

13

Heuristics and Biases in Charity

JONATHAN BARON* and EWA SZYMANSKA

INTRODUCTION

Altruistic behavior often leads to desirable social outcomes. We can thus assume that more altruism is better than less, other things being equal. But altruism tends to be already widely encouraged, so efforts to promote it even further may produce little noticeable change. Instead, it might be easier to do more good by improving efficiency of the altruistic behaviors already in place.

Because altruism is a precious but scarce resource, its wise allocation should be a priority. That is, given the amount of money a person would be willing to give to charity, and the relevant information available, a donor should allocate the donated funds in as efficient a manner as possible.[†] There are several obstacles in the way of achieving such efficiency, however. For example, some donors make decisions that are simply not as well informed as they could be. But an even bigger threat to optimality is posed by overly simplistic decision rules, also referred to as *heuristics* or *cognitive biases*. Biases are systematic errors in thinking that are not necessarily a result of misinformation or ignorance, but rather are brought about by an overgeneralization of some decision rule that might be useful in one context but is ill-suited or even harmful when applied in another. Cognitive biases often lead to the systematic misallocations of funds and waste of resources.

* For conference on experimental approaches to the study of charitable giving, Princeton, July 13–14, 2007. Supported by a grant from the U.S.-Israel Bi-national Science Foundation (Ilana Ritov, co-PI). Email: baron@psych.upenn.edu.

† We focus here on monetary donations, but we acknowledge that altruism also includes many other behaviors, such as donating time. Two other forms of altruism are worth mentioning. One is political action designed to serve the ends of helping others in much the way that charity does. The other is the choice of work. Arguably, people want to contribute to the well-being of others through their work, to the point where they are willing to accept lower salaries in order to do jobs that serve this end, e.g., being a public defender instead of a tax lawyer. These other competing modes of altruism are important and should not be ignored. We shall return to this point, especially with respect to political behavior.

We propose utilitarianism, or the totality of good that comes about from a choice, as our gold standard for assessing the effectiveness of choices made in the context of charitable contributions. The utilitarian model we apply does not attempt to tell people how altruistic they “should be,” but simply accepts their willingness to donate as is. (Later, we discuss extensions of this model.) From a utilitarian perspective, the best measure of efficiency is the amount of good done per some monetary unit (e.g., a dollar) donated. For example, we ask how many lives are saved for how many dollars? In what follows, we discuss several examples of how people’s contribution decisions may fail to meet this standard of economic efficiency.

We next present our own empirical findings, which further demonstrate the systematic deviations from the predictions of a normative model. Finally, we return to the subject of utilitarianism, and we introduce the theoretical foundations of the efficiency standard we have proposed. We also provide an overview of the other possible versions of the utilitarian model, and we comment on ways in which people could be assisted in their selection of the optimally altruistic choices.

In sum, our general perspective is that optimal altruism should be viewed as a value maximization problem. In the upcoming section, we examine the most common decision-making errors that stand in the way of maximizing efficiency of donations. In the next section, we ask how people may depart from this standard of efficiency. We believe that, if we can correct the cognitive biases or work around them, then we can improve the efficiency of altruistic behavior and thus do more good for the sacrifices that people are willing to make.

Possible Nonutilitarian Heuristics

Evaluability Hsee (1996) argued that people pay more attention to attributes that are easy to evaluate. He manipulated the ease of evaluation by presenting unfamiliar information with or without a comparative context. For example, imagine being asked how much you would pay for a music dictionary with a torn cover and 50,000 entries. The cover is easy to evaluate, and torn covers are not a desirable feature to most buyers. But unless music theory is your profession or a hobby, you are probably not sure how to evaluate the number of entries. Consequently, you would likely pay considerable attention to the condition of a book, and would not be willing to pay much for the dictionary with a torn cover. On the other hand, if you were simultaneously presented with another dictionary that has an intact cover but only 25,000 entries, you would attend more to the number of entries and likely pay more for the 50,000 dictionary, despite its torn cover.

In the case of donating to charities, donors often correctly wonder about efficiency. As we have noted, the best measure of efficiency from a utilitarian perspective is the amount of welfare-enhancing results. But this dimension is often difficult to evaluate. Few organizations devoted to saving lives bother to even try to calculate just how many lives are saved for how many dollars. Fewer publicize such figures. And when life saving is not the only goal, the task of measuring benefits becomes more difficult still. (For health outcomes, it is possible to estimate benefits in terms of “quality adjusted life years,” as many agencies now do (see Baron, 2008).

Such analyses are used by governments and insurance companies, but not, to our knowledge, by charitable organizations.)

Instead, what is more evaluable than the lives saved per dollar of contribution is the operating cost per dollar. How much money is spent on administration, advertising, and other functions not related directly to a charity's main function? Ratings are available for this sort of efficiency. For example, the Charity Navigator site providing online rankings of over 5,400 American charities (see www.charitynavigator.org), publishes lists such as "10 Inefficient Fundraisers" or "10 Charities Overpaying Their For-Profit Fundraisers." It is, however, unclear that focusing on the operating costs maximizes the sort of efficiency that ought to matter, namely, the amount of good done per dollar spent. For example, Charity A may spend 10% of donors' contributions on operating costs, but need \$10 to save a single life. Charity B may have higher operating expenses, using 25% of donations to cover the costs, but require as little as \$5 to save a human life. In this scenario, the Charity Navigator would reliably single out the percentage of resources used to pay for the operating costs, thus suggesting to the potential donors that contributing to Charity A would be the more efficient use of their money. However, this is evidently not the case. Suppose that charities A and B both receive \$100. Charity A spends \$10 on overhead and saves 9 human lives with the remaining \$90. Charity B, on the other hand, allocates \$25 toward its operating expenses, but manages to save 15 lives with the remaining \$75. Clearly, Charity B is generating significantly more welfare with the \$100 donation than Charity A.

Of course we are not arguing that keeping the costs down is irrelevant. It is undeniably of great importance to manage the expenses wisely and to ensure that operating costs are not exceeding acceptable limits. After all, each dollar cost reduction frees up funds for the actual charitable programs. But, as in the case of judging a music dictionary by its torn cover, if we look at the operating costs in isolation, we might not be making allocation decisions that are as good as we might think we are. An efficient organization should manage costs wisely, while carefully watching the bottom line. In the case of charities, the bottom line is an increase in the economic welfare or the number of lives saved. The bottom line ought to matter to donors, but it is a number that is hard to obtain. Instead, donors, reinforced by third-party evaluators such as the Charity Navigator, base their allocation decisions on easy-to-evaluate information about operating costs.

Average vs. Marginal Benefit We believe that, if the information about the average benefit per dollar was made available and presented in a way that allows for meaningful comparisons, people would likely attend to this important statistic. In the case of charitable contributions, knowing and appropriately taking into account the average benefit per dollar is a definite improvement in the decision-making process. But what is even more relevant to an optimal allocation of funds is the *marginal benefit per marginal dollar*. The marginal benefit per marginal dollar is the benefit-to-cost ratio for every *new* contribution. For example, a popular charity might be relatively well funded, so that the marginal benefit per dollar of a new donation is much lower than the average benefit. In other words, the organization may lack good ways to use additional money. On the other hand, a small

organization that has not yet reached its target funding level might be able to make better use of additional funds, even if its average benefit-to-cost ratio is lower for all the funds it has. This might be the most difficult statistic to obtain, however, as it is not in the well-funded charity's best interest to advertise the ratio.

Diversification In general, variety seeking is a good thing to do for several reasons. For goods that involve experiences, like listening to music or watching movies, variety reduces the effects of adaptation to repeated presentations of the same thing. For investments, diversification reduces risk. But people continue to seek variety also when they have no reason to do so. In fact, as we are going to demonstrate, people apply the *diversification heuristic* (Read & Lowewenstein, 1995; Fox, Ratner, & Lieb, 2005) even to charitable giving.

Diversification in charitable giving means that an individual donor prefers to contribute money to many different charities, instead of allocating all the funds to a single, most efficient organization. It may extend to the idea of giving equal contributions to several charities (Fox et al., 2005).

Diversification has many possible explanations. For example, people may fail to understand the relation between allocation and maximization. To illustrate, Ubel, DeKay, Baron, and Asch (1996) told subjects (people waiting for jury duty), "All else equal, kidney transplants are more likely to succeed if they are transplanted into patients who 'match' the kidney donor. The number of successful transplants one year after transplant is eight out of ten for patients with a *complete match* and seven out of ten for patients with a *partial match*." The subjects were asked to allocate 100 available kidneys among two groups: 100 patients with a complete match (hence 80% success rate) and 100 patients with a partial match (70% success rate). Most subjects allocated 50 kidneys to each group. When asked how to maximize the success rate, fewer than 20% of the subjects gave the correct answer (all 100 to the 80% success group), and most of these subjects still allocated the kidneys equally. Thus, people may diversify because they allocate so as to match rather than maximize effectiveness.

Another reason for giving to many different charities, instead of just the best one, might be the desire for a "warm glow," (Andreoni, 1990) or a positive feeling that arises from every act of helping each organization. The magnitude of the glow may be roughly constant for each act of contributing (Margolis, 1982). Of course this is a self-interested motivation, although not necessarily a bad one if it increases altruism.

Diversification might also be explained in terms of the donor's inclination to think of charities like investments. For the average donor, they are not the same. The difference can be understood in terms of the principle of *declining marginal utility*. According to this well-known economic principle, when an individual, or a charitable organization, receives more of something, each extra unit of the added thing contributes progressively less and less utility (does less and less good). The principle applies to money as well as it does to, for example, ice cream (would you be as excited about a tenth consecutive scoop of ice cream on your plate as you were about the first, the second, or even the third?). For investments, given a fixed expected level of income, it is more important to avoid losses from that level than it is to obtain gains of the same amount, because the losses have a greater effect on total utility (i.e., the pain after losing some amount is greater than the pleasure after

AU: Add page number for quote here.

gaining the same amount). Thus, it pays to reduce one's risk by spreading out investments, so that if one investment does badly, the investor does not lose too much.

The marginal benefit per dollar (recall the difference between the average vs. marginal benefit) of some support is usually much higher than the marginal benefit of support in addition to some other previous support. In the context of charitable donations, however, a single donor does not have much influence over the total level of the charity's income (unless the donor is a major philanthropist). Thus, the usual arguments that favor diversification in investments do not necessarily apply to donations by typical donors. For charities, the focus should be on maximizing the expected increase in the economic welfare per dollar of every contribution.

In other words, a donor may err by thinking and acting as if she were the only donor to the charities that she supports. If she were, then it would make sense to diversify because of the principle of declining marginal utility just described. In reality, however, more often than not, each contributor has only a tiny effect on the total funding required by a charitable organization, so the argument based on the principle of declining marginal utility is unlikely to apply.

We should, therefore, identify a charity with the highest expected benefit per dollar, and make our entire contribution to that one charity. Of course, this is a risk. We may be wrong. But if people were willing to take the risk, more good would be generated in the long run.

For donors of small amounts of money, diversification has another negative feature, which is that charities spend money on mailings. As a consequence of this fact, and the arguments made earlier, the first author resolved last year to donate \$10 each to most of the charities to which he had given money for years, telling them that this was the last contribution and that he should be removed from their mailing lists. Unfortunately, he could not bring himself to follow his own advice and narrow it down to one, or two, or three. The diversification intuition is strong.

Prominence People tend to pay attention to a single prominent attribute, or an attribute they view as the most important (Tversky, Sattath, & Slovic, 1988), to the exclusion of all other attributes they deem as less important. For example, people may care first and foremost about reducing the number of casualties on the highways. This is a perfectly valid goal; however, treating the number of casualties as a strictly prominent attribute makes wise decision trade-offs very difficult. If two safety programs differ in cost and in the number of lives saved, people tend to choose an option that saves more people. They do this even if the difference in lives is small and the difference in cost is large. The prominence effect relates to the lexicographic choice strategy, in which the option with "the best value on the most important attribute is simply selected" (Bettman, Luce, & Payne, 1998).

The trouble with basing decisions on prominence is that sometimes the prominent attribute should be outweighed by large differences in other attributes. Accepting large surges in cost in exchange for only a small increase in goal achievement represents yet another bias leading to the inefficient allocations.

Parochialism Parochialism, formally defined, is an in-group bias in which people weigh the welfare of their own group more heavily than those of outsiders,

AU: Add page number for quote.

to the point where they favor policies that help their group but do greater harm to outsiders. Charity will, we assume, never do this; it will always do more good than harm. But it can be inefficient if the out-group charities are more efficient than the in-group ones, yet receive less funding simply because they provide aid to the outsiders. A situation in which the out-group charities are able to stretch a dollar further than the organizations helping fellow citizens is often the case in the United States, as the cost of life saving in, for example, Africa is far lower than that in the United States.

If we assume that people would not reallocate their contributions to the more efficient external charity because their altruism is limited to insiders, then we cannot consider parochialism a bias, as any attempt at encouraging the reallocation would bring about more harm than good. This argument depends on the assumption that parochialism is a result of some strong human intuition, much like our deeply ingrained commitment to close family members. In other words, parochialism may be seen as analogous to family loyalty, which can be explained in terms of a biological adaptation aimed at increasing the individual's fitness indirectly—through assisting one's close relatives. However, it is not clear how the biological tendency to help the family members would generalize to our commitment to, for example, strangers in Hawaii who suffered from a tsunami. But given that the U.S. citizens do seem to experience a sense of commitment to tsunami victims in Hawaii, we have grounds to expect that the same cognitive mechanism may affect their commitment to New Zealanders who suffered the same fate. This assumption, in turn, suggests that parochialism results from cognitive illusions that might be corrected.

AU: Update available?

Baron (in press) argues that parochialism is greater for harms of inaction than for action. That is, people are reluctant to harm outsiders through acts but much more willing to do nothing to help them. If people can be brought to see this as a kind of inconsistency, perhaps there will be more obligation to help outsiders. Failure to give to charity is, in fact, an omission that results in great harm.

Parochialism can also be reduced by encouraging people to think of outsiders as individuals rather than as members of an abstract group. Parochialism is at least in part the result of thinking abstractly about people in terms of, for example, “nations” and “the enemy.”^o

AU: Add Schick 1991 to ref list.

If the option of modifying parochialism is open either to the decision maker or to others, then in-group bias in charity could be counted as a bias against even the sophisticated utilitarian view. In particular, the view that takes moral education

^o Schick (1991, p. 1) tells the following story: “Writing about his experiences in the Spanish Civil War, George Orwell tells this story. He had gone out to a spot near the Fascist trenches from which he thought he might snipe at someone. He waited a long time without any luck. None of the enemy made an appearance. Then, at last, some disturbance took place, much shouting and blowing of whistles followed, and a man: jumped out of the trench and ran along the parapet in full view. He was half-dressed and was holding up his trousers with both hands as he ran. I refrained from shooting at him. I did not shoot part because of that detail about the trousers. I had come here to shoot at ‘Fascists’; but a man holding up his trousers isn’t a ‘Fascist’, he is visibly a fellow-creature, similar to yourself, and you don’t feel like shooting at him.”

costs into account might be relevant, in that the costs of reducing in-group bias might be relatively low.

Identifiability We shall not examine the identifiability effect in the studies we report; however, since four articles in this volume discuss the finding that altruism is greater when single victims are identified (Loewenstein & Small; Kogut & Ritov; Ritov & Kogut; Slovic), we will briefly discuss how the identifiability effect violates the efficiency standards we proposed earlier. Recently, Charness and Gneezy (2008) found that “giving” in a dictator game (where one subject has control of some money and can give some to another subject who has none) was higher when the recipient was identified by last name. This effect results in part from the fact that the identified victim evokes more empathy.

Is this a bias? From the simple utilitarian perspective, it is. By increasing altruism toward the identifiable victims, it almost necessarily reduces altruism toward the unidentified ones, who are often the ones most in need of help (because they tend to be masses of people who are undergoing some great deprivation). Perhaps, though, the emotion changes the benefit/cost threshold for altruism, meaning that people are willing to incur a greater personal cost of helping, even when the amount of benefit to others does not change. This shift in the benefit/cost threshold might possibly even increase overall altruism.

On the other hand, it is clearly a kind of framing effect, in which choices are affected by how they are presented, holding the reality constant. Victims all have names. The fact that we are aware of one of them is an accident. We could make up names for the others, or even tell ourselves that our donation to some relief fund is going to help someone named Zhang. Thus, the identifiability effect may fail a simpler test of rationality.

Voluntary Versus Tax A final violation of the efficiency standards we discuss is that people tend to oppose the government aid programs supported by taxes—often referred to as “forced charity.” Baron (1997) argued that contributions (of time or money) to political causes that support helpful programs could sometimes be more efficient than direct contributions to the beneficiaries of the proposed programs. This is because although the probability of having some effect through political action is very low, the benefits of a successful initiative are potentially very high. In economics, the expected utility, or the anticipated amount of generated welfare, equals the size of the possible benefits times the probability of obtaining those benefits. If we apply this formula separately to the two modes of helping—an individual donation and a political action—the expected utility of the two modes might be similar, sometimes higher for individual contributions but sometimes higher for political action, depending on the details. After all, little benefit times high probability of making at least a small difference (as in an individual donation), might well be roughly the same as a great benefit times a small probability of doing any good at all (as in political action). Moreover, the cost of political action might be so much lower than the cost of direct contribution that it might be required by the efficiency standards to combine political action with direct donation.

In part, the bias against “forced charity” may arise from a belief in freedom, the belief that government should not force us to help others but should, more or less, provide us with services from which we all benefit and pay for collectively, such as roads, military defense, and protection of our property. (Some libertarians would not even go that far.) Insofar as this is true, it may represent a kind of cognitive inconsistency. Some people benefit very little from roads or property protection, so paying taxes for these things is a way of forcing them to sacrifice for the benefit of others. It is a matter of degree.

EXPERIMENTS

Here we report the results of four experiments designed to look for evidence of some of the biases just described.*

All studies were done as questionnaires on the World Wide Web. Subjects were part of a panel of about 1,500. We sent emails to 500 at a time, indicating the availability of a study and giving the URL. We paid \$3 (to be added to earnings from other studies and then paid through PayPal) for each study. We aimed for 80 responses to each study and then withdrew it. The subjects were mostly Americans. Their median income and education was typical of the U.S. population. Their ages ranged from 20 to 80, with a median of about 42 (varying from study to study). About 80% were female.

Each study began with an introductory page, which was followed by 14 to 16 additional pages, each with one item and 4 to 5 questions following the item. The full sets of questions for each study are shown in Appendix A, with one item used for illustration for each set. The introductory page was similar for all the studies reported, and read as follows:

The first part (7 pages, 4 questions each) is about personal contributions to worthy causes, such as helping people in need who cannot help themselves, people such as young children or those who are ill. Imagine that you are not one of these needy people and that, in fact, you have enough money so that it is easy for you to give some away for such causes without seriously hurting your quality of life. Imagine that you are willing to contribute some of your annual income to such causes.

The second part (7 pages, 4 questions each) is about your evaluation of programs to cure and prevent serious diseases that affect children, such as AIDS, malaria, worms, polio, and whooping cough. Many children in Africa suffer from these diseases and often die. But pockets of curable and preventable childhood diseases exist in rich countries too. Other questions in this part concern health insurance.

Suppose that all of these programs are effective and efficient. We are interested in your views about how such programs should be funded.

Here we report results according to the issues addressed, identifying the study that produced them.†

* The experiments themselves are available at <http://finzi.psych.upenn.edu/~baron/ex/ewas/>, as char1, char2, char3, and char4.

† The data and R scripts used for analysis are at <http://finzi.psych.upenn.edu/~baron/R/ewas/>.

Waste

The following item was included in our first study to test the idea that subjects will focus on waste even with efficiency held constant (recall our discussion of the evaluability heuristic, and the donor's preoccupation with the operating costs, to the exclusion of other relevant information): "A and B help prevent deaths in children. Both of them can prevent 5 deaths for every \$1,000 of donations. A spends \$200 out of every \$1,000 of donations on advertising. B spends \$100." (Note that despite higher operating costs, B generates exactly the same amount of welfare as A). Subjects were asked, "How much would you allocate to A/B?" The mean response was 75% to A, which was significantly higher ($p = .0000$, rounded) than the 50% we might expect of an unbiased responder, given equal efficiency.

In the second study, we put "of donations" in bold font; but again, the mean allocation was 68%, still clearly different from 50%.

Since subjects might have thought that the \$1,000 referred to what was left after advertising, in the third study, we tried to make the distinction clearer still: "A and B help prevent deaths in children. Both of them can prevent 5 deaths for every \$1,000 of donations. A spends \$200 out of every \$1,000 of donations on overhead expenses, but manages to save 5 lives with the remaining \$800. B spends \$100 out of every \$1,000 on overhead, and saves 5 lives with the remaining \$900." The mean response was 3.38 on a 1–5 scale in which 1 meant "all to A" and 5 meant "all to B" ($t_{83} = 3.39$, $p = 0.001$, vs. 3, which would mean equal allocation).

In sum, waste (or the operating cost) matters to people, even in cases when the reason for worrying about waste—the resulting lack of efficiency—is absent.

Average cost

We can similarly ask whether subjects attend to average benefit per dollar even when the marginal benefit turns out to be equal. In our first study about this, we asked, "A and B help prevent deaths in children. A prevents 5 deaths for every \$1,000 of donations, on the average, and B prevents 6 deaths for every \$1,000. Given the donations they have received so far, and the opportunities for expansion, A will prevent 5 deaths for each additional \$1,000 beyond its current level of spending and B will also prevent 5 deaths." The mean allocation to B was 72% (again, significantly higher than 50%). Subjects apparently attended to the average benefit even when the marginal benefit was the same.

In a subsequent study, we asked, "A and B will each prevent 5 deaths for every \$10,000 of new donations. A was much more expensive to get started. Thus, the cost per life saved on the average is higher for A, because A has spent more money in total." The mean allocation to B was 61%, still higher than 50% ($t_{76} = 3.66$, $p = 0.0005$).

As in the case of waste, subjects attend to a dimension that would be relevant, average benefit per dollar, except that the reason for its relevance is now absent.

Diversification

We examined the diversification heuristic using multiple approaches.

Unequal Efficiency We find that people allocate money to the less efficient charity just for the sake of diversifying. In the first study we stated, “A can save one life for \$10,000. B can save one life for \$12,500. The people helped are from the same groups, with the same problems.” On average, subjects allocated 18% to B, the less efficient charity; 40% of the subjects allocated at least some proportion of funding to B.

In the second study we stated, “A can save 5 lives for \$50,000. B can save 4 lives for \$50,000. The people helped are from the same groups, with the same problems.” Here, the mean allocation to B was 14%, with 38% of the subjects allocating more than nothing.

In the same study we stated, “A and B are both involved in preventing death in people with AIDS. A uses a method with a 75% chance of success over 5 years. B uses a method with a 50% chance of success over 5 years, with the same patients.” Here, the mean allocation to B was 12%, with 43% allocating more than nothing.

In the fourth study, we included three relevant items: “A and B are both involved in preventing death in people with AIDS. A uses a method with a 75% chance of success over 5 years. B uses a method with a 50% chance of success over 5 years, with the same patients”; “A can save one life for \$10,000. B can save one life for \$12,500. The people helped are from the same groups, with the same problems”; “A can save 5 lives for \$50,000. B can save 4 lives for \$50,000. The people helped are from the same groups, with the same problems.” We consider these items together, as they did not differ in any meaningful way. Each item was followed by several questions, given here with the responses to each. In all cases, the options were as follows:

1. All to A
2. More to A, some to B
3. Equally to A and B
4. More to B, some to A
5. All to B.

“What is the right allocation between A and B, ignoring your own feelings?”

Mean 1.46 (that is, between 1 and 2 on the above scale), 28% of subjects allocated some to B.

“What allocation would you feel best about making?”

Mean 1.48, 29% of subjects allocated some to B.

“What allocation between A and B would be the most efficient use of your money?”

Mean 1.33, 21% of subjects allocated some to B.

“What between A and B would do the most good for each \$1,000 spent?”

Mean 1.33, 22% of subjects allocated some to B.

Note that the responses to “right allocation” and “feel best” were more equal than those to the last two questions, which were about efficiency. The difference was significant ($t_{77} = 3.22$, $p = 0.0019$, for the means; $t_{77} = 3.61$, $p = 0.0006$, for the proportion of responses in which something was allocated to B).

In sum, as found by Ubel et al. (1996) in a different situation involving allocation of organs for transplantation to two groups, one group with a higher success rate than the other, people do not allocate so as to achieve the most efficient (or the utilitarian best) outcome. That is, people fail to allocate organs first to the group with the higher success rate. Ubel et al. found two reasons why some people systematically fail to maximize efficiency. Both of these are consistent with our results. Some subjects thought that some allocation to the group with a lower success rate would in fact maximize the outcome. Others understood how to maximize but thought that it was unfair to give nothing to the group with the lower success rate.

Unequal Efficiency, Several Projects Versus One We also looked at the bias toward supporting multiple projects, at the expense of overall efficiency.

In the first study we asked, “A puts all the money into one project, which has a 75% chance of helping many children, and a 25% chance of doing no good at all. B puts the money into several different projects, each of which has a 70% chance of helping some children, but a 30% chance of doing no good.” The expected utility of A is higher than B ($75\% > 70\%$), and so, our goal of maximizing welfare clearly implies allocating *all* funds to A. But 32% (vs. the optimal 0%) of the allocations went to B, and 60% of the subjects allocated something to B.

In the third study, we attempted to be more specific: “A puts \$1,000,000 into one project, which has a 75% chance of helping 10,000 children, and a 25% chance of doing no good. B puts \$200,000 into each of 5 projects (\$1,000,000 total). Each of the 5 has a 70% chance of helping 2,000 children and a 30% chance of doing no good. (If all 5 succeed, then the total benefit is 10,000 children, the same as A.)” The mean allocation was 2.73 on 1–5 scale where 1 is all to A, 5 is all to B; 81% of the subjects allocated some funds to B.

We asked the same question in our last study, and the mean allocation was 2.46 (with 75% allocating some to B). In this case, the allocation judgment was not significantly higher than the efficiency judgment (2.39), meaning that people actually believed the efficiency of both charities to be roughly equal. Thus, unlike in the previous case of two charities, concern for fairness leads people to go against what they think is most efficient. They just get that incorrect.

(Note that these results are unaffected by the number of groups benefiting from the projects. When the projects are equated for efficiency, their number has no effect on allocation decisions. Specifically, in the first two studies, we asked, “A puts all the money into one project, which will help 100,000 children. B puts the money into five different projects, each of which will help 20,000 children. [The benefit per child will be the same.]” The mean allocation to A was 49% and 48% in the two respective studies, neither of which differed from equal allocation.)

Nationalism

We asked some simple questions about nationalism. In the first study we asked, “A helps children who are in your own country. B helps children around the world. The children are equally needy.” The mean allocation was 33% to B, where 50% would be equal allocation. In a follow-up study we added, for emphasis, “and the benefits are the

same for each child.” The mean was unchanged at 32% to B. We also specified what countries the children supported by B were in, using India, Africa, and Latin America. The mean allocation to B was 30% overall, and the regions did not differ. Subjects had an opportunity to comment on each item. A typical comment was, “There are just as many needy children in THIS country and I would help them FIRST.”

In our third study, we tried to make it clear that the children were all strangers: “A cures a disease in children who are in a distant part of your own country. B cures the same disease in children in India. A and B are equally efficient. You do not know any of the affected children, or any children who have had this disease.” National bias was still present. The mean allocation was 2.55 (1 = all to own country, 5 = all to other country; $p = .0000$ compared to 3.00, i.e., equal allocation). The same held for Eastern Europe, China, and Africa, and the regions did not differ in allocation.

However, 58% of the subjects had means of 3 (equal allocation) or higher. One of these subjects commented, “Children are children, no matter where they are from.” Thus when the need and benefits are made explicit, some people do not feel a duty to help their co-nationals first.

Forced Charity

Comments on early studies suggested a general antipathy to any sort of compulsory assessment for charitable purposes:

“Nobody has the right, weather [sic] it is voted upon or not, can force employees to donate money.”

“While I believe in helping others, have a hard time being forced to help.”

“No matter how important a program is we cannot give up freedom of choice.”

“The citizen should have the choice as to whether he wants to contribute or not.”

“I would not like being forced to pay for any—it should be your choice.”

In our fourth study, we presented the following scenario:

Your country requires everyone to buy health insurance. The fee is fixed at about \$2,500 per person.

Case A: To help those who have trouble paying, the government levies a special income tax. Families earning less than \$7,500 per person pay no tax. Others pay a fixed percentage of their income above \$7,500/person.

Case B: To help those who have trouble paying, charities collect **voluntary donations**. The charities distribute the funds to try to come as close as possible to the situation described in Case A. That is, they provide a subsidy for families earning less than \$7,500/person, they reduce the subsidy gradually as income increases, and they solicit contributions from rich people who would pay less in Case B than in Case A. Suppose that the charities succeed, so that the bottom line is the same for each family as in Case A.

We asked four questions:

1. Which case would you favor if you had a choice? [A, B, equal]
2. Which case is more fair in distributing the cost and benefits?

3. Which case provides more freedom of choice?
4. "Which is more important in this case? [fair distribution, freedom of choice, equal]."

Answers to the questions were highly correlated. The mean response to them was 0.12 (-1 = forced, 1 = voluntary), which did not differ significantly from zero ($p = .09$). On the whole, in this case, there was no particular bias against forced charity. Another similar case yielded similar results. However the following case yielded a small but significant bias:

Workers in your country who make widgets (imaginary goods) are getting lower wages because of competition from foreign imports. The price of widgets has gone down, and the workers have accepted wage cuts to avoid layoffs.

Case A: The government puts a tax on widgets. The proceeds from the tax are used to help the domestic workers by restoring their wages to their original level.

Case B: A voluntary charity collects funds to help the domestic workers. The funds are sufficient to restore their wages to their original level.

A second version of Case B read: "A campaign to induce buyers to buy domestic widgets voluntarily succeeds in allowing their price to go up. Wages go back up to their original level."

For both versions, subjects favored the voluntary mechanism, with means of .17 ($t_{77} = 2.84$, $p = 0.0057$) and .20 ($t_{77} = 5.96$, $p = 0.0000$), respectively. But 33% and 26%, respectively, favored the tax. This contrasts with another case:

A **new** epidemic disease threatens to infect 10% and kill 1% of the children in **your nation**. It can be cured by a treatment that costs too much for any insurance company to cover it, including current government insurance, without raising premiums or taxes. Additional funds are needed.

Case A: The funds are raised with an increase of income taxes by a percentage of the current tax, sufficient to cover the costs.

Case B: The funds are raised by voluntary donations. Treatment is free to those who get the disease. Donations are sufficient to cover the costs.

Here the mean was .37, and only 14% favored the tax.

In sum, we have some evidence for a labile preference for voluntary mechanisms. Of course, in stating these cases, we have asked subjects to accept the implausible assumption that voluntary mechanisms can raise as much as taxes.

DISCUSSION

Whether a charitable donation is inspired primarily by altruism, reputational concerns, psychological benefits to the self, or a combination of other motivational factors,[°] we argue that a decision to financially self-sacrifice should be accompanied by a straightforward goal of maximizing the total expected utilities of others. Given the enormous economic needs in developing nations coupled with the relatively small costs of saving human lives in the poorest countries (World Bank, 2004), maximizing

AU: Add to ref list.

[°] For a comprehensive review of the reasons why people donate, see Bekkers and Wiepking (2007).

the utility of each donation is of crucial importance. We recognize that one problem with utility maximization in charitable giving is the current lack of adequate measures to assess the efficacy of philanthropic organizations worldwide (World Bank, 2007). What concerns us more, however, are our findings from a series of studies on charitable donations where the subjects made systematically suboptimal allocations despite being provided with clear diagnostic indicators of the charitable organizations' performance.

AU: Add to ref list.

The most pronounced deviations from the utilitarian standards included: (1) the waste heuristic, or preoccupation with the overhead expenses that led subjects to favor less efficient charities with low operating costs over more efficient organizations with slightly higher operating costs; (2) average cost bias, or attention to average benefit per dollar when, in fact, the marginal benefit should guide a decision; (3) the diversification heuristic, or a tendency to donate to many charities even in cases when contributing to only one organization would save more lives; (4) parochialism, or support for organizations helping co-nationals even when donating to international causes would help more people; and finally (5) a framing effect, or preference for either a system of voluntary donations or forced charity (i.e., taxes) depended on the type of presentation.

One idea for helping people to generate the most good with the resources they are willing to donate is to encourage charities to disclose information about their current state of funding by regularly updating the benefit-to-cost ratio for new contributions (i.e., the ratio of benefit to beneficiaries over cost to a donor). Such a single indicator of efficacy, if calculated properly, would enable donors to select the most efficient charities at the time of contribution. A skeptic might argue that the underperforming charities, as well as the organizations with ample funding, would have little or no incentive to publish their benefit-to-cost ratios. A possible solution may include educating donors about the biases that subvert the cost-effectiveness of their decisions. We doubt whether education in and of itself would be capable of eliminating the common biases. Nevertheless, education might awaken public demand for access to standardized information on charities, turning a disclosure of the benefit-to-cost indicator into a norm.

More beneficial still might be maximization of the efficacy of tax-supported initiatives, accompanied by an increase in public support for collective action. Sachs (2005) argues that even though "introductory economics textbooks preach individualism and decentralized markets, our safety and prosperity depend at least as much on collective decisions to fight disease, promote good science and widespread education, provide critical infrastructure, and act in unison to help the poorest of the poor." In other words, the most efficient use of the human capacity for caring may not be an annual charitable donation, but rather support for a broader social movement or an engagement in social entrepreneurship.

AU: Add page for Sachs quote.

Utilitarian Models of Altruism

The utilitarian model that we have applied, assuming a fixed willingness to sacrifice self-interest, is simplistic. Future research should consider extensions of this model, or alternatives to it, even within the utilitarian framework.

Maximize Total Utility The simplest utilitarian model for decision making requires that all choices maximize the welfare of everyone affected by them. The model provides a measure of the total expected goodness of consequences and holds that more self-sacrifice is always better, as long as the individual costs of sacrificing are less than the benefits provided to others. Such a simple utilitarian formula is especially useful in the context of third-party decision making, in which those who decide remain unaffected by their decisions. However, it is less helpful in understanding and improving individual decisions, since in real life, people find it difficult to sacrifice their self-interest.

An alternative normative model evaluates the goodness of consequences of different options (i.e., their utility) in terms of the decision maker's own goals. These goals include altruistic goals, but the welfare of other's carries less weight than the well-being of a decision maker.

Thus, those who want to know how rational their decisions are can get two different answers, depending on which model they apply—simple utilitarianism or expected-utility based on the personal goals.

Limited Self-Sacrifice We can think of self-sacrifice as a limited resource, subject to an external constraint. In the simplest form of this model, each person begins with a simple expected utility approach and then reduces personal expected utility by some fixed limited amount, so long as this reduction leads to greater benefit for others, for each unit of time.

For example, suppose that we consider only money. People work to make money. They have some utility function for annual income; increases in income have utility for them. Presumably this utility is marginally declining with income, e.g., their utility in “utiles” is the square root of thousands of dollars. Then, having made their money, they consider how much to give to charity. They are willing to sacrifice, say, 1 utile, period. So, someone who makes \$64,000 is willing to sacrifice \$15,000, because 1 utile is the difference between the utility of \$64,000 (8 utiles) and the utility of \$49,000 (7 utiles). This form of the model allows us to consider questions of efficiency very easily. We ask, given the limitation, how we can best maximize the utilities of others. For example, what charitable causes will do the most good, for \$15,000? This model is basically what we have applied here, although our results are relevant to some other models too.

An advantage of this model is that it is compatible with a decision maker's point of view. Instead of asking an abstract question, “What is the most rational thing to do?” a decision maker can ask, “How can I do the most good given the limited self-sacrifice that I am willing to make?” The latter question assumes that the level of self-sacrifice is constrained and not solely under the individual's control, which is, in fact, how many people may view it.

It is also possible to reframe the model, and instead of focusing on the limited ability to sacrifice, take into account the generated benefits for others. For example, people might be willing to reduce their own welfare if the gain for others is greater by some constant multiplier, such as 10. That is, they would be following a rule of thumb that if you can give up 1 utile in order to increase 11 people's utility by 1 utile each, then you would do it.

AU: Add to list.

This form corresponds to proposals for *weighted utilitarianism* that have often been discussed (e.g., Baron, 1993). The difference between the present approach and the earlier proposals is that they were viewed as recommendations made to an individual—almost a way of capturing some concept of “duty” within a utilitarian framework. It is as if they said, “If you always sacrifice your own good when you can do 10 times as much good for others, then you have done your duty and you don’t need to feel guilty anymore.” The trouble with this is that it doesn’t really answer any question. Not “what is best?” and not “what is rational?” Those are already answered in other ways: what is best is to maximize utility for all, which could require more self-sacrifice than most people are willing to make; and what is rational (arguably) is to best achieve your own goals, including your goal of being altruistic. Our idea of an external limitation on altruism is an answer to the second question. Our normative model does not attempt to tell people how altruistic they should be, but it simply accepts their willingness to sacrifice as it is.

A more realistic form of this model is that self-sacrifice is indeed limited but can be “stretched.” As the amount of self-sacrifice increases (per unit time), the benefit-to-cost ratio—the ratio of benefit to others over cost to the self—must be higher before a self-sacrificing option is chosen. This model specifies the same optimal behavior in our experiments as the basic model.

Limited Altruism We could extend the model further to allow altruism to depend on the beneficiaries. In the most general case, the function implied in the last case would depend on who benefits. That is, we would behave as if we multiplied the marginal benefit to person P by some factor (usually less than 1) that depended both on who P is and on the marginal cost to us of providing the marginal benefit (with lower weight for greater cost), and then did a utilitarian calculation taking into account everyone’s utility. This is the most realistic model, because, for example, people are much more willing to sacrifice for their children and spouses than for others. (We could imagine that most people think of self-interest as the interest of their immediate family.) Another example is that altruism for unrelated others may be greater for people in one’s group, such as a nation.

In the present studies, we tried to avoid this issue by asking subject to put aside their own feelings.

Moral Education A final model is relevant if we take the position of the moral educator rather than the individual decision maker. Arguably, the position of the educator is the most basic one (Baron, 1996).

If we had no moral principles, we could derive them by asking what principles we would rationally endorse as principles for everyone to follow. With an answer to this, we are in a position to know how to educate others. This argument is a theoretical device (in the spirit of the “veil of ignorance”).

Yet moral education is relevant if we consider the costs and benefits of efforts to inculcate or encourage altruism in others, as done by Kaplow and Shavell (2007). It might be easier to increase altruism toward distant strangers than toward family members, because the latter would be naturally strong and not susceptible to much further increase. Because of this natural strength, it would likewise be

futile (that is, prohibitively costly) to try to encourage altruism toward distant strangers *at the expense* of family members. An optimal moral system for the purpose of public norms, norms that were generally endorsed and supported, would thus take into account these natural constraints, wherever they come from. It is not obvious, though, that national loyalty is one of these constraints (in contrast to loyalty toward spouses and children, which has clear biological roots).

Implications

Of course, charities may capitalize on the above biases by, for example, suggesting that donors allocate some money to a particular charity even if they contribute to others, or advertising favorable statistics, even if they are not the most relevant ones. In turn, one lesson for a donor is to be more open minded. Stretch your imagination. See below the surface. Be aware of your own cognitive shortcuts that can undermine your goal of helping.

More generally, though, donors should put charitable contributions in a broader context that includes all their other altruistic behavior, including what they do with their work, and their political actions. Charitable institutions might also consider political action as an alternative to direct aid. For example, it may (or may not) be more helpful to Africans to lobby against trade restrictions on their exports, or in favor of medical research on tropical diseases, than to provide direct aid. Of course, such efforts must be considered in the light of tax laws.

Efforts at donor education might be seen as parallel to investor education, which is now carried out (to some degree, with some effectiveness) by investment advisors and by various institutions. It might even make sense for this function to be filled by current financial advisors, given that charitable donations are part of people's total financial picture. Estate lawyers will tell you how to minimize your tax burden, but not, yet, how to maximize the benefits that you leave behind.

REFERENCES

- Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal*, 100, 464–477.
- Baron, J. (1996). Norm-endorsement utilitarianism and the nature of utility. *Economics and Philosophy*, 12, 165–182.
- Baron, J. (1997). Political action vs. voluntarism in social dilemmas and aid for the needy. *Rationality and Society*, 9, 307–326.
- Baron, J. (2008). *Thinking and deciding* (4th ed.). New York: Cambridge University Press.
- Baron, J. (in press). Parochialism as a result of cognitive biases. In A. K. Woods, R. Goodman, & D. Jinks (Eds.), *Understanding social action, promoting human rights*. Oxford: Oxford University Press.
- Bekkers, R., & Wiepking, P. (2007, October 28). Generosity and philanthropy: A literature review. Available at the Social Science Research Network (SSRN): <http://ssrn.com/abstract=1015507>
- Bettman, J. R., Luce M. F., & Payne, J. W. (1998). Constructive consumer choice processes. *Journal of Consumer Research*, 25, 187–217.

AU: Update available?

- Charness, G., & Gneezy, U. (2008). What's in a name? Anonymity and social distance in dictator and ultimatum games. *Journal of Economic Behavior & Organization*, 68, 29–35.
- Fox, C. R., Ratner, R. K., & Lieb, D. S. (2005). How subjective grouping of options influences choice and allocation: Diversification bias and the phenomenon of partition dependence. *Journal of Experimental Psychology: General*, 134, 538–551.
- Hsee, C. K. (1996). The evaluability hypothesis: An explanation of preference reversals between joint and separate evaluation of alternatives. *Organizational Behavior and Human Decision Processes*, 46, 247–257.
- Kaplow, L., & Shavell, S. (2007). Moral rules, the moral sentiments, and behavior: Toward a theory of an optimal moral system. *Journal of Political Economy*, 115, 494–514.
- Margolis, H. (1982). *Selfishness, altruism, and rationality: A theory of social choice*. New York: Cambridge University Press.
- Read, D., & Loewenstein, G. (1995). Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices. *Journal of Experimental Psychology: Applied*, 1, 34–49.
- Sachs, J. D. (2005). *The end of poverty*. New York: Penguin Press.
- Tversky, A., Sattath, S., & Slovic, P. (1988). Contingent weighting in judgment and choice. *Psychological Review*, 95, 371–384.
- Ubel, P. A., DeKay, M. L., Baron, J., & Asch, D. A. (1996). Public preferences for efficiency and racial equity in kidney transplant allocation decisions. *Transplantation Proceedings*, 28, 2997–3002.

APPENDIX A

The questions asked in each study are shown below, with an example of one item.

CHAR 1 AND 2

A and B both help thousands of children. In A, each contribution goes to one child. In B, the contribution is spread out over all the children. Each child gets the same help.

How much would you allocate to A/B?

- 100%/0%
- 90%/10%
- 80%/20%
- 70%/30%
- 60%/40%
- 50%/50%
- 40%/60%
- 30%/70%
- 20%/80%
- 10%/90%
- 0%/100%

Which choice would give you a better feeling?

A or B: no difference

AU: CE changed and rearranged some of your data to make it easier to understand. Please review carefully.

Which choice is a more effective use of your money?

A or B: no difference

With which choice are you more sure that your contribution will do some good (as opposed to none)?

A or B: no difference

CHAR3, PART 1

What is the right allocation between A and B, ignoring your own feelings? All to A
More to A, some to B Equally to A and B More to B, some to A All to B

Which would give you a better feeling?

Contributing to A or contributing to B: no difference

What allocation between A and B would be the most efficient use of your money?

All to A More to A, some to B

Equally to A and B

More to B, some to A

All to B

CHAR3, PART 2

Case A: Your employer collects individual voluntary contributions by deducting them from paychecks. It is expected that 50% of the employees will contribute an average of \$20 each.

Case B: The employees have voted, by a margin of 55% to 45%, for a system in which the employer deducts \$10 from everyone's paycheck, leading to the same total contribution as in Case A.

Which Case would you favor if you had a choice?

A and B equal

Which case is more fair in allocating the cost among potential contributors?

A and B equal

Which case provides more freedom to potential contributors?

A and B equal

Which is more important in this case?

Fair cost allocation

Freedom of choice equal

What if Case A, the voluntary-contribution case, would raise half as much money as Case B (rather than the same amount). Which Case would you favor?

A and B equal

CHAR4, PART 1

A can save one life for \$10,000. B can save one life for \$12,500. The people helped are from the same groups, with the same problems.

What is the right allocation between A and B, ignoring your own feelings?

- All to A
- More to A, some to B
- Equally to A and B
- More to B, some to A
- All to B

What allocation would you feel best about making?

- All to A
- More to A, some to B
- Equally to A and B
- More to B, some to A
- All to B

What allocation between A and B would be the most efficient use of your money?

- All to A
- More to A, some to B
- Equally to A and B
- More to B, some to A
- All to B

What allocation between A and B would do the most good for each \$1,000 spent?

- All to A
- More to A, some to B
- Equally to A and B
- More to B, some to A
- All to B

CHAR4, PART 2

Your country requires everyone to buy health insurance. The fee is fixed at about \$2,500 per person.

Case A: To help those who have trouble paying, the government levies a special income tax. Families earning less than \$7,500 per person pay no tax. Others pay a fixed percentage of their income above \$7,500/person.

Case B: To help those who have trouble paying, charities collect voluntary donations. The charities distribute the funds to try to come as close as possible to the situation described in Case A. That is, they provide a subsidy for families

earning less than \$7,500/person, they reduce the subsidy gradually as income increases, and they solicit contributions from rich people who would pay less in Case B than in Case A. Suppose that the charities succeeds, so that the bottom line is the same for each family as in Case A.

Which case would you favor if you had a choice?

A and B equal

Which case is more fair in distributing the cost and benefits?

A and B equal

Which case provides more freedom of choice?

A and B equal

Which is more important in this case?

Fair distribution

Freedom of choice

Equal

