could be an important mechanism linking underestimates of the difference in effectiveness between charities and ineffective giving. Our hypothesis, therefore, was that if people knew that the difference was 100 instead of 1.5, they would be more willing to invest resources to find the most effective charity.

Method
Participants. We recruited 302 US American participants from MTurk. Sixty-six were excluded because they failed an attention check question. The final sample consisted of 236 participants (123 females, age M = 40.53, SD = 12.55). Our sample size was determined by a power analysis, which suggested that we would need at least 272 participants to detect a small to medium effect of f = 0.17, taking an alpha of .05, a power of .80, and df of 1. We aimed to recruit 300 participants to account for exclusions. This study was pre-registered at https://osf.io/dz243/. Since a relatively large number was excluded due to failing the attention check question, we conducted the analysis with and without any exclusions. The findings remained the same and are reported in the supplementary materials.
Procedure. Participants were divided into two conditions (1.5 vs. 100 ratio) and were informed that "in a recent report, researchers concluded that the most cost-effective charities are [1.5/100] times more effective than a charity of average level of cost-effectiveness". Next, participants were told that they would be asked which out of five named charities they would donate $100 to. We used names of actual (non-fictional) charities and donated in accordance with the participants' choices after the study. (After the study was completed, we donated a total amount of $100 in proportion to how much participants on average wanted to donate to the different charities. The same is true for Studies 6 and 7) Participants were told that one of the charities is among the most cost-effective charities, whereas the other four are of average cost-effectiveness. They were then asked whether they would want to pay $50 out of the $100

to find out which of the charities is among the most cost-effective charities, or whether they would prefer not to do so, in order to be able to donate the full $100 to one of the charities. Next, we asked the participants which charity they would actually donate to. Finally, participants were given a post-experimental questionnaire. This included an attention check question, asking participants to re-state how many times more effective the most effective charity is compared to a charity of average cost-effectiveness according to the information they have just read.

The five charities mentioned in the materials and method section were (the order was randomized):
- Miracle Feet
- H2O for Life
- Schistosomiasis Control Initiative
- EcoAgriculture Partners
- African Solutions to African Problems

After the list of charities was given, participants were reminded that the most effective charity was [1.5/100] more effective than the others. After participants had made their decision to pay or not pay to know which charity was the most effective one, they were asked to give a free text response explaining their decision. Participants were also given an attention check asking how much more effective the most effective charity was than the average charity. In addition, participants completed the Cognitive Reflection Test and a number of demographic questions.

Results and Discussion
A Chi-squared test revealed that significantly more participants ($\chi2(1) = 49.66$, $p < .001$, $d = 1.03$, 95% CI [0.74, 1.32]) indeed did choose to pay $50 to find out which charity is among the most cost-effective in the 100 ratio condition (60.0%, 69 out of 115 participants) than in the 1.5 ratio condition (14.8%, 18 out of 121 participants). Thus, people who are informed about the median expert's cost-effectiveness ratio estimate are more willing to invest resources to find the most cost-effective charities, than people who are given the median lay estimate. We also found a corresponding difference between the conditions with regards to which charity the participants donated to. Participants in the 100 ratio condition were significantly more likely to donate to the most effective-charity (55.7%) than participants in the 1.5 ratio condition (12.4%; $\chi2(1) = 47.62$, $p < .001$, $d = 1.00$, 95% CI [0.72, 1.29]). This is explained by the fact that 82.8% of participants who paid to identify the most effective charity also donated to it, whereas 95.3% of participants who did not pay to identify the most effective charity donated to an ineffective charity. In sum, the results suggest that informing people of how large the difference in effectiveness is did lead to more effective giving, via a greater tendency to invest resources to identify the most effective charities. However, this donation task is fairly complex and it may be argued that it does not reflect real-world donation contexts. It is possible that we found strong effects of the manipulation because it is relatively easy to compare the costs of the investments with the expected benefits of finding the most effective charity in this task.

A binary logistic regression found that CRT was not a significant predictor of deciding whether or not to pay for information when accounting for all other demographics and for variance condition. However, income was a significant predictor, $\chi^2(1) = 3.9$, $p = .048$. Additional logistic regressions for each variance condition separately showed that participants with higher income were more likely than participants with lower income to pay for

information when the variance in charity effectiveness was a ratio of 100, $\chi^2(1) = 6.7$, $p = .009$, but not when the ratio was 1.5, $\chi^2(1) = 0.3$, $p = .581$. No other demographic variable was a significant predictor.

Results without any exclusions. The sample without any exclusions consisted of 302 participants. Similar to results reported in the main manuscript, a Chi-squared test revealed that significantly more participants ($\chi^2(1) = 48.04$, $p < .001$, $d = 0.87$, 95% CI [0.62, 1.12]) did choose to pay \$50 to find out which charity is among the most cost-effective in the 100 ratio condition.

Continuous Splitting study (not reported in the main manuscript)

In this study, we tested whether correcting the misconception would make giving more effective in a task where participants are presented with two charities which differ in cost-effectiveness. We knew from previous studies (Baron & Szymanska, 2011; Berman et al., 2018; Caviola et al., 2014) that when presented with multiple charities, people have an inclination to split their donations—even if that reduces the number of lives saved—and wanted to see if this tendency would be reduced if the participants learned that the cost-effectiveness ratio is large.

Method

Participants. We recruited 601 US American participants from MTurk. Eighty-six were excluded because they failed at least one of two attention check questions. The final sample consisted of 515 participants (266 female, age $M = 37.07$, $SD = 12.28$). Our sample size of 601 was determined by a power analysis which suggested we would need at least 540 participants to detect a small effect of $f = 0.15$, taking an alpha of .05, a power of .80, and five groups. We aimed to recruit 600 participants to account for exclusions. All findings remained the same when conducting the analysis without any exclusions.

Procedure. Participants were divided into five conditions. In the first four conditions, participants were told that "in a recent report, researchers concluded that the most cost-effective charities are [1.05/1.5/10/100] times more effective than a charity of average level of cost-effectiveness". In the fifth (control) condition, no information about cost-effectiveness was given. Participants were then asked to allocate \$100 between a charity of the highest level of cost-effectiveness and a charity of average cost-effectiveness. We randomly presented different charities, used names of actual (non-fictional) charities and donated in accordance with the participants' choices after the study. Finally, participants were given a post-experimental questionnaire.

Results and Discussion. While there was an overall effect across conditions $F(4, 510) = 6.41$, $p < .001$, $\eta_p^2 = .05$, 95% CI [.014, .082], post-hoc comparisons showed no significant differences regarding how much was donated to the highly effective charity between the 1.05 (cost-effectiveness) ratio (M = \$64.21, SD = 21.13), 1.5 ratio ($M = \$65.42$, $SD = 22.91$), 10

ratio (M = \$70.81, SD = 21.21), and 100 ratio (M = \$70.35, SD = 24.30) conditions. The overall effect was due solely to differences between the control group (M = \$57.09, SD = 24.14) and the 10 ratio (p < .001) and 100 ratio (p < .001) conditions. Hence, correcting misconceptions about the cost-effectiveness ratio did not substantially affect the effectiveness of participants' giving in this study.

However, in all conditions where participants had been informed about the cost-effectiveness ratio, their beliefs reflected this information. Participants who were told that the cost-effectiveness ratio was 1.05 [1.5/10/100] on average thought that the cost-effectiveness ratio was 1.05 [1.5/10/100] (median responses). Generally, participants' beliefs were predicted by their condition $F(4, 510)$ = 28.59, p < .001, η_p^2 = .18, 95% CI [.12, .24]. These findings show that participants believed the information they were given. Taken together, they suggest that correcting people's false beliefs about the difference in charities' effectiveness is straightforward. However, such corrections does not necessarily translate into more effective donations in all contexts.

Continuous Splitting study with explicit comparison to misconception (not reported in the main manuscript)

Based on a pilot study, we hypothesized that the manipulation would yield stronger effects if participants were informed both about the true cost-effectiveness ratio (100) and the fact that most people underestimate this number (1.5). Our assumption was that even though, as revealed in our previous studies, people implicitly believe that the difference in effectiveness between the most effective and the average charity is around 1.5-2, they do not necessarily explicitly think about this belief unless prompted. Therefore, when simply presented with the true cost-effectiveness ratio, people may not necessarily realize that this number is much higher than their previously implicitly held belief. To test that hypothesis, we included a non-comparative information condition in which participants were informed only about the true cost-effectiveness ratio, and a comparative information condition in which they were informed about both the true cost-effectiveness ratio and the cost-effectiveness ratio most people believe is true. Our hypothesis was that participants would give more effectively in the comparative than non-comparative condition if told that the true cost-effectiveness ratio is 100.

Method

Participants. We recruited 393 US American participants from MTurk. Fifty-seven were excluded because they failed the comprehension check question. The final sample consisted of 335 participants (160 females, age M = 40.66, SD = 12.64). Our sample size was based on a power analysis, which suggested that we would need a total sample of 351 to detect a small effect of f = 0.15, relying on an alpha of 0.05, power of 0.8, and four groups. We aimed to recruit 380 to account for exclusions. This study was pre-registered at https://osf.io/2tfjd/. All findings remained the same when conducting the analysis without any exclusions (reported in the supplementary materials).

Procedure. Participants were divided into four conditions in a 2 ratio (1.5 vs. 100) x 2 information (non-comparative vs. comparative) between-subjects design. Participants were asked to distribute \$100 between one of the most cost-effective charities and a charity of

average cost-effectiveness. We used names of actual (non-fictional) charities and donated in accordance with the participants' choices after the study. Participants in the non-comparative information condition were only informed that the cost-effectiveness ratio between the most effective and average charities is 1.5 [100]. Participants in the comparative information condition were informed that the cost-effectiveness ratio is 1.5 [100], and also that most people believe that the ratio is 1.5. Thus, we made the average participant's beliefs about the ratio explicit.

After having chosen their allocations, participants were asked two further questions with a response scale from 1 (Much less likely) to 7 (Much more likely) about their intended donation behavior in the future. First, they were asked whether learning about the cost-effectiveness ratio makes them more likely to invest time into finding the most effective charity next time they donate. Second, they were asked whether learning about the cost-effectiveness ratio makes them more likely to donate only to the most effective charities, even if they feel more strongly about a charity of average level cost-effectiveness.

Participants in the control conditions were given an unrelated text about what car color people with a zest for life choose, instead of information about cost-effectiveness.

The descriptions of the three top charities mentioned in the paper were as follows:

- **Against Malaria Foundation:** This charity saves children's lives by distributing anti-malarial bed nets. Studies suggest it is one of the **most cost-effective charities**.

- **Fortify Health:** This charity saves children's lives by adding iron and folic acid to staple foods such as flour and rice. Studies suggest it is one of the **most cost-effective charities**.

- **Results for Development:** This charity saves children's lives by stocking hospitals with amoxicillin, treatment for childhood pneumonia. Studies suggest it is one of the **most cost-effective charities**.

The descriptions of the three average charities mentioned in the paper were as follows:

- **Children's AIDS Fund:** This charity saves children's lives by teaching children how to remain HIV free in high-risk areas. Studies suggest it has an **average level of cost-effectiveness.**

- **WaterAid:** This charity saves children's lives by providing safe drinking water and sanitation to - remote communities. Studies suggest it has an **average level of cost-effectiveness.**

- **Forever Angels:** This charity saves children's lives by providing formula milk to severely malnourished infants. Studies suggest it has an **average level of cost-effectiveness.**

Participants also responded to an attention check, asking about the difference in effectiveness between the most effective and average charities, or, in the control condition, what car color people with a zest for life choose. In addition to the question how they would donate $100,

participants were also asked which of the two charities they would donate to if they could not split the money. Participants responded on a scale from 1 (definitely [top charity]) to 6 (definitely [average charity]).

Then participants were asked how important they thought charity work in each of the six charity areas (food fortification, HIV prevention, clean water and sanitation, treating pneumonia, malaria prevention, providing baby milk formula) is (1 = not important at all, 7 = very important) and how strongly they felt about charity work in each area (1 = not at all, 7 = very strongly). Participants were then asked five questions assessing how important effectiveness was to them when donating (irresponsibility to donate without researching effectiveness, effectiveness is the most important factor, invest a lot of time picking the best charities, it's not very important to find out effectiveness, one should always consult charity evaluators comparing effectiveness). They were then asked to rate eight factors on how important they are in their donation decisions. These factors were: caring deeply about the cause, personal experience with the cause, the charity fits with their religious beliefs, the charity being respected and established, the charity works in their own community, cost-effectiveness, low overhead costs and the charity's cause being neglected.

Participants were asked to report in a free text field how many times more cost-effective they believe the most cost-effective charities that focus on helping the world's poorest people are in comparison to charities of average level cost-effectiveness in this area. Participants in the control condition were asked to explain the effectiveness ratio (between the top and an average charity) they had estimated in a free text field. Participants in the other conditions were asked how surprised they were by the effectiveness ratio they were informed of.

Participants then responded to the Empathic Concern Scale (1), consisting of seven items and one attention check. Participants were then given a three-item version of the Cognitive Reflection Test (2) and a number of demographic questions.

Results and Discussion

We tested our hypotheses using a 2x2 between-subjects ANOVA (see Figure 1). There was an interaction between the cost-effectiveness ratio and information conditions $F(1, 331) = 8.49$, $p = .004$, $\eta p2 = .03$ 95% CI [.003, .066]. Participants in the non-comparative information condition donated about the same amount to the most effective charity if told that the cost-effectiveness ratio is 1.5 (M = $48.60, SD = 29.70) and 100 (M = $50.50, SD = 30.80), respectively, $p = .97$. However, participants in the comparative information condition donated more to the more effective charity if told that the cost-effectiveness ratio is 100 (M = $63.70, SD = 30.30) rather than 1.5 (M = $42.50, SD = 29.60), $p = < .001$. Thus, as predicted, information about the cost-effectiveness ratio did make a difference to people's giving behavior in the comparative information condition, but not in the non-comparative information condition (Figure S2). There was a main effect of cost-effectiveness ratio on average amount donated to the more effective charity, $F(1, 331) = 11.26$, $p < .001$, $\eta p2 = .03$, 95% CI [.006, .079], but no main effect of information type, $F(1, 331) = 0.95$, $p = .33$, $\eta p2 = .003$, 95% CI [0, .025]. However, note that even in comparative information, 100 ratio condition, there was substantial splitting between the two charities, potentially because the participants have a preference for being "fair" to the less effective charity (Sharps & Schroeder, 2018).
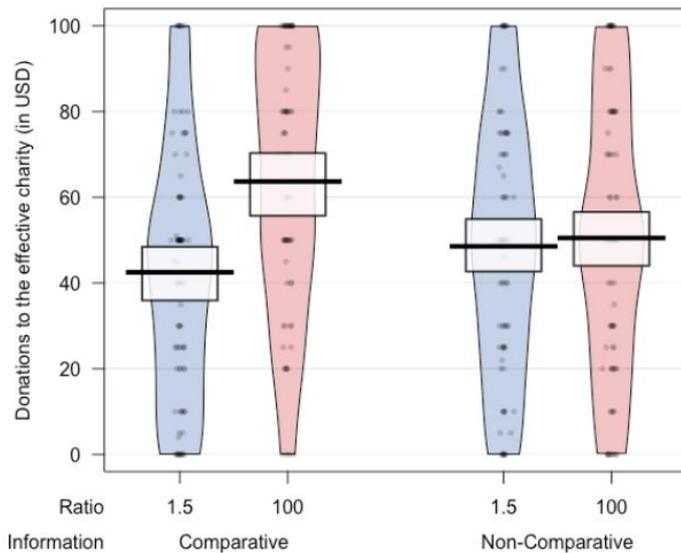
Figure S2. Informing participants that the most effective charities are 100 times more effective (red plots) instead of 1.5 times more effective (blue plots) than an average charity can increase donations to the most effective charity. This is then the case if the true cost-effectiveness ratio of 100 is contrasted with lay people's belief of the ratio only being 1.5. The thick lines represent the means and the bands represent the 95% CI. Dot represents individual responses.

We found that participants were more likely to say that they intended to exclusively donate to the most effective charities in the future, and that they intended to invest time to find the most cost-effective charity next time they if they believed the ratio is 100 instead of 1.5 (main effects of ratio), $F(1, 331) = 12.79$, $p < .001$, $\eta p2 = .04$, 95% CI [.008, .087] and $F(1, 331) = 13.19$, $p < .001$, $\eta p2 = .04$, 95% CI [.009, .087]. There were two significant interaction effects, revealing that the cost-effectiveness ratio only made a difference in the comparative information conditions, $F(1, 331) = 5.31$, $p = .02$, $\eta p2 = .02$, 95% CI [0, .022] and $F(1, 331) = 11.93$, $p < .001$, $\eta p2 = .03$, 95% CI [0, .040].

The findings of this study show that informing people about the fact that charities differy very widely in their cost-effectiveness can increase their willingness to give to effective charities. This suggests that ineffective giving observed in the real world may at least partly be caused by the fact that people are unaware of the vast difference in charities' effectiveness. However, we found that effective giving only increases if people are informed about both the true cost-effectiveness ratio and the ratio most people believe is true. Future research could investigate in more depth why just presenting people with the true cost-effectiveness ratio is not sufficient. One possible explanation is that people do not have an explicitly formed belief about the cost-effectiveness ratio unless prompted, and therefore are not sufficiently surprised when presented simply with the true cost-effectiveness ratio.

There were significant differences in beliefs about importance of charity effectiveness across conditions $F(4, 504) = 2.61$, $p = .035$, $\eta p^2 = .02$. Post-hoc comparisons with a Bonferroni correction revealed this was due to participants in the 1.05 variance condition finding charity effectiveness significantly more important ($M = 43.43$, $SD = 20.04$) than participants in the control condition ($M = 34.68$, $SD = 21.6$). No other conditions differed significantly on this measure.

There were significant differences in level of surprise across conditions $F(3, 394) = 18.35$, $p < .001$, $\eta p^2 = .12$ (see Table S1). Post-hoc comparisons with a Bonferroni correction revealed that participants were significantly more surprised when informed that top charities are 100 times more effective than the average charity than when informed that top charities are 10 times ($p = .026$), 1.5 times ($p < .001$), or 1.05 times ($p < .001$) more effective than the average charity. Participants were also significantly more surprised when informed that top charities are 10 times more effective than the average charity than when informed that top charities are 1.5 times more effective than the average charity ($p < .001$). Participants were also significantly more surprised when informed that top charities are 1.05 times more effective than the average charity than when informed that top charities are 1.5 times more effective than the average charity ($p = .011$). We found no difference in surprise between being informed that top charities are 1.05 and 10 times more effective than the average charity, respectively ($p = 1.00$).

Table S1. Mean and standard deviations of surprise when informed that the top charities are 1.05, 1.5, 10 and 100 times more effective than the average charity, reported from 1 (not at all surprised) to 7 (incredibly surprised)

	1.05	1.5	10	100
Surprise	4.5(1.85)	3.73(1.76)	4.84(1.91)	5.58(1.62)

Several measures correlated with amount donated. Participants donated more to the more effective charity, the more important they considered charity effectiveness to be, $r(599) = .09$, $p = .04$, and the higher they scored on cognitive reflection, $r(599) = .17$, $p < .001$. Conversely, the more empathetic people were, the less they donated to the more effective charity (i.e., the more they split), $r(599) = -.11$, $p = .02$. There were no significant correlations of amount donated with age, education (although this approached significance with p = .06) or income.

Results without any exclusions. The sample without any exclusions consisted of 601 participants. While there was an overall effect across conditions $F(4, 596) = 4.86$, $p < .001$, $\eta_p^2 = .03$, 95% CI [.01, .06], *post-hoc* comparisons showed no significant differences regarding how much was donated to the highly effective charity between the 1.05 (cost-effectiveness) ratio ($M = \$63.45$, $SD = 20.86$), 1.5 ratio ($M = \$64.07$, $SD = 22.42$), 10 ratio ($M = \$68.67$, $SD = 21.93$), and 100 ratio ($M = \$70.35$, $SD = 24.30$) conditions. The overall effect was due solely to differences between the control group ($M = \$66.79$, $SD = 24.95$) and the 10 ratio ($p < .001$) and 100 ratio ($p < .001$) conditions. This is the same pattern of results as the ones reported in the main manuscript.

Single-Charity study (not reported in the main manuscript)

Participants. We recruited 563 US American participants from MTurk. Forty-five were excluded because they failed at least one of two attention check questions. The final sample consisted of 518 participants (253 female, age $M = 35.79$, SD = 11.37). Our sample size was based on a power analysis suggesting that we would need a total sample of 547 to detect a small effect of $f = 0.12$, relying on an alpha of 0.05, power of 0.8, and four groups. We aimed to recruit 560 participants to account for exclusions. This study was pre-registered at https://osf.io/kvy57/.

Procedure. Participants were divided into four conditions in a 2x2 between-subjects design (high vs. low cost-effectiveness ratio x top vs. average-effectiveness charity). Participants read the same explanation of the concept of cost-effectiveness as in Study 1 and were then told that the most effective charities are either 1.5 or 100 times more cost-effective than an average charity. Participants were then presented with a charity, Forever Angels, said to save "children's lives by providing formula milk to severely malnourished infants", and either told that it is among the most cost-effective charities (top condition), or that it is of average cost-effectiveness (average condition). They were informed that one randomly chosen participant would receive a $100 Amazon gift voucher. We asked them to decide how much of this $100 gift voucher they would like to keep for themselves and how much they would like to donate to the presented charity, should they be the chosen participant. Participants were also given a post-experimental questionnaire.

Method. The post-experimental questionnaire included four measures of the "warm glow" associated with the donation. Participants were then asked four questions assessing how important effectiveness was to them when donating, and were given an attention check. They were then asked to report in a free text field how many times more cost-effective they believe the most cost-effective charities are in comparison to charities of average level cost-effectiveness. Participants then completed the Empathic Concern Scale (1) and answered a number of demographic questions.

Results. A 2x2 ANOVA showed that participants donated significantly more to the charities with top cost-effectiveness than to those with average cost-effectiveness, $F(1, 514) = 6.411$, $p = .01$, $\eta p^2 = .018$. Surprisingly, however, there were no significant differences in donations depending on whether the cost-effectiveness ratio was high or low, $F(1, 514) = 0.02$, $p = .89$, $\eta p^2 < .01$. Nor was there a significant interaction effect between the top/average factor and the cost-effectiveness ratio factor, $F(1, 514) = 0.27$, $p = .60$, $\eta p^2 < .01$. Notably, though our manipulation had no effect on the participants' donations, it was successful in affecting their beliefs. Participants who were told that the cost-effectiveness ratio was 100 on average thought that the cost-effectiveness ratio was 100 (median response), whereas participants who were told that the cost-effectiveness ratio was 1.5 on average thought that the cost-effectiveness ratio was 1.5 (median response). The difference in estimated cost-effectiveness ratio between the conditions was significant, U = 6690, p < .001.

Results of exploratory analyses. Amount donated correlated positively and significantly with the two warm glow measures, i.e., how positive participants felt about their donation, $r(516) = .48$, $p < .001$, how happy they felt about their donation, $r(516) = .44$, $p < .001$, and

with the two do good measures, i.e., how much good they thought their donation would do, $r(516) = .54$, $p < .001$ and how many lives they thought their donation would save, $r(516) = .52$, $p < .001$. Empathy, $r(516) = .14$, $p < .001$, education, $r(516) = .09$, $p = .04$, and income, $r(516) = .09$, $p = .04$, were correlated positively and significantly with amount donated, while age was not.

Impact-Salience study (not reported in the main manuscript)

Participants. We recruited 383 US American participants from MTurk. Eighty one were excluded because they failed the comprehension check question. The final sample consisted of 302 participants (168 females, age M = 36.03, SD = 11.78). Our sample size was based on a power analysis, which suggested that we would need a total sample of 351 to detect a small effect of f = 0.15, relying on an alpha of 0.05, power of 0.8, and four groups. We aimed to recruit 380 to account for exclusions. This study was pre-registered at https://osf.io/4n2q6/.

Procedure. Participants were divided into four conditions in a 2 (1.5 vs. 100 ratio) x 2 (high vs. low salience) between-subjects design and again read about the concept of cost-effectiveness, this time measured by saved units of health (due to reduced suffering or saved lives) as measured by experts. They were told that "in a recent report, researchers concluded that the most cost-effective charities are [1.5/100] times more effective than a charity of average level of cost-effectiveness". They were then presented with a charity said to be among the most cost-effective, focusing in this case on arthritis research, and a charity said to be of average cost-effectiveness, focusing in this case on cancer research. We chose arthritis and cancer research because we thought that we might observe ceiling effects unless participants in the low salience condition chose to donate substantial amounts to the less effective charity, and because we thought that people would have a preference for the cancer research charity, which would reduce their donations to the arthritis research charity. Participants were then asked to allocate $100 between the two charities, and were informed that $100 would actually be distributed based on the average allocation across all participants.

At this point (right before they chose their allocation), participants in the high salience conditions (but not those in the low salience conditions) were presented with the following information. First, the results of their choices were explained in five sentences at $25 intervals. One sentence read: "If you donate $25 to Arthritis research and $75 to Cancer research, you will save 25 [1.125] units of health". Another read: "If you donate $100 to Arthritis research and $0 to Cancer research, you will save 100 [1.5] units of health". Below this text a graph was shown that displayed the same information visually. The graph displayed a straight line indicating how many units of health different allocations to the arthritis and the cancer research charities saved. Allocations to the arthritis and cancer research charity were on the X-axis (with higher amounts to the arthritis research charity further to the right) and saved units of health on the Y-axis. The line was much steeper in the 100 ratio compared to the 1.5 ratio condition. After having chosen their allocations, participants were given a post-experimental questionnaire.

Method. The graph referred to in the material and method section is shown below (Figure S1).
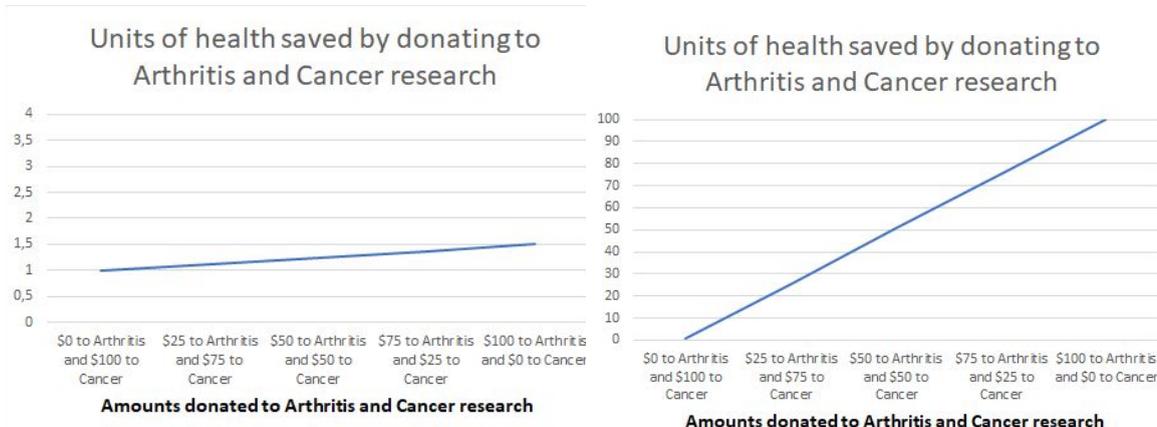
Figure S1 Graphs depicting how many units of health could be saved with different allocations when the top charities are 1.5 times (left) and 100 times (right) more effective than the average charity. The graphs were only shown to participants in the conditions where the implications of differences in effectiveness on one's impact were made salient.

The post-experimental questionnaire referred to in the materials and methods section was as follows. First, participants were asked a question about the importance of arthritis and cancer. Participants were asked if it is true or false that in the preceding donation task, the more one donates to the cancer research charity, the fewer units of health will be saved. Participants who gave the wrong answer were not excluded, because we suspect participants misunderstood the question to mean: "Everything else equal, the cancer research charity is harmful". As a result, 118 out of 303 participants answered with no instead of yes. What we meant to ask, however, was if distributing more of the 100 dollars to the cancer research charity and thus less to the arthritis charity saves less lives.

Participants were asked how important they found charity work in the two charity areas, cancer and arthritis research (1 = not important at all, 7 = very important), and how strongly they felt about charity work in each area (1 = not at all, 7 = very strongly). Participants were then asked five questions assessing how important effectiveness was to them when donating and a number of demographic questions.

Results. We tested our hypothesis using a 2x2 between-subjects ANOVA. There was an interaction between the cost-effectiveness ratio and salience conditions $F(1, 305) = 6.01$, $p = .015$, $\eta p^2 = .019$. In line with the results from Study 4, participants in the low salience condition donated about the same amount to the most effective charity if told that the cost-effectiveness ratio is 1.5 ($M = \$53.01$, SD = 26.97) and 100 ($M = \$48.95$, SD = 30.35), respectively, $p = .81$. Participants in the high salience condition donated more to the more effective charity if told that the cost-effectiveness ratio is 100 ($M = \$68.03$, SD = 28.8) rather than 1.5 ($M = \$56.26$, SD = 26.94), $p = .05$. Thus, in the high salience condition, information about the cost-effectiveness ratio did make a difference to people's giving behavior, whereas that was not the case in the low salience condition. However, note that even in high salience condition, there was substantial splitting, potentially because the participants have a preference for being "fair" to the less effective charity. Note also, that this result is only barely significant and that we are not confident in how robust the effect is. There was no main

effect of cost-effectiveness ratio on average amount donated to the more effective charity $F(1, 305) = 1.42$, p = .23, $\eta p^2 = .005$, and also no main effect of salience $F(1, 305) = 0.43$, $p = .43$, $\eta p^2 = .036$.

Results of exploratory analyses. Correlations (see Table S2 below) suggest that feeling good about donating to the arthritis charity, and viewing the arthritis charity as doing good, are fairly consistent predictors of donating more to the arthritis charity. However, not feeling good about donating to the cancer charity, and viewing the charity as not doing good was a larger and more consistent predictor of donating more to the arthritis charity (and donating less to the cancer charity by implication). This is in line with research that suggests people are more interested in knowing that their chosen charity is not corrupt or wasteful, than they are interested in knowing that their charity is particularly cost-effective (3).

Table S2. Descriptive statistics of effectiveness estimates in two effectiveness variance conditions (1.5 vs 100) and the two salience conditions (low vs high)

Measure	Charity	Total	1.5 Low	1.5 High	100 Low	100 High
Feel Good	Arthritis (Top)	.24**	.20	.28**	.27*	.28*
	Cancer (Average)	-.41**	-.39**	-.28*	-.57**	-.35**
Do Good	Arthritis (Top)	.23**	.29*	.17	.27*	.19
	Cancer (Average)	-.41**	-.24*	-.43**	-.50**	-.43**

In addition, older participants were more likely to donate more effectively, $r(296) = .13$, $p = .03$. Participants who considered effectiveness an important determinant of charitable giving were also more likely to donate more effectively, $r(296) = .16$, $p = .01$. There were no significant correlations between income and and amount donated effectively ($p = .73$) nor between education and amount donated effectively, although the latter approached significance ($p = .08$).

References

1. Davis MH (1980) *A Multidimensional Approach to Individual Differences in Empathy* (American Psychological Association).

2. Frederick S (2005) Cognitive Reflection and Decision Making. *J Econ Perspect* 19(4):25–42.

3. Bagwell S, de Las Casas L, van Poortvliet M, Abercrombie R (2013) Money For Good UK: Understanding donor motivation and behaviour. Retrieved from:

https://www.bl.uk/collection-items/money-for-good-understanding-donor-motivation-and-behaviour

4. Berman JZ, Barasch A, Levine EE, Small DA (2018) Impediments to Effective Altruism: The Role of Subjective Preferences in Charitable Giving. *Psychol Sci 29*(5):834–844.