Boosting the impact of charitable giving

Lucius Caviola^{1*}, Joshua D. Greene^{1,2}

¹ Department of Psychology, Harvard University; Cambridge, MA, USA

² Center for Brain Science, Harvard University; Cambridge, MA, USA

* Correspondence: lcaviola@fas.harvard.edu

Keywords: altruism, moral motivation, charitable giving, charity bundling, effective altruism

Abstract

The most effective charities are hundreds of times more impactful than typical charities. Yet most donors favor charities with personal/emotional appeal over effectiveness. We gave donors the option to split their donations between their personal favorite charity and an expert-recommended highly effective charity. This bundling technique increased donors' impact without undermining their altruistic motivation, boosting effective donations by 76%. An additional boost of 55% was achieved by offering matching donations with increasing rates for allocating more to the highly effective charity. We show further that matching funds can be provided by donors focused on effectiveness through a self-sustaining process of micro-matching. We applied these techniques in a new online donation platform (GivingMultiplier.org) which fundraised over 1.5 million USD in its first 14 months. While prior applied research on altruism has focused on the quantity of giving, the present results demonstrate the value of focusing on the effectiveness of altruistic behavior.

Each year hundreds of thousands of people die from malaria¹; millions suffer from debilitating parasitic worm infections^{2,3}; and over a million lose vision due to trachoma⁴. These diseases can be prevented at low cost. Consequently, charities focused on these problems can be hundreds of times more effective than typical charities, saving or significantly improving many more lives per dollar^{5,6}. For example, it costs ~50,000 USD to train a guide dog to help a blind person in a rich nation, but 100 USD can save a person in a less wealthy nation from trachoma-induced blindness⁵. Americans alone donate ~450 billion USD (2% of GDP) each year⁷, but relatively little goes to the most effective charities. Why don't people give more effectively? And how might giving become more effective?

Applied research on charitable giving has focused primarily on strategies for increasing the *amount* given^{8,9}. By contrast, here we focus on the *effectiveness* of giving, consistent with the ideals of effective altruism^{10,11}. Charitable giving is less effective than it could be for several reasons¹². First, many people are unaware of the massive differences in cost-effectiveness across charities⁶. However, informing donors about the relative effectiveness of charities often has little effect on their choices¹³. Most donors favor charities that are personally meaningful and emotionally appealing^{9,13}. While such feelings play a critical role in motivating altruism in the first place^{9,14,15}, they limit the positive impact of charitable giving^{11,12,16}. One might conclude from prior research that few donors value effectiveness. Here we propose an alternative theory emphasizing mixed altruistic motivations. More specifically, we hypothesize that most donors have at least two, often conflicting, motivations: Donors care primarily about supporting personally/emotionally appealing charities, but they also have a significant secondary motivation to support effective causes. Thus, we hypothesize that many donors would be willing to support expert-recommended effective charities, as long as they can also support their personal favorite charities.

This suggests a strategy for increasing the effectiveness of giving using *bundled* donations: providing donors with the option to split their donation between a personal favorite charity and an expert-recommended highly effective charity ("favorite-effective" bundles). This donation bundle strategy aims to increase the effectiveness of giving by working *with* people's

motivations rather than attempting to displace them¹⁷. We hypothesize that favorite-effective bundles will be appealing because they satisfy donors' desire to support a personally meaningful cause, while also satisfying a secondary desire to engage in altruism that is highly effective. Studies 1-5 test our theory of mixed altruistic motivation. Study 1 demonstrates that offering favorite-effective donation bundles can increase effective giving. Studies 2-5 examine the psychological mechanisms behind the distinctive appeal of favorite-effective bundles. Our findings suggest that favorite-effective bundles offer donors an appealing combination of giving with the "heart" and giving with the "head".

Studies 6-7 and a real-world proof-of-concept demonstration aim to answer three questions concerning the real world applicability of the bundle technique: First, could offering matching incentives^{18,19} increase donors' willingness to make favorite-effective bundled donations? If so, this raises a second question: Who would provide matching incentives, and how would their provision be sustained? Finally, there is the question of generalizability beyond the lab: Would real-world donors choose to make favorite-effective bundled donations if they were alerted to the opportunity and offered matching incentives?

Our proposed practical solution is a technique that we call *micro-matching*. Matching funds are typically provided by a single large donor working directly with a charity. Micro-matching aims to democratize the matching process, allowing many ordinary donors to provide matching incentives for other donors. Micro-matching leverages variation in donors' priorities. Among donors willing to make favorite-effective bundled donations, some may care primarily about their personal favorite charities, while others may care primarily about effectiveness. We propose micro-matching as a way of coordinating the actions of donors that differ in this way. Donors focused primarily on effectiveness may be willing to supply matching funds, especially if this multiplies their own effectiveness: Instead of directly supporting effective charities, supplying matching funds could incentivize others to donate even more to effective charities. In turn, donors who are primarily focused on their personal favorite charities may be incentivized to toward greater effectiveness by the offer of matching funds for favorite-effective bundled donations.

Thus, micro-matching aims to produce a sustainable, win-win exchange between donors with different, but compatible, altruistic priorities. Study 6 shows that matching incentives, when combined with favorite-effective bundles, can indeed increase the *effectiveness* of donations (as distinct from the *quantity* of donations^{18,19}). Study 7 provides evidence that many donors are willing to provide matching funds, enough that a micro-matching system could be financially self-sustaining.

Although the two key techniques introduced here—favorite-effective donation bundling and micro-matching—may seem unrelated, they in fact rely on closely related insights: Favorite-effective bundling leverages mixed motivations *within* people, while micro-matching leverages mixed motivations *between* people. Our final step is to combine these techniques in a real-world proof of concept. For this purpose, we custom-built a new donation platform (GivingMultiplier.org) which remains publicly available.

Design and Results

Bundling to increase effective giving

Study 1 (N = 895) asked whether donation bundling can increase effective giving. All participants could donate up to \$100 to charity and decide where to allocate the funds. We emphasize that in all donation decisions, here and subsequently, participants kept funds not donated, making decisions incentive compatible. All decisions were executed by lottery^{20–22}, enabling us to use more substantial donation amounts. This study had five conditions. In the baseline giving condition, participants could only donate to their personal favorite charity (or keep the funds). This reflects the typical real-world context in which donors consider their personal favorite charities but not charities backed by research on effectiveness^{6,12}. In the control condition, participants had two donation options: donating exclusively to their personal favorite charity (all-favorite) or exclusively to a highly effective charity (all-effective). This condition follows previous research on (in)effective giving in which participants have two charity options and cannot

split between them^{13,23}. It also reflects the options typically considered by effective altruists who recommend supporting highly effective charities instead of personal favorite charities^{10,11}. In the primary experimental condition (the three-option bundle condition), participants had the two options available in the control condition (all-favorite and all-effective) plus the option to donate with a 50/50 favorite-effective bundle. Study 1 included two alternative experimental conditions. The two-option bundle condition lacked the all-effective option. The free-split bundle condition replaced the 50/50 bundle option with one in which participants could adjust the favorite/effective ratio.

In the baseline giving condition, 0% of participants gave to the effective charity, as this was not possible. In the control condition, 82.3% gave everything to their favorite charity, and 17.7% gave everything to the effective charity. In the primary experimental condition (the three-option bundle condition), 45.6% gave everything to their favorite charity, 51.6% chose the 50/50 bundle, 2.7% gave everything to the effective charity (Figure 1A). In the two-option bundle condition, 49.5% gave everything to their favorite charity and 50.5% chose the 50/50 bundle. The median response in the free-split bundle condition was to give 75% to the favorite and 25% to the effective charity. In addition to choosing an allocation proportion (e.g. 50/50 bundle), donors chose a total donation amount (e.g. \$80 out of \$100), and thus an amount for each charity (e.g. \$40 favorite/\$40 effective). Adding the bundle option increased total donations to the effective charity by 76% (three-option bundle: $M = \$14.51$, $SD = 19.4$ vs. control: $M = \$8.23$, $SD = 22.2$; $\chi^2 = 177.31$, $df = 4$, $p < .001$). The results from the two other experimental conditions did not significantly differ from the three-option bundle condition (two-option bundle: $M = 13.59$, $SD = 17.16$; free-split bundle: $M = 11.96$, $SD = 14.15$). Total donation amounts (across charities) did not differ across all five conditions. (For detailed results see Supplementary Materials.) In short, Study 1 demonstrated that simply allowing donors to split their donation between their favorite charity and a highly effective charity, rather than forcing them to choose between the two, can substantially boost effective giving. A subsequent study produced similar results using other effective charities as the designated effective charity. (See Supplementary Materials.)

Figure 1. The distinctive appeal of bundling a personal favorite charity with a highly effective charity. (A) In Study 1 few participants chose the highly effective charity over a personal favorite charity in a forced choice, but a majority chose a 50/50 favorite-effective bundle when given the option (three-option bundle condition), increasing funds to the effective charity by 76%. (B) In Study 2, bundling a favorite charity with a charity described as effective was more appealing than bundling it with one described as popular or described neutrally.

Study 2 (N = 227) asked whether people find favorite-effective bundles distinctively appealing, or whether they simply like donation bundles, consistent with prior research^{24,25}. Participants chose among three 50/50 bundles, each including the participant's favorite charity plus a second charity. All three second charities were in fact highly effective but were described (counterbalanced) as either highly effective (favorite-effective), highly popular (favorite-popular), or without additional description (favorite-neutral). Participants also rated each charity for perceived effectiveness and popularity. Participants preferred the favorite-effective bundle, which was chosen by 56% of participants, with 21% choosing the favorite-popular bundle ($\chi^2 = 36.36$, $df = 1$, $p < .001$) and 22% choosing the favorite-neutral bundle ($\chi^2 = 33.12$, $df = 1$, $p < .001$). Only effectiveness ratings predicted participants choices ($\beta = .82$, $p < .001$; popularity, $\beta = -.13$, $p = .51$; see Supplementary Materials). A follow-up study using a familiar popular charity produced comparable results (see Supplementary Materials). This indicates that favorite-effective bundles are distinctively appealing (Figure 1B).

Figure 2. Explaining the appeal of favorite-effective donation bundles. (A) Study 3 shows diminishing returns for increasing allocations to one's personal favorite charity within a bundle. * indicates $p < .001$. (B) In Study 4 bundle donors (49%) viewed bundled donations as an appealing compromise, supporting a personally meaningful cause (red) while giving effectively (blue), and also being fair (orange). (C) In Study 5 bundle donors were perceived as more warm (red) than all-effective donors and more competent (blue) than all-favorite donors. All error bars show SEM.

Study 3 (N = 242) asked why favorite-effective bundles are appealing, testing the hypothesis that giving exclusively to one's favorite charity yields diminishing returns of value for the donor. All participants made four decisions, in each case choosing between an all-favorite donation and a favorite-effective bundle. The bundle options varied in their proportions (favorite/effective) as follows: 10/90, 40/60, 60/40, 90/10. This facilitated a simple test for diminishing returns, comparing a 30% increase in the proportion allocated to the favorite charity in the lower range (10% vs. 40%) vs. the upper range (60% vs. 90%). As expected, participants preferred bundles with higher allocations to the favorite charity (Figure 2A). But crucially, participants were more sensitive to a 30% shift in the lower range vs. the upper range, consistent with diminishing returns. An additional 15% of participants chose the favorite-effective bundle when it was split 40/60 vs. 10/90 (46% vs. 31%). By contrast, only an additional 5% chose the bundle when it was split 90/10 vs. 60/40 (64% vs. 59%)—a significant difference of differences ($b = -.96$, $z = -2.48$, $p = .01$). Overall, most donors wanted to give at least half of their donation to their favorite charity, but after that point, most were willing to support the expert-recommended effective charity.

Giving to one's favorite charity succumbs to diminishing returns, making room for a second charity. But why are effective charities distinctively appealing as second charities, as shown in Study 2? In Study 4 (N = 265) participants rated all three options for effectiveness, supporting a personal cause, fairness, and for overall appeal. As in Study 1, many participants (49%) chose the favorite-effective bundle, while 47% chose all-favorite, and 4% chose all-effective. Here we focus on the 49% who chose the bundle. Unsurprisingly, the donors who chose the favorite-effective bundle rated it as overall more appealing than the all-effective option ($t(220) = 7.28$, $p < .001$, $d = .62$). However, they rated the bundle option as only slightly (and not significantly) more appealing than the all-favorite option ($t(220) = 2.04$, $p = .11$, $d = .20$). This is consistent with our hypothesis that bundle donors are *primarily* motivated to support a personally meaningful cause. Consistent with the mixed motivations hypothesis, the overall appeal of the bundle was independently predicted by feelings of supporting a personal cause ($t(107) = 3.25$, p

= .002, $\beta = .30$), perceived effectiveness ($t(107) = 2.24$, $p = .03$, $\beta = .19$), and also perceived fairness ($t(107) = 3.13$, $p = .002$, $\beta = .28$). Critically, when rating their options for supporting a personally meaningful cause, the bundle donors saw the bundle as closer to the all-favorite option (diff = 10.41) than to the all-effective option (diff = 20.50), a significant difference of differences ($t(110) = 10.96$, $p < .001$, $d = 1.51$). (Figure 2B, red bars) This is despite the fact that the bundle option is objectively a 50/50 mix of the all-favorite and all-effective options. This subjective asymmetry is consistent with the diminishing returns for supporting the favorite charity observed in Study 3: Giving half to one's personal favorite charity feels almost as personally meaningful as giving all to one's personal favorite charity (Fig 2B red bars; compare Fig 2A red lines). When bundle donors rated their options for effectiveness, the pattern was critically different. The bundle-donors saw the bundle as intermediate in effectiveness, exactly halfway between all-effective (diff = 11.3) and all-favorite (diff = 11.3), with no significant difference of differences ($t(110) = 0.01$, $p = .99$, $d < .01$). (Figure 2B, blue bars). Bundle donors also saw the bundle as the most fair (orange bars), consistent with prior research^{24–26}. (As expected, these patterns were not observed in donors who didn't choose the bundle. See Supplementary Materials.) Thus, consistent with our mixed motivations hypothesis, those who choose the favorite-effective bundle saw it as offering an appealing compromise, with much greater effectiveness (and fairness) than the all-favorite option, but little sacrifice of personal meaning.

Moving from psychological motivations to corresponding traits, Study 5 (N = 394) examined perceptions of different donors. Participants read a hypothetical vignette about three people who donated either exclusively to a personally meaningful favorite charity (all-favorite), exclusively to a highly effective charity (all-effective), or with a 50/50 bundle (favorite-effective). Participants rated the warmth and competence²⁷ of all three people. The favorite-effective bundle donor was seen as particularly high in both warmth and competence, while the all-effective donor was perceived as less warm ($F(2, 786) = 56.87$, $p < .001$), and the all-favorite donor as less competent ($F(2, 786) = 61.14$, $p < .001$)²⁸ (Figure 2C). Thus, favorite-effective bundle donations may be appealing because they are suggestive of good overall character, to others and perhaps

to the donors themselves²⁹. More specifically, the bundled donations indicate an appealing balance of traits, paralleling the appealing balance of motivations observed in Study 4: Those who choose the favorite-effective bundle are seen as highly competent, but with no sacrifice of apparent warmth. (Fig 2 B-C, red vs. blue bars)

Micro-matching to incentivize bundling

Next, we examined the use of matching incentives as a way to advertise and incentivize favorite-effective bundled donations. Study 6 (N=421) asked whether and to what extent an offer of matching funds can induce donors to give more effectively through bundled donations (or all-effective donations). To increase the incentive to give effectively, we offered donors a higher matching rate for allocating more to the effective charity (5% for all-favorite, 25% for 50/50 favorite-effective bundle, and 50% for all-effective). The control condition offered no matching, following Study 1's three-option bundle condition.

In the control condition, 13% chose all-effective, 40% chose all-favorite, and 47% chose the favorite-effective bundle. In the matching condition, 28% chose all-effective, 21% chose all-favorite, and 51% chose the favorite-effective bundle. The offer of matching funds produced a 55% increase in funds donated to the effective charity (no matching, $M = \$17.77$, $SD = 22.67$; matching, $M = \$27.47$, $SD = 29.32$). Total amounts donated did not differ significantly between the two conditions ($t(345) = 1.04$, $p = .30$, $d = 0.11$).

A second goal of Study 6 was to examine the feasibility of donor coordination through micro-matching. Could a dollar invested in a matching fund cause more than an additional dollar to be donated to a highly effective charity? We found that in the matching condition, the average matching funds required per participant was \$3.73 ($SD = 4.06$). But in the matching condition, compared to the control condition, donations to the effective charity increased by an average of \$9.70 per participant. Thus, the rate of return for a hypothetical matching funder focused on effectiveness, who would otherwise simply donate to an effective charity, is 2.6 times the initial investment ($9.70/3.73$), or a 160% rate of return. This high rate of return suggests that a micro-

matching system could be self-sustaining, provided that enough donors are willing to provide matching funds.

Study 7 (N = 145) further tested the feasibility of micro-matching by giving donors the option to support a matching fund for subsequent donors. Here we asked whether donors' support of the matching fund would be sufficient to cover the costs of matching. Similar to the free-split bundle condition of Study 1, participants decided whether to donate and how to allocate their donation between their favorite charity and a highly effective charity. Next, they were asked whether they would like to direct the part of their donation that they allocated to the effective charity into the matching system. We found that 34% of participants were willing to provide matching funds. For every dollar donated, 34 cents were provided as matching funds. On average, 8 cents in matching funds were required for each dollar donated. Thus, there was a matching surplus of 26 cents (34 minus 8) for each donated dollar. Put another way, participants in this sample provided four times the matching funds needed to cover the costs of matching in this sample.

Proof of concept in a real-world application

Could these techniques work in the real world? Would real donors be willing to make favorite-effective bundled donations? Likewise, would real donors be willing to supply matching funds sufficient to incentivize such donations? To find out, we created a publicly accessible online donation platform, Giving Multiplier (<https://GivingMultiplier.org>), designed to encourage effective giving through bundling and micro-matching. To be clear, this proof-of-concept demonstration is not a controlled experiment, as we could not publicly advertise Giving Multiplier's essential features and then randomize would-be donors into a control condition in which those features are absent. Our aim here was not to compare Giving Multiplier's strategy to an alternative strategy, but, more modestly, to determine whether its strategy is at all viable in the real world.

Giving Multiplier users select their favorite charity—any US 501(c)(3)—and one of nine expert-recommended highly effective charities. They select a total donation amount and can use

a slider to allocate that amount between the two charities (Figure 3). The slider mechanism indicates the variable matching rate and matching amount. Prior to finalizing their donation, donors have the option to redirect the designated effective part of their donation into the matching fund.

Figure 3. Giving Multiplier slider mechanism for displaying increasing matching rates and amounts.

We launched Giving Multiplier in November 2020, advertising it primarily through social media and unpaid media coverage. By December 31st 2021, Giving Multiplier had fundraised over 1.5 million USD from 3,450 donations, with over 1 million USD going to expert-recommended highly effective charities³⁰. Analyses reported here cover 2,422 donations made before 01/01/22 from donors who consented to the use of their data for research. The median donation was \$100 ($M = \494, $SD = 2,881$). The median split ratio was 60% to the effective charity. 73% of the total went to recommended effective charities. 65% of the total went to effective charities counterfactually, meaning that donors indicated that they would not have supported these charities without Giving Multiplier. 38% of donations included support for the matching system.

The mean amount provided into the matching system was \$185 (median = \$0). The mean amount required from the matching system was \$49 (median = \$10). Thus, Giving Multiplier was financially self-sustaining during this period.

Discussion

In seven experiments and a real-world proof-of-concept demonstration, we developed and deployed a new strategy for increasing the effectiveness of charitable giving, one of the most widespread and important forms of human altruism⁷. Our strategy is based on our hypothesis that donors typically have mixed motivations for charitable giving. People primarily want to support charities that are personally meaningful and emotionally appealing, but they also care about the effectiveness of their giving, although this is not always apparent^{9,12–14,16}.

By offering donors the option to make a bundled donation, split between a personal favorite charity and a highly effective charity recommended by experts, donors can dramatically increase their effectiveness without undermining their primary motivation for giving. These favorite-effective bundled donations strike an appealing balance between the “heart” and the “head”. They allow donors to support highly effective charities while doing little to diminish their sense that they have supported a personally meaningful cause.

Whereas favorite-effective donation bundling appeals to mixed motivations *within* people, donor coordination through micro-matching appeals to mixed motivations *between* people. Donors focused on effectiveness can offer matching funds for others, encouraging them to give more effectively by making favorite-effective bundled donations (or all-effective donations). Other donors can benefit from receiving these matching funds, which allow them to direct more resources to the charities they’ve selected—both personal favorite charities and expert-recommended highly effective charities. In our proof-of-concept demonstration, we found that many donors chose to play both roles, receiving matching funds for the two charities that they selected and then choosing to convert their donation to the expert-recommended charity into a matching fund donation. It is possible that such donors were motivated by indirect reciprocity³¹, a

desire to “pay it forward” for other donors by supporting the matching fund after receiving matching funds. This possibility may be explored in future research. About a third of donors were willing to support the matching fund with some or all of their donation. This provided enough matching funds to cover the matching funds received by donors, making the micro-matching system self-sustaining.

Despite a long history of altruism, including centuries of organized philanthropy, humans have only recently attempted to systematically measure the cost-effectiveness of altruistic endeavors with the goal of doing as much good as possible^{10,11}. The effective altruism movement is growing and has been notably successful in securing large commitments from relatively few people^{32,33}. Effective altruism’s potential for more widespread adoption is unknown. The seven studies and proof-of-concept demonstration presented here are cause for optimism, grounded in a more detailed understanding of altruistic motivation. Today, relatively few donors prioritize effectiveness. But our results suggest that effective giving can be a satisfying complement to giving based on personal feelings, adding a “competence glow”²⁷ to the proverbial “warm glow” of giving. Some donors are willing to incentivize bundle donations in others, promoting a chain of giving that is both personally meaningful and effective. The stakes are high, as ordinary people have the power to do enormous good. The limited proof-of-concept demonstration reported on here raised funds sufficient to provide 100,700 deworming treatments and 17,500 malaria nets, among other benefits. (See Supplementary Materials.) A better understanding of moral motivation, and how to channel it, could dramatically increase the impact of human altruism.

Methods

All reported studies, including the final proof of concept, were pre-registered, except for Study 7 (which was a pre-test for the proof of concept). For more detailed descriptions of the methods and results, please refer to our Supplementary Materials available at https://osf.io/zu6j8/?view_only=28050663bd6b4b5cae0a73ad8068bc34. Across Studies 1-7 we recruited participants using Amazon Mechanical Turk. All participants were taken from different

samples. The financial stakes in Studies 1-7 were probabilistically implemented (executed by lottery), enabling the use of higher stakes²⁰⁻²². As per our pre-registrations, participants in Studies 1-7 were excluded for either failing the attention check or (where applicable) for indicating that they did not believe that the financial stakes were real. All key results held when no participants were excluded (see Supplementary Materials). Unless otherwise specified, all conducted t-tests were two-sided. Our studies were approved by Harvard's Committee on the Use of Human Subjects.

Study 1

We recruited 1,039 US American participants online via MTurk. They received \$0.47 in payment for their participation. 144 were excluded for either failing the attention check or for indicating that they did not believe that the financial stakes were real, leaving a final sample of 895 (432 female, 463 male, $M_{age} = 40.5$, $SD_{age} = 12.84$). This study was pre-registered at <https://aspredicted.org/blind.php?x=kr2s85>.

There were five conditions: (i) the favorite-only baseline condition in which participants simply had the option to donate to their favorite charity, (ii) the control condition in which participants could donate exclusively to their favorite charity (all-favorite) or exclusively to the expert-recommended highly effective charity (all-effective), (iii) the three-option bundle condition, which included the option to make a 50/50 bundle donation (favorite-effective) along with the all-favorite and all-effective options, (iv) the two-option bundle condition, which included the 50/50 favorite-effective option and the all-favorite option, and (v) the free-split bundle condition, which included the all-favorite option, the all-effective option, and a favorite-effective option in which participants can freely choose the allocation proportion.

First, participants identified their favorite charity by entering its name and website URL. In all conditions apart from the favorite-only condition, participants were presented with a short and accurate description of one of the world's most effective charities: Evidence Action's Deworm the World Initiative. (See Supplementary Materials.) Participants were also informed that this charity

is, according to experts, approximately one hundred times more effective than typical charities (see <http://GiveWell.org>).

Next, it was explained to participants that they will be given the option to donate up to \$100 to charity and that, at the end of the study, we will randomly select one participant and execute their decision. Participants were then presented with the donation allocation options, which varied by condition. For example, in the three-option bundle condition, they were presented with three options: donate all to the effective charity, split 50/50, donate all to their favorite charity. In the favorite-only condition, participants skipped this part.

On the next page, participants chose a donation amount (\$0 - \$100) to be allocated in the proportion previously specified. Critically, participants had the option to keep the money for themselves. It was explained that any amount not donated will be allocated to the participant in the form of an Amazon voucher, should they be the randomly chosen participant. Thus, this task has real stakes for the donor, and not just for the charities.

After responding to some exploratory questions (see Supplementary Materials), they responded to an attention check question, a question about whether they believed that the financial stakes were real, and to a set of demographic questions.

Study 2

We recruited 294 US participants online via MTurk. They received \$0.90 in payment for their participation. 67 were excluded for either failing the attention check or for indicating that they did not believe that the financial stakes were real, leaving a final sample of 227 people (123 female, 104 male, $M_{age} = 38.84$, $SD_{age} = 12.47$). This study was pre-registered at <https://aspredicted.org/blind.php?x=rn5hg9>.

As in Study 1, participants first entered the name and URL of their favorite charity. On the next page, they were presented with three new charities. All three charities are real charities, and they were described by their names and a short description of their activities. In addition to the basic description, one charity (the effective charity) was said to be one of the world's most cost-effective charities according to experts. Another charity (the popular charity) was said to be a

particularly popular charity, endorsed by many people and organizations. The third charity (the neutral charity) was not described as having an additional positive feature. Which of the three charities was presented as effective, popular, or neutrally was randomized. All three charities were in fact highly effective charities, and when they were presented as effective their effectiveness was described accurately. Likewise, when the charities were described as popular, their popularity was described accurately.

Participants were told that they could donate \$100 to charity and that their decision could have real-world consequences. They were presented with three donation options: 50/50 split between their favorite charity and the charity described as effective, 50/50 split between their favorite charity and the charity described as popular, 50/50 split between their favorite charity and the charity with no additional description (neutral).

On the next page, participants indicated the perceived effectiveness of the three charities, relative to their favorite charity. That is, participants were asked three times, “Which charity do you think is more effective according to experts?” and responded on a 7-point scale (1 [favorite charity] is much more effective, 4 Both charities are equally effective, 7 [alternative charity] is much more effective). Participants also indicated the perceived relative popularity of the three charities compared to their favorite charity. That is, participants were asked three times (once for each bundle type), “Which charity do you think is more popular?” and responded on a 7-point scale (1 [favorite charity] is much more popular, 4 Both charities are equally popular, 7 [alternative charity] is much more popular).

We conducted a follow-up study (see Supplementary Materials) in which we used real charities that differ strongly in their actual levels of effectiveness and popularity, thus enabling us to ensure that the popular charity is perceived as more popular, while the effective charity is once again perceived as more effective. This study yielded the same patterns of results.

Study 3

We recruited 299 US participants online via MTurk. They received \$0.85 in payment for their participation. 67 were excluded for either failing the attention check or for indicating that they

did not believe that the financial stakes were real, leaving a final sample of 242 people (118 female, 124 male, $M_{age} = 40.24$, $SD_{age} = 13.36$). This study was pre-registered at <https://aspredicted.org/blind.php?x=sw7m4z>.

Similar to Study 1, participants first entered the name and URL of their favorite charity. On the next page, participants were presented with a short description of one of the world's most effective charities: Evidence Action's Deworm the World Initiative. Participants were also informed that this charity is, according to experts, approximately one hundred times more effective than typical charities. Next, it was explained to participants that they will be given the option to donate \$100 to charity. They were told that they will be presented with four questions asking how they would like to divide the donation between their favorite charity and the highly effective charity. They were also informed that, at the end of the study, we will randomly select one participant and execute one of their decisions, also chosen at random.

Next, participants were presented with the four decision questions. For each question, participants had to choose between donating exclusively to their favorite charity (all-favorite) or donating with a favorite-effective bundle option (favorite-effective). The bundle options varied for each of the four decision questions. More specifically, we examined the effect of increasing the proportion going to the favorite charity by 30% in the lower range (from 10% to 40%) and in the upper range (from 60% to 90%). This implements a 2 (Range: lower vs upper) x 2 (Within-range proportion: low vs high) within-subjects design. Thus, the four decisions were as follows: (1) All to favorite vs. bundle with 10% favorite / 90% effective (low within lower range), (2) All to favorite vs. bundle with 40% favorite / 60% effective (high within lower range), (3) All to favorite vs. bundle with 60% favorite / 40% effective (low within upper range), (4) All to favorite vs. bundle with 90% favorite / 10% effective (high within upper range).

As indicated above, the proportion going to the favorite charity is always below 50% in the lower range, and the proportion going to the favorite charity is always above 50% in the upper range. The proportion is never 0%, 50%, or 100% to avoid these special values.

Study 4

We recruited 301 US participants online via MTurk. They received \$0.47 in payment for their participation. 36 were excluded for either failing the attention check or for indicating that they did not believe that the financial stakes were real, leaving a final sample of 265 people (94 female, 171 male, $M_{age} = 38.50$, $SD_{age} = 11.69$). This study was pre-registered at <https://aspredicted.org/blind.php?x=3my8rf>.

Study 4 focused on donors who chose the favorite-effective bundle and aims to understand their motivations. We tested the hypothesis that these donors primarily want to support their personal favorite charities, but that the bundle option allows them to satisfy their secondary motivation for effectiveness with relatively little sacrifice of their primary motivation. The study had only one condition, and therefore its statistical methods are descriptive and correlational. Its purpose was to test predictions about how bundle donors perceive their options and how those perceptions predict their overall satisfaction. As in Study 1's primary experimental condition, participants in Study 4 chose among three donation options: all-favorite, all-effective, and a 50/50 favorite-effective bundle.

The materials and procedure were identical to the three-option bundle condition of Study 1. The only differences were the follow-up questions. After participants made their donation decisions, those who donated more than \$0 were asked, (1) how good they feel overall about the donation, (2) how effective they believe the donation will be according to experts, (3) to what extent they feel that they have supported a cause they personally care about, and (4) how fair they think the decision was. Each of these questions was asked three times: once for the donation option that the participant chose (e.g., bundle) and then for the other two donation options (e.g., donating exclusively to their favorite charity or the effective charity). The effectiveness question (#2) asks about effectiveness "according to experts" rather than simply asking participants for their own assessments of effectiveness, as we expected that this to yield a less biased assessment of effectiveness, where "biased" is understood as relative to the

assessments of experts. In the Supplementary Materials, we report an otherwise identical follow-up study in which participants first rate the option attributes and then make their donation choices.

Study 5

We recruited 495 US participants online via MTurk. 101 were excluded because they failed the attention check, leaving a final sample of 394 (177 female, 217 male, $M_{age} = 39.96$, $SD_{age} = 12.77$). They received \$0.50 in payment for their participation. The study was pre-registered at <https://aspredicted.org/blind.php?x=3t7yj8>.

Participants read a hypothetical vignette about three people, all of whom were planning to donate \$100 to charity. They have similar backgrounds and do not know each other. Each feels a strong emotional pull towards supporting a local orphanage housing a small number of children ("Harmony Home"). All three of the soon-to-be donors know that Harmony Home is not particularly effective according to independent charity experts. They have heard about "The Deworm the World Initiative", which is considered by such experts to be one of the world's most effective charities. Each of the three people is unsure which of these two charities to support. One person decides to donate \$100 to Harmony Home and nothing to The Deworm the World Initiative. Another person decides to donate \$100 to The Deworm the World Initiative and nothing to Harmony Home. And the third person decides to donate \$50 to Harmony Home and \$50 to The Deworm the World Initiative (i.e., the bundle option). Participants rated each of these three people on six character traits using a 7-point scale: Caring, Warm, Empathic, Competent, Rational, Effective. The first three items were aggregated to form a Warmth score and the last three items were aggregated to form a Competence score.

Study 6

We recruited 421 US participants online via MTurk. They received \$0.45 in payment for their participation. 62 were excluded for either failing the attention check or for indicating that they did not believe that the financial stakes were real, leaving a final sample of 359 people (158 female, 201 male, $M_{age} = 39.47$, $SD_{age} = 12.25$). This study was pre-registered at <http://aspredicted.org/blind.php?x=xk84hh>.

The study had two between-subjects conditions: matching and control. Apart from minor changes in wording, the materials and procedure of the control condition was identical to the three-option bundle condition of Study 1. In the matching condition, participants were informed that we will match their donations. The matching rate was 5% for donating exclusively to the favorite charity, 25% for choosing the 50/50 bundle, and 50% for donating exclusively to the effective charity. For the bundled donations, matching funds were applied to both charities at the 25% rate. In the control condition, no matching was offered.

Study 7

We recruited 172 US participants online via MTurk. They received \$0.50 in payment for their participation. 27 were excluded for not entering the correct completion code at the end of the survey, leaving a final sample of 145 people (59 female, 86 male, $M_{age} = 35.74$, $SD_{age} = 10.82$).

Participants were informed that they would be directed to a website that allows users to make donations to charity. Participants were told that they could donate any amount between \$0 and \$10 and that we would randomly pick ten participants and execute their decisions. For the selected participants, any portion of the \$10 not donated would be given to the participant in the form of an Amazon voucher. Participants were then directed to a preliminary version of our website, GivingMultiplier.org, and asked to follow the instructions there (see Proof of concept section).

After participants submitted the donation form, they were asked whether they would like to support the matching system by directing the part of their donation that they allocated to an effective charity into our matching system, so that it can be used to incentivize others to make donations to effective charities. It was explained to participants that doing this would allow them to double their impact because a dollar spent on matching indirectly raises two dollars for effective charities (as indicated by the results of Study 6). It was explained that donations directed to the matching system won't be matched but will instead be used to match other people's donations. Participants could choose between the options yes or no.

Proof of concept

We refer to this study as a "proof of concept" because it is not controlled experiment. Given our need to publicly advertise Giving Multiplier and its essential features, we could not randomize prospective donors into a control condition in which Giving Multiplier's essential features are absent. Our aim, then, is not to compare Giving Multiplier's performance to a pre-specified benchmark. Rather, our more modest aim is to determine whether the psychological principles behind Giving Multiplier are at all applicable to the real world.

The proof-of-concept demonstration was pre-registered at <https://aspredicted.org/blind.php?x=jk4j8u>. We promoted the donation website GivingMultiplier.org through word of mouth, primarily on social media and through unpaid media coverage. For example, the website was promoted in articles in the LA Times and Vox.com and in episodes of the podcasts Happiness Lab, Waking Up, and Mindspace. Note that the aim of this proof of concept is not to attract a nationally representative sample, as few charities (if any) have a pool of donors that is representative of the national population. We note that only 27% of our donors indicated that they had heard of their chosen effective charity before visiting Giving Multiplier.

Once on the website, donors were first asked to specify their favorite charity using a search field that contained all legally recognized US charities. Next, they were asked to select one out of nine recommended highly effective charities. Next, they specified the amount they

want to donate (with a minimum of \$10). Next, they were asked to specify how they would like to divide their donation between the two previously selected charities using a slider. The matching rates increased linearly with the proportion allocated to the effective charity. At the time of launch, the rate was 30% for allocating everything to the effective charity, 15% for allocating 50% to the effective charity, and 3% for allocating 10% to the effective charity. (Matching rates were subsequently increased.) Donors had to allocate at least 10% to their selected effective charity.

After donors submitted the donation form, they were asked a set of follow-up questions. First, they were asked whether they had heard of the effective charity before and how much they would have donated to it had they not known about our website. This allows us to more precisely estimate the impact of the website, providing information about which amounts directed to highly effective charities are “counterfactual” (i.e., pass a counterfactual test for incremental value). Second, donors were asked whether they would like to support the matching system by redirecting the part of their donation that they allocated to an effective charity into our matching system. It was explained that these funds could be used to incentivize others to make donations to effective charities. It was explained to donors that doing this would allow them to have a greater impact because a dollar spent on matching indirectly raises more than one dollar for effective charities. It was explained that donations directed to the matching system won’t be matched but will instead be used to match other people’s donations. Donors could choose between the options yes or no.

References

1. Malaria. <https://www.who.int/news-room/fact-sheets/detail/malaria>.
2. Soil-transmitted helminth infections. <https://www.who.int/news-room/fact-sheets/detail/soil-transmitted-helminth-infections>.
3. Schistosomiasis. <https://www.who.int/news-room/fact-sheets/detail/schistosomiasis>.
4. Trachoma. <https://www.who.int/news-room/fact-sheets/detail/trachoma>.
5. Ord, T. The moral imperative toward cost-effectiveness in global health. *Centre for Global Development* (2013).
6. Caviola, L. *et al.* Donors vastly underestimate differences in charities' effectiveness. *Judgm. Decis. Mak.* **15**, 509–516 (2020).
7. Giving USA Foundation. Giving USA 2020: Charitable giving showed solid growth. *Giving USA Report* (2020).
8. Karlan, D. & List, J. A. Does Price Matter in Charitable Giving? Evidence from a Large-Scale Natural Field Experiment. *Am. Econ. Rev.* **97**, 1774–1793 (2007).
9. Small, D. A., Loewenstein, G. & Slovic, P. Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organ. Behav. Hum. Decis. Process.* **102**, 143–153 (2007).
10. Singer, P. *The most good you can do*. (Yale University Press, 2016).
11. MacAskill, W. *Doing good better*. (Guardian Faber Publishing, 2015).
12. Caviola, L., Schubert, S. & Greene, J. D. The Psychology of (In)Effective Altruism. *Trends Cogn. Sci.* (2021) doi:10.1016/j.tics.2021.03.015.
13. Berman, J. Z., Barasch, A., Levine, E. E. & Small, D. A. Impediments to Effective Altruism: The Role of Subjective Preferences in Charitable Giving. *Psychol. Sci.* **29**, 834–844 (2018).
14. Batson, C. D. *Altruism in Humans*. (Oxford University Press, 2011).
15. Bekkers, R. & Wiepking, P. A Literature Review of Empirical Studies of Philanthropy: Eight Mechanisms That Drive Charitable Giving. *Nonprofit and Voluntary Sector Quarterly* **40**, 924–973 (2011).

16. Bloom, P. *Against empathy*. (HarperCollins, 2016).
17. Milkman, K. L., Mazza, M. C., Shu, L. L., Tsay, C.-J. & Bazerman, M. H. Policy bundling to overcome loss aversion: A method for improving legislative outcomes. *Organ. Behav. Hum. Decis. Process.* **117**, 158–167 (2012).
18. Meier, S. Do subsidies increase charitable giving in the Long Run? Matching donations in a field experiment. *SSRN Electron. J.* (2006) doi:10.2139/ssrn.960959.
19. Eckel, C. C. & Grossman, P. J. Do donors care about subsidy type? An experimental study. in *Research in Experimental Economics* 157–175 (Emerald (MCB UP), 2006). doi:10.1016/s0193-2306(06)11007-8.
20. Reeck, C., Wall, D. & Johnson, E. J. Search predicts and changes patience in intertemporal choice. *Proc. Natl. Acad. Sci. U. S. A.* **114**, 11890–11895 (2017).
21. Huang, K., Greene, J. D. & Bazerman, M. Veil-of-ignorance reasoning favors the greater good. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 23989–23995 (2019).
22. Riedl, A. & Smeets, P. Social Preferences and Portfolio Choice. (2017) doi:10.2139/ssrn.2334641.
23. Caviola, L., Schubert, S. & Nemirow, J. The many obstacles to effective giving. *Judgment and Decision Making* **15**, 159 (2020).
24. Sharps, D. L. & Schroeder, J. The preference for distributed helping. *J. Pers. Soc. Psychol.* (2019) doi:10.1037/pspi0000179.
25. Baron, J. & Szymanska, E. Heuristics and biases in charity. *The science of giving: Experimental approaches to the study of charity* 215–235 (2011).
26. Weisz, E. & Cikara, M. Merely Increasing Action Options Increases Charitable Donation. (2020) doi:10.21203/rs.3.rs-59021/v1.
27. Fiske, S. T., Cuddy, A. J. C., Glick, P. & Xu, J. A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *J. Pers. Soc. Psychol.* **82**, 878–902 (2002).

28. Montealegre, A., Bush, L., Moss, D., Pizarro, D. & Jimenez-Leal, W. Does maximizing good make people look bad? *PsyArXiv* (2020) doi:10.31234/osf.io/2zbax.
29. Burum, B., Nowak, M. A. & Hoffman, M. An evolutionary explanation for ineffective altruism. *Nat Hum Behav* **4**, 1245–1257 (2020).
30. Giving Multiplier Impact. <https://givingmultiplier.org/impact>.
31. Nowak, M. A. & Sigmund, K. Evolution of indirect reciprocity. *Nature* **437**, 1291–1298 (2005).
32. Open Philanthropy. <https://www.openphilanthropy.org/>.
33. Giving What We Can. <https://www.givingwhatwecan.org/>.

Acknowledgments

We thank Fabio Kuhn and Daniel Rüthemann for developing the donation platform, every.org for helping with the donation processing, EA Funds for supporting the research, and Fiery Cushman Ben Grodeck, Joshua Lewis and Stefan Schubert for their helpful comments.

Author Contributions

L.C. and J.D.G designed the experiments, L.C. carried out the experiments and statistical analyses, and L.C. and J.D.G. wrote the paper.