Warm Glow, Information, and Inefficient Charitable Giving

Clair Null*

Job Market Paper

Abstract

I investigate the efficiency consequences of donors who simultaneously give to multiple charities using revealed preference data from a lab experiment in which more than 200 real-world donors decide how to divide a gift between a charity they currently support and a set of international development charities. Most subjects simultaneously give to multiple development charities that have similar mission statements. This is true even when the social benefit of gifts, proxied by the matching rates received by the charities, are not equal. Taking preferences for charities as given, these choices result in substantial inefficiencies. Subjects forfeited social surplus (matching funds) equal to 25% of the value of their gifts. Two-thirds of donors who split their gifts are motivated by "warm glow", personal satisfaction derived from the act of giving that leads to a love of variety even among charities that have similar missions. The rest appear to do so because of risk aversion over the social benefit of their gifts, which leads them to diversify their charitable portfolios. Few subjects were willing to pay for information that could have enabled them to increase the social benefit of their gifts. This might help explain why there are so few rigorous evaluations of international development programs: such evaluations are costly to charities and not highly valued by donors.
JEL codes: O19, D64, D61, L31

*Department of Agricultural and Resource Economics, University of California, Berkeley. I would like to thank my committee members and participants of the Berkeley Development Lunch and ARE Development workshop for extremely helpful discussions of this project; Brenda Naputi of the XLab for technical support; and certainly the Kiwanis, Lions, and Rotary clubs who hosted me as a guest speaker and whose members participated in the experiment. This material is based upon work supported by a National Science Foundation Graduate Research Fellowship. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the National Science Foundation. I also gratefully acknowledge funding from the Institute for Business & Economic Research (Dissertation Research Award) and the XLab (Graduate Student Mini-Grant).

1 Introduction

Charitable giving in the U.S. is big business with major consequences for the world's poor, many of whom rely on goods and services provided by non-profit organizations. In 2007, Americans are estimated to have made donations totaling $306 billion, or 2% of GDP (Giving USA Foundation 2008). Although gifts to international causes comprise a relatively small fraction of total charitable giving, Americans gave roughly $30 billion for overseas projects in 2005, the most recent data available (Rollins 2007). In fact, private development aid is currently estimated to be approximately equal to official development assistance from the U.S. government when measured in terms of money that is actually reaching the poor, as opposed to debt relief or technical cooperation (Desai and Kharas 2008).

While the total sum of money being donated is quite large, it is divided across myriad charities in a very disaggregated way by many individual donors. This raises the possibility that resources will be inefficiently allocated in the absence of a social planner or market mechanism that could solve coordination problems. While small donations and small charities are not a priori inefficient, the complexity of so many donors allocating gifts across so many charities makes it more likely that any single donor will give to multiple charities that serve the same purpose but are not equally productive. I focus on this very specific form of inefficiency rather than a more general love of variety for different types of charities.

Unfortunately, we know relatively little about *how* donors choose which charities to support. We can infer that any donor who simultaneously gives to more than one charity does so because her expected marginal utility from each charity is equal. In the next section I argue that there are two aspects of the donation allocation decision that could lead to equal marginal utilities, even between charities that serve the same purpose and thus might otherwise be substitutes. First, a donor might be motivated by something other than the output her gift to a charity produces, deriving private utility directly from the act of making the gift; in the literature such utility is typically referred to as warm glow or joy-of-giving (Andreoni 1990). Second, charitable contributions can be thought of as credence goods since the donor never knows the true value of her gift in terms of what the charity produced (Darby and Karni 1973). As in the standard investment decision framework, donors who are risk averse over the social value of contributions might choose a portfolio of charitable giving that has a lower expected productivity in exchange for a reduction in the variance of charitable output. As a consequence, risk aversion, like warm glow, can lead to an inefficient allocation of resources across charities since a donor who is motivated by either of these factors considers something other than the expected marginal social productivity of her gifts.

In order to design policies that encourage donors to give generously and efficiently, it is necessary to first understand the factors that influence their decisions about how much and to whom to give. Inefficiencies due to risk aversion can be mediated by reducing the uncertainty associated with donations whereas the policy implications of inefficiencies from warm glow are less clear. In this paper, I emphasize that warm glow can lead to an inefficient allocation of gifts across charities, but it also has potentially beneficial consequences in terms of increasing total gifts either by reducing free-riding on others' gifts (because they are no longer perfect substitutes for one's own contributions) or by increasing the marginal utility of giving. Ideally, policies should encourage generosity by exploiting warm glow while at the same time safeguarding against inefficient allocations as much as possible. To the extent that warm glow donors do not pick charities based solely on their productivity, it is even more important that government regulation ensures a minimum level of quality among charities, in particular by preventing fraud (Greenlee, Fischer, Gordon and Keating 2007).

In observational data it would be impossible to determine if a donor gives to multiple charities because doing so maximizes her warm glow utility or because she is diversifying her charitable portfolio as a result of risk aversion, or because of some combination of the two factors. Instead, I use an experimental setting to investigate whether real-world donors are solely motivated by the social benefit of their gifts (proxied by changes in the matching rates received by the charities) or whether other considerations such as warm glow and risk aversion lead to inefficient allocations, taking as given each individual's preferences over a set of international development charities that have similar mission statements. By exogenously varying the marginal social benefit of a gift and the risk associated with this marginal social benefit, my experimental design allows me to distinguish subjects who give to multiple charities because they are perfect substitutes from those who split their gifts inefficiently, and to identify the source of these inefficiencies as either being attributable to warm glow or risk aversion.

In a modified dictator game, when asked to divide a gift between a charity the subject currently supported and three development charities which are arguably substitutes (CARE, Mercy Corp, and/or Oxfam America), 70% of the experimental subjects gave money to more than one of the development charities. When one of the charities that a subject supported became exogenously more productive (in the sense that it had a higher matching rate), very few subjects perfectly substituted into that charity, giving it their entire gift. Rather, the majority of subjects weakly substituted, moving a larger share of their total giving into a higher-valued charity but continuing to allocate at least some of their total giving to the now

lesser-valued charities. In the process, subjects who did not perfectly substitute contributed to major inefficiencies in the allocation of gifts across charities (taking initial preferences as given), forfeiting matching funds equal to 25% of total un-matched giving. Weak substitution could be consistent with either risk aversion or warm glow, but decisions that involve risk over matching rates identify only one-third of these subjects as risk averse. Instead, warm glow seems to explain the majority of inefficient giving.

An alternate test of how much subjects cared about the social benefit of their gifts indicates that relatively few subjects were willing to pay for information about matching rates that could have enabled them to increase the value of their gifts. When they were told the distribution of matching rates but not how the rates would be assigned to charities, only 40% of subjects were willing to give up a small portion of their endowments in order to find out which charity would receive the highest rate; the rest preferred to allocate their gifts without knowing what they would be worth to the charities. This result should be concerning for development economists who have argued that rigorous evaluations of aid projects are too rare (Duflo 2004, Savedoff and Levine 2006). The possibility that donors might place so little value on information about the relative social benefits of different programs could help explain why so few charities are willing to undertake costly evaluations of their projects.

In a complementary set of scenarios, I find suggestive evidence that warm glow is one of the key reasons it is profitable for charities to allow donors to earmark their gifts for specific programs, an option I will refer to as "donor control". In this case, donors are choosing between multiple programs offered by a single charity, rather than between numerous charities, but the allocation problem the donor solves is quite similar. Donor control could increase the utility-maximizing donation either because it increases the marginal altruistic utility by resolving agency problems between the donor and the charity or because it increases the marginal warm glow utility. However, charities that offer donor control risk facing binding budget constraints as a result; this is particularly common in the wake of natural disasters that lead to outpourings of generosity for one specific cause, sometimes beyond what the charity would have chosen to spend. Understanding why gifts increase in response to donor control is an important first step in designing institutions that foster generosity without running the risk of inefficient allocations. For example, charities might be able to avoid binding budget constraints due to restrictions placed on gifts, especially after natural disasters, if they were to ask for permission to reallocate restricted gifts at the same time that they offer the option of donor control. Indeed, the introduction of donor control in the experiment increased total giving by 8% but more than half the subjects who gave more did so out of

warm glow rather than the need to solve an agency problem, since they agreed to let the charity reallocate their restricted gifts.

The inefficiencies that this paper investigates in a laboratory setting are quite likely to exist in the real world as well, where millions of donors give to multiple charities in small increments. An estimated 65% of households with incomes less than $100,000 gave to charity in 2006 (Giving USA Foundation 2007) and a recent random sample of Americans shows that most gifts are relatively small, with two-thirds of the reported gifts less than $100 and a median gift of $50 (The Center on Philanthropy at Indiana University 2007). Approximately 4000 U.S.-based non-profits operate international development programs, and the vast majority of these are small operations with less than $500,000 in annual revenue (Kerlin and Thanasombat 2006). With so many donors and charities, the potential for inefficient allocations is a serious concern.

This paper is related to the large literature on the efficiency of official foreign aid allocations[1] and bridges the broad research fields of charitable giving and generosity in experimental dictator games.[2] In that regard, this paper complements two recent publications that also used dictator games with unconventional recipients (welfare beneficiaries and diabetic smokers) in order to study dictators' sensitivity to how their gifts would be spent, though neither of these papers tested as rigorously for motivations as I do (Fong 2007, Jacobsson, Johannesson and Borgquist 2007). Several authors have investigated substitution in charitable giving, but this work has been more frequently concerned with crowding out of private donations by public contributions (Andreoni 1989, Ribar and Wilhelm 2002, Andreoni and Payne 2008) rather than with a single donor's choice between gifts to multiple charities. Reinstein (2006) is the only other study, aside from this one, to consider one individual's substitution patterns between charities. He uses both observational data from the Panel Study of Income Dynamics and experimental data from a laboratory experiment, but both of his approaches are confounded by variations in the total amount an individual gives and thus he can not assess the magnitude of inefficiencies or the cause of imperfect substitution

[1] While official foreign aid inefficiencies usually stem from political considerations rather than warm glow (Hoeffler and Outram 2008), the two themes are nonetheless related.

[2] Camerer (2003) provides an overview of dictator games such as those used in laboratory tests of altruism and social preferences by Andreoni and Miller (2002), Charness and Rabin (2002), and Fisman, Kariv and Markovits (2007). Eckel and Grossman (2003) find that subjects in dictator games are more generous to charities when contributions are matched rather than subsidized, with important implications for income tax policy. Karlan and List (2007) and Meier (2007) provide evidence from field experiments that donors are responsive to the "price" of their gifts in terms of matching rates. Other field experiments have compared various fundraising mechanisms such as lotteries, challenge grants, and matching grants (Landry, Lange, List, Price and Rupp 2006, List and Lucking-Reiley 2002, Rondeau and List 2008).

as my experiment allows me to do. Finally, another closely related paper reports evidence from a field experiment in which the introduction of donor control had no effect on either the intensive or extensive margins of giving among households solicited for donations by a Dutch NGO that provides aid programs focused on children in Africa and Asia, although the author acknowledges that the option to exercise control was not obvious on the donation form (Siegel 2006).

This paper extends the existing literature by exploring the importance of risk aversion and agency problems, in addition to warm glow, as factors that influence donors' decisions of which charities to support and whether or not to earmark their gifts for specific programs. The experimental nature of my data allows me to probe reasons *why* donors do or do not substitute between charities and/or restrict their gifts, and the fact that all of the experimental subjects are donors outside the laboratory strengthens the relevance of their experimental choices.

In the next section, I explain a simple theoretical framework for understanding how donors with different types of preferences will allocate their gifts in the experiment, which is described in Section 3. In Section 4, I discuss the results relating to substitution between charities. Section 5 presents results from the same experiment carried out with professional subjects as a robustness check. Section 6 provides a second test of the importance of warm glow in the related context of donors' decisions of whether or not to earmark gifts for a specific purpose. Section 7 concludes.

2 Theoretical Framework

The basic goal of this study is to characterize donors' preferences over multiple charities which serve the same purpose. These charities use the donations they receive in order to produce the same public good Y which is an argument of the donor's utility function.[3] Ultimately, a donor chooses the bundle of consumption goods and charitable donations that maximizes her utility, but here I abstract away from the question of private consumption versus charity and focus on only the allocation of donations across charities, as in the second step of a two-stage budgeting process, in keeping with the experimental setup.

Suppose that a dollar given to charity 1 produces m_1 units of the public good while a

[3]For example, CARE, Mercy Corps, and Oxfam America (the charities used in the experiment) all have very similar mission statements and could be considered to "produce" poverty alleviation, which is a public good for everyone who cares about the welfare of the world's poor (i.e. whenever *anyone* makes a gift to one of these charities, *everyone* who cares about the welfare of the poor is better off).

dollar given to charity 2 produces m_2 units of the public good. A donor who cares only about producing the public good will view the charities as perfect substitutes. In this case the donor gives to both charities only if the budget constraint she faces has the same slope as her indifference curves. However, if charity 1 becomes more productive than charity 2 ($m_1' > m_2$), this donor would maximize her utility by switching to a corner solution in which she allocates her entire gift to the more productive of the two charities. This situation is shown in Illustration 1 below, with the donor picking an interior solution of $(m_1 g_1^\star, m_2 g_2^\star)$ when her indifference curve has the same slope as her budget constraint but choosing the corner solution $(m_1' g_1^{\star\star}, 0)$ when charity 1 becomes more productive, such that the budget constraint (in bold) is less steep than her indifference curves.

Illustration 1: Charities as Perfect Substitutes

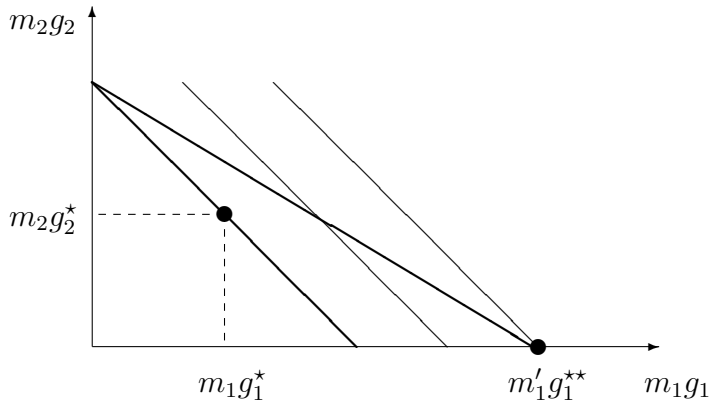

However, the special case of indifference curves for perfect substitutes that have the same slope as the budget constraint is not the only scenario in which a donor might choose an interior solution to the problem of allocating her gift between charities. Rather, warm glow and risk aversion over production of the public good would both lead the donor to have strictly convex indifference curves and therefore divide her gift between several charities even if they do not have the same marginal productivity in terms of public good production (formal definitions of warm glow and risk aversion are discussed in the following paragraphs). Though these two explanations of convexity are caused by different aspects of the utility function, they are manifest in the same way and are both shown in Illustration 2 below. An increase in charity 1's productivity relative to charity 2 ($m_1' > m_2$), will not be sufficient to lead the donor to give exclusively to charity 1, and inefficiency in the production of the

public goods ensues. A donor who initially chose the interior solution $(m_1 g_1^\star, m_2 g_2^\star)$ might now choose a different interior solution such as $(m_1' g_1^{\star\star}, m_2 g_2^{\star\star})$, substituting partially, but not perfectly, toward the more productive charity 1.

Illustration 2: Charities as Imperfect Substitutes

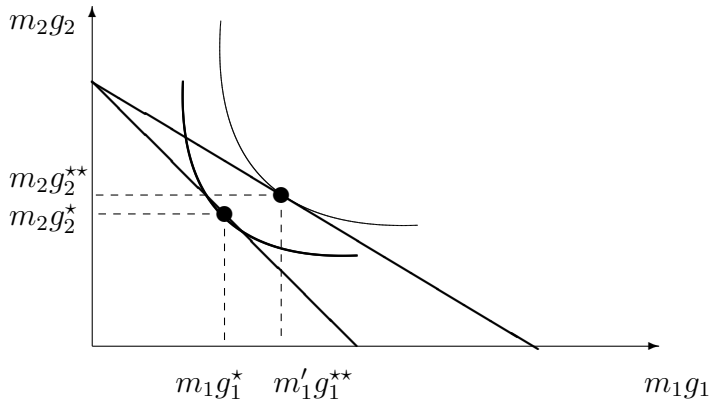

Most research on charitable giving considers contributions to charity as simply an element of the consumption commodity space. Alternatively, due to uncertainty over what gifts will be worth in terms of public good production, the donor's decision can be modeled as the choice of a portfolio of risky assets. However, the donation allocation decision differs from the standard consumption and investment problems in an important way. As Andreoni (1990) suggested, donors might derive utility directly from the act of making a donation, in addition to any altruistic utility experienced from increasing production of the public good. Thus, a pure altruist's utility function over private consumption c_i and the public good Y is $u(c_i, Y)$ whereas a purely warm glow donor's utility function is $u(c_i, g_i)$ such that the donor cares only about her gift to the charity g_i and not at all about the public good. Specifically, let $Y = f(g_i + G_{-i})$ where gifts from individual i and all others, G_{-i}, are converted into the public good according to the production function $f(\cdot)$. Then warm glow implies that someone else's donation is not a perfect substitute for one's own donation and helps explain the fact that private contributions to public goods are not perfectly crowded out by public contributions, as would be the case if all donors were purely altruistic.[4] As I demonstrate

[4]In perhaps the most direct possible test, Crumpler and Grossman (2008) designed an experiment in which subjects' contributions to charities perfectly crowded out contributions made by the experimenters. Nonetheless, over half of the experimental subjects chose to contribute from their own experimental payouts, even though they could not increase the net amount the charities received from the experiment. Previously,

below, warm glow can have similar consequences in terms of making donors reluctant to substitute between charities solely on the basis of their social productivities.

We can write the donor's problem of allocating her total donation D across gifts g_{ij} to the n charities as

$$\max_{g_{i1},\cdots,g_{in}} \mathbb{E}[u(Y, g_{i1}, \cdots, g_{in})]$$

$$\text{s.t.} : \sum_{j=1}^{n} g_{ij} \leq D \text{ and } \; g_{ij} \geq 0 \text{ for } j = 1, \cdots, n$$

(see Appendix I for a more explicit example). Each dollar given to charity j produces m_j units of the public good. While it is likely that the charities' production functions exhibit diminishing marginal returns to donations globally, it is not unreasonable to assume that they are approximately linear over the range of gifts that any one donor would contribute. However, donors do not know the true value of m for any given charity, either because they do not know the shape of the charity's production function or because they have imperfect knowledge about others' gifts and thus do not know where the charity is currently located along its production function. Instead, donors maximize their expected utility based on their their priors about the distribution of m for each of the n charities.

To help fix ideas, assume that utility is additively separable in altruism and warm glow, so that we can write the donor's utility function as

$$\mathbb{E}[u(Y, g_{i1}, \cdots, g_{in})] \equiv \underbrace{\alpha \mathbb{E}[u_1(Y)]}_{\text{altruism}} + \underbrace{(1-\alpha) u_2(g_{i1}, \cdots, g_{in})}_{\text{warm glow}}$$

since the uncertainty is only over Y and not over the gifts themselves. Let $u_1(Y) = Y^{(1/\rho)}$ where $\rho = 1$ for a risk neutral donor, $\rho > 1$ for a risk averse donor, and $0 < \rho < 1$ for a risk loving donor. The α parameter ranges from 1 (pure altruist) to 0 (pure warm glow).

Note that the expected marginal public good productivity of a dollar to charity j is constant ($\frac{\partial Y}{\partial g_{ij}} = m_j$). In contrast, I assume that the marginal warm glow utility of donations to charity j is decreasing ($\frac{\partial u_2}{\partial g_{ij}} > 0$ but $\frac{\partial^2 u_2}{\partial g_{ij}^2} < 0$).[5] Thus, in the absence of risk a purely

in another laboratory study, Andreoni (1993) found that private contributions to a public good were not perfectly crowded out by taxes.

[5] In the classic Andreoni definition, a donor experiences warm glow over her total charitable donation, rather than over gifts to individual charities. To that extent, the warm glow I consider in this paper is a subset of the broader definition. Since warm glow over total donations would not affect the donor's choice of which charities to give to, such a donor would be indistinguishable from a purely altruistic donor in terms of substitution patterns between charities. Because there are no efficiency consequences of such preferences, I do not devote further attention to the more general case of warm glow over total contributions.

altruistic donor ($\alpha = 1$) will make gifts only to the charity(ies) with the highest value of m. A purely warm glow donor ($\alpha = 0$) will consider only the marginal warm glow utility she would derive from giving to each of the n charities when making her allocation decision, and is likely to make smaller gifts to several charities because the marginal utility of a gift to any single charity is decreasing in the size of the gift. More realistically, a donor who is motivated by both altruism and warm glow ($0 < \alpha < 1$) will weakly substitute toward a more productive charity but persist in giving to multiple charities even if their marginal public good productivities are not equal, because for such a donor the marginal warm glow utility she derives from multiple gifts outweighs the inefficiencies in public good production that arise from supporting less productive charities.

By definition, a donor who is risk averse over the production of the public good ($\rho < 1$) will be willing to accept a lower expected level of public good production as long as the variance of public good production is also lower. This reduction in risk can be achieved by giving to charities whose marginal social productivities are not positively correlated. A risk averse donor who gives to a charity with a lower expected marginal social productivity in the interest of diversifying her portfolio contributes to inefficiency in public good production.

In an experimental setting, it is possible to exogenously shift the budget curve and observe whether or not donors perfectly substitute between charities. In the notation of the framework above, donors who are risk neutral ($\rho = 1$) and purely altruistic ($\alpha = 1$) will perfectly substitute. Determining whether donors who only weakly substitute do so because of warm glow ($\alpha < 1$) or risk aversion ($\rho > 1$) is not as straightforward. As described in the next section, I use variation in the degree of risk associated with the social benefit of gifts to different charities in order to identify the nature of these donors' preferences.

3 The Experiment

In this study, I do not seek to explain *why* a donor initially chooses the charities she does, though I do briefly discuss characteristics that correlate with the allocation decision in the next section. Rather, I am interested in the degree to which allocation decisions are motivated by production of the public good versus other factors such as warm glow or risk aversion. While the external validity of an experiment can be drawn into question, I attempted to minimize this concern by recruiting a subject pool of known donors rather than using professional undergraduate subjects as is common in the literature. In particular, subjects in this experiment were members of service clubs (Kiwanis, Lions, and Rotary),

who by their very membership in such clubs self-identify themselves as donors of both time and money with an interest in international causes.[6] In order to make the decisions salient, prior to the experiment subjects at each session were informed that after the experiment one decision for one subject would be randomly selected to be paid (anonymously) in real money according to the chosen subject's decisions in the experiment. More details on the logistics of the experiment are provided in Appendix II.

To investigate choices between charities, the experiment exogenously varied the marginal public benefit and degree of risk associated with donations. This variation allows me to observe how allocations between the charities change, taking as given each subject's preferences for the charities in a baseline scenario. Panel A of Table 1 summarizes how responses from the experiment are used to classify subjects on the basis of the parameters of their utility functions as described in the previous section. It should be noted that while I will be able to differentiate warm glow ($\alpha < 1$) from perfect altruism ($\alpha = 1$) in the absence of risk aversion ($\rho \leq 1$), the experimental design does not allow me to differentiate between these two for risk averse subjects (i.e. if $\rho > 1$, I can not say anything about α). Thus, my results put a lower bound on the importance of warm glow.

In the first stage of the budgeting process, subjects divided a \$100 gift from the author's grant between their service club and a group of three development charities (CARE, Mercy Corps, and Oxfam). Subjects were told that each dollar given to the development charities would be matched at a rate of at least 50¢ and that they would be able to specify which of the charities in particular they wanted to give the money to in following decisions.[7] Subjects were given the charities' mission statements as listed on the independent charity evaluator CharityNavigator.org (see Appendix III) and assured that all three of these secular charities had received the highest ranking for being "fiscally responsible and financially healthy".

After having set the total amount that they wanted to give to the development charities, subjects were asked to specify how they wanted to divide the gift between the three charities in eight situations with varying matching rates and degrees of risk. As promised, the matching rate was always at least 50¢ and in some cases was more. Specifically, each subject

[6] All three of these clubs are global organizations, each with its own international foundation. While Rotary is perhaps best known among the three for its campaign to eradicate polio, Kiwanis activities to prevent iodine deficiency disorders and the Lions' commitment to eye health have each had global impact.

[7] The author explained that the availability of matching funds for gifts to the development charities did not reflect a judgment about the worthiness of the development charities relative to the host club, but rather was intended to compensate for the fact that clearly everyone present was comfortable giving money to the club whereas they might not be familiar with the development charities. The matching funds were meant to make the decision between the development charities and the club non-trivial.

allocated her gift under the following scenarios:

1. equal $1 matching rates for all three charities
2. matching rates of 50¢, $1, and $1.50 (order randomly assigned to the three charities)
3. matching rates of 50¢, $1.25, and $1 (paired with the random assignment in the prior decision)
4. equal matching rates as in decision 1 but with 20% risk that each charity would not receive the matching funds (risks across charities were independent, determined by 3 separate rolls of a 10-sided die)
5. unequal matching rates as in decision 2 but with 20% risk that each charity would not receive the matching funds
6. equal matching rates with risk as in decision 4 but with the option to purchase delivery guarantees (i.e. accept a guaranteed but lower matching rate)
7. unequal matching rates with risk as in decision 5 but with the option to purchase delivery guarantees (i.e. accept a guaranteed but lower matching rate)
8. matching rates of $1.50 for two charities and $3 for one charity, unspecified assignment of rates to charities; information about which matching rate was assigned to each charity could be purchased for $5 deducted from the total amount to be given to the charities (alternately, the subject could opt not to purchase the information and divide her gift without knowing the assignment of matching rates to charities)

Decision 1 defines each subject's set of "preferred charities" (those that the subject gives to when matching rates are the same for all three charities). Decisions 2 and 3 were designed to test whether subjects who initially gave to multiple charities would treat their preferred charities as perfect substitutes when the matching rates changed. Among subjects who did not perfectly substitute, these decisions allow me to check that preferences are consistent with standard utility maximization theory as well as gauge the extent of the inefficiencies that arise when subjects do not perfectly substitute. Decisions 4-7 were designed to explore the degree to which allocations were motivated by risk aversion. In decision 4, subjects who are risk averse will spread their gifts equally between their preferred charities. In decision 5, subjects who are sufficiently risk averse will forfeit higher matching rates in order to allocate

their gifts more equally. In decisions 6 and 7, the delivery guarantee option allows sufficiently risk averse subjects to revert to their preferred allocations as in decisions 1 and 2. Finally, decision 8 tests whether subjects care about the social benefit of their gifts.

A follow-up questionnaire elicited basic demographic data (age and sex), and asked subjects if they had ever "heard of / supported with time or money / or had a personal connection to" the three charities prior to the experiment.

4 Results

A total of 205 subjects participated in the study over the course of 16 sessions at Kiwanis, Lions, and Rotary club meetings in the East Bay region of California (details provided in Appendix II). Participating clubs ranged in size from fewer than ten to more than thirty members, but were generally demographically similar with the exception of one club which was a college campus affiliate of Rotary and had much younger members than the other clubs. Two-thirds of the 151 subjects who provided demographic information were male. Ages ranged from 18 to 87 with a mean and median of 52 and 56 years, respectively.

4.1 Budget first-stage and initial allocations

Subjects in the experiment were quite generous to the group of development charities. Panel A of Figure 1 shows a histogram of gifts to the development charities from the first stage of the budgeting process. Around half of the subjects gave either \$50 or \$100 to the development charities, with almost all of the rest choosing gifts that were multiples of either \$10 or \$25.

Table 2 investigates correlates of the decision to give to a particular charity, reporting marginal effects from probit regressions of a binary variable for whether or not a subject gave to each specific charity on a set of subject and charity characteristics. As would be expected, subjects who had previously heard of a charity were significantly more likely to give to that charity during the experiment. In particular, about 60% of subjects had previously heard of CARE relative to only 30% of subjects who recognized the names of the other two charities (p-values for the difference between CARE and the other two charities are less than 0.01), and this translated into subjects being significantly more likely to give to CARE during the experiment.[8] Prior personal connections to these three charities were relatively rare, but

[8]The statistically significant negative coefficient on age is an unstable result that disappears when members of the undergraduate club are excluded from the regression. These 11 subjects, who were much younger than the rest of the pool, were slightly more likely to give to all three clubs; 9 out of 11 (82%) gave to all

perfectly predicted giving to the charities in the experiments.[9]

Only 35 subjects (17% of the total sample) gave their entire gift to a single development charity when the matching rates were equal. As shown in Panel B of Figure 1, the majority of subjects (70% of the total sample) split their gifts across all three of the development charities.

4.2 Substitution and inefficiencies

Subjects' responses to changes in matching rates are categorized in Table 3. While perfect substitution was quite rare, weak substitution was quite common, though in some cases subjects actually substituted away from a higher-valued charity.[10] The extent of substitution is shown in Figure 2 which compares the share of each subject's total gift given to a particular charity before and after the matching rates change.[11] In the figure, perfect substitution appears as observations for which the highest-valued preferred charity lies along the $y = 1$ line. Of the 146 subjects who initially gave to multiple charities, only 11 (4%) treated their preferred charities as perfect substitutes ($\alpha = 1$, $\rho = 1$). Another three subjects who initially

three, relative to 62% of the rest of the subject pool.

[9]Including indicators for the different experimental sessions also does not yield any notable trends in terms of popularity of certain charities with members of certain clubs.

[10]It is worth noting that 44 subjects (21% of the total sample) maintained the exact same allocations even when matching rates changed. There are three possible explanations for this behavior: first, the subjects could have been exerting minimal effort in the experiment and not seriously considering the different scenarios; second, they could have been confused by the experimental protocol; or third, they could have been truly impervious to changes in the matching rates and degree of risk associated with gifts to the three charities (an extreme form of warm glow). Unfortunately, it is impossible to distinguish between these three hypotheses with the available data, although trends from the pool of professional subjects suggests that confusion might be an important factor (see Section 5 for details). Demographic characteristics of these subjects also provide few clues as to why they did not respond to changes in the matching rates. These subjects were drawn from almost every club, with the notable exception of the undergraduate club. On average, subjects who maintained constant allocations were somewhat older and more likely to be male than those who responded to the matching rates, though these differences are only marginally statistically significant if the undergraduate club members are excluded from the comparison.

Regardless of the factors that led these subjects to be unresponsive to changes in matching rates, it is clear that such behavior is inefficient from the point of view of maximizing production of the public good (in this case, matched contributions to the charities). Moreover, it seems quite plausible that in real-world decisions about charitable giving, donors might also be effort-minimizers, confused, or motivated entirely by warm glow. To be conservative, I exclude these 44 subjects when trying to determine why subjects did not perfectly substitute so that my estimates are a lower bound on the degree of warm glow (and the resulting aggregate inefficiencies) in this sample of donors.

[11]Data on charities with the highest matching rates is both necessary and sufficient to show perfect substitution. For weak substitution among subjects with three preferred charities, data on charities with the lowest matching rates are also needed since the fraction received by the charity with the highest matching rate could remain unchanged while the subject reallocated part of the gift from the charity with the lowest matching rate to the one with the middle-valued matching rate.

gave their entire gift to a single development charity later substituted perfectly into their preferred charity with the highest matching rate in decisions 2 and 3, bringing the total number of risk neutral pure altruists to 14 (7% of the total sample).

Though perfect substitution was quite rare, the majority of subjects responded at least somewhat to the changes in matching rates ($\alpha > 0$). This is shown in Figure 2 by the frequency of observations between the 45° line and the $y = 1$ line. Sixty-two subjects (30% of the total sample) reallocated a portion of their gift in order to capture a higher matching rate while still leaving at least some of the gift for their other preferred charity(ies). Twenty-three subjects (11% of the total sample) actually reallocated some of their gifts into charities that were *less* valuable when the matching rates were changed (seen in the figure as observations for highest-valued preferred charities that are below the 45° line). In general, the subjects' responses were quite consistent between decisions 2 and 3.

Comparing subjects' allocations between the three budget sets, only four subjects' choices violate of the Weak Axiom of Revealed Preference (WARP).[12] All other subjects' choices were consistent with WARP so there is no reason to believe the data were generated by a process other than standard utility maximization. Thus, the weak substitution observed in the experiment is a rational choice and can be interpreted as contributing to true inefficiency.

The relative inelasticity of gifts with respect to changes in matching rates led to large sums of money being left on the table in the form of unclaimed matching funds. Taking as given preferences from the first decision in which the matching rates were equal and excluding subjects who did not respond to changes in matching rates, imperfect substitution resulted in forfeited matching funds ranging from \$2.50 to \$112.50 per subject. When measured as a percent of each subject's un-matched gifts to the three development charities these losses are quite notable, with half of the subjects who responded in some way to changes in matching rates nonetheless giving up matching funds equal to 25% or more of their un-matched gift to the development charities. Summing up the losses over both decisions, the total of over \$3600 represents almost 25% of the total un-matched gifts to development charities. Clearly, the magnitude of this inefficiency is cause for concern. In the rest of this section, I discuss evidence from the remainder of the experiment that allows me to identify the characteristics of preferences that led to these inefficient choices. Ultimately, by distinguishing between warm glow and risk aversion, it will be possible to recommend policies that could help prevent such inefficiency in real world charitable giving.

[12]This is similar to the rate of violations in Andreoni and Miller (2002) but much lower than the rate of violations in Fisman et al. (2007), likely because the former paper observed choices under 8-11 budget sets whereas the latter observed choices under 50 different budget sets.

4.3 Risk aversion

What causes subjects to substitute weakly but not perfectly to the highest-valued of their preferred charities? One possibility is that even though observable matching rates have changed, subjects are reluctant to invest too heavily in a single charity when they have incomplete knowledge of its effectiveness. The experiment did not attempt to measure subjects' priors over the distributions of the three charities' relative productivities, but the exogenous risk that matching funds would not be delivered to the charities, introduced in decisions 4 and 5, makes it possible to observe how they respond to additional risk. From this we can infer whether or not their reluctance to substitute could be at least partially driven by their original perceptions of the real-world risk associated with the three charities. Specifically, in decision 4, when the matching rates are equal, risk averse subjects will divide their gift as evenly as possible between their set of preferred charities since the risks associated with matching funds for each charity are independent. Subjects who do not adjust their allocations in response to the additional risk are unlikely to have originally split their gifts because of risk aversion.[13]

Unfortunately, decision 4 was not very informative since so many of the subjects who did not perfectly substitute had initially split their gifts equally as shown in Panel A of Table 4.[14] In decision 5, when matching rates differed, even risk averse subjects might not spread their gifts more equally, since doing so would mean sacrificing higher matching rates. However, when matching rates differed there was more scope for subjects to change their allocations since fewer subjects had split their gifts equally in the absence of risk. In fact, 19 subjects (20% of those who did not perfectly substitute) gave up higher matching rates in order to spread their gifts more equally in response to the exogenous risk in decision 5 (Table 4, Panel

[13] Over 80% of subjects who gave to more than one development charity divided their gift equally between whichever charities they supported, consistent with use of the "$\frac{1}{n}$" heuristic that has been observed in other contexts, such as retirement fund portfolio choice (Benartzi and Thaler 2001). This behavior simplifies the analysis since it is less likely that extreme risk aversion will be mistaken for warm glow. If a subject has strong priors about one of the charities being more efficient but nonetheless gives to another charity out of extreme risk aversion, changes in either the level or associated risk of matching rates might not be enough to affect her allocations and she could be incorrectly classified as warm glow. Based on the fact that so few subjects seem to clearly prefer one of the charities that they support relative to the others, there is little concern that risk aversion will be mistaken for warm glow in my experimental results. To be cautious, for 16 subjects who gave to all three charities but gave one charity more or less than the other two, I categorize substitution patterns on the basis of the two charities which received equal gifts in the baseline case. There were only nine subjects who gave different amounts to all of the charities they supported.

[14] To be conservative, I assume that subjects perceived \$10 as the smallest increment in which gifts could be divided. For example, I would count a \$100 gift that was divided \$30/\$30/\$40 between the three charities as being equally divided. While some subjects did divide gifts in \$5 increments and a few even divided in \$1 increments, \$10 increments were much more common in subjects' allocations.

A). In both risky decisions there were also subjects who consolidated their gifts (10-13% of those who did not perfectly substitute), indicating risk loving preferences over donations.

Panel B of Table 4 summarizes the results from decisions 4 and 5. Among those who did not perfectly substitute when matching rates changed in the first three decisions, I find evidence of risk averse preferences ($\rho > 1$) for 18 subjects (19%), while 12 subjects (13%) exhibited risk loving preferences ($\rho < 1$).[15] Eleven subjects (11% of those who did not perfectly substitute) could have spread their gifts more equally when matching rates were the same for all charities but chose not to and thus appear to be risk neutral ($\rho = 1$). I conclude that while some subjects who did not perfectly substitute between multiple charities in the first three decisions might have been motivated by risk aversion, there is also a group of subjects who show no evidence of risk aversion (or even exhibited risk loving preferences) and who therefore must have chosen to support multiple charities because of warm glow.[16]

For the remaining subjects who did not perfectly substitute but for whom decisions 4 and 5 were uninformative, the delivery guarantee purchase decisions provide an indication of risk aversion regardless of the initial allocation of gifts.[17] Of course, finding that a subject chooses not to purchase a delivery guarantee does not necessarily mean that she is not risk averse, only that she is not sufficiently risk averse for the purchase to be utility maximizing given the cost of the guarantee and the degree of risk that matching funds will not be delivered. With that caveat in mind, based on the decision to purchase delivery guarantees, I find evidence of risk aversion for an additional 11 subjects for whom decisions 4 and 5 were uninformative (12% of those who did not perfectly substitute). Though I only use the delivery guarantee purchase data for subjects whose imperfect substitution I am still trying to explain after the uninformative tests of decisions 4 and 5, among all subjects (including those who gave to only one charity and those who maintained constant allocations), demand for delivery guarantees is quite high, with a 65% purchase rate.

Thirteen subjects (6% of the total sample) did reallocate gifts into more valuable chari-

[15] This is less than the number of subjects who spread gifts more equally in decision 5 (quoted in the above paragraph) because of the seven subjects who had mixed evidence on risk preferences, consolidating their gift in one of the two risky decisions and spreading it more equally in the other.

[16] Among the 14 subjects who perfectly substituted in response to changes in matching rates and therefore appeared to be risk neutral pure altruists, the extra risk was enough to induce seven (50%) of these subjects to spread their gifts more equally and sacrifice higher matching rates.

[17] In practice, subjects "purchased" the delivery guarantee by agreeing to accept a lower matching rate that was a sure thing. The price of a guarantee against the 20% risk of matching funds not being delivered was 12.5¢ for each 50¢ of matching funds, so that only sufficiently risk averse subjects should want to purchase a guarantee. A table in the experimental protocol clearly listed the price of the guarantee and the value of a $1 gift with the guarantee and without the guarantee, contingent upon the die roll that would determine whether matching funds would be delivered.

ties after insuring themselves against the risk that matching funds might not be delivered. However, seven subjects (3% of the total sample) who purchased delivery guarantees then allocated less efficiently, and some subjects who purchased delivery guarantees maintained their allocations from decision 5 (in which risk was introduced) rather than reverting back to their more efficient allocations from decision 2 (in which there was no risk). Finally, several subjects changed their allocations even though they did not purchase delivery guarantees, contrary to the prediction that decisions 6 and 7 would be the same as decisions 4 and 5 for subjects who did not purchase delivery guarantees. While institutions that reduce the risk associated with charitable giving might help improve the efficiency of resource allocation in some cases, there is only weak evidence that this occurred in the experiment.

4.4 Willingness to pay for information

The final decision regarding simultaneous giving to multiple charities investigated whether or not subjects were willing to pay for information about the assignment of matching rates to charities. This decision mimics the real-world trade-off between spending money on programs and spending money on evaluations in order to improve program efficiency, a topic that has garnered considerable attention from development economists who argue that charities and governments under-invest in evaluations (Duflo 2004, Savedoff and Levine 2006).

Subjects were given the option to spend $5 of their total gift to the development charities in order to find out which of the three would receive a matching rate of $3 (the other two would receive matching rates of $1.50). Altruistic subjects whose donation was at least $20 and gave to all three charities, or whose total gift is greater than $35 and gave to two charities, would find it profitable to purchase the information.[18] Among the 166 subjects who participated in this part of the experiment, 140 (84%) met these criteria on gift size and number of charities supported. Consistent with the economists' arguments that donors choose to purchase less information than would be socially optimal, only 58 subjects (41% of those who met the criteria) actually purchased the information.[19]

Unfortunately, many subjects have missing data on their allocations in decision 8 so it is difficult to say conclusively whether those who purchased the information would have

[18] Of course, other subjects might also have been willing to purchase the information despite having not previously given to certain charities since the matching rates in this decision were much higher than those in prior decisions.

[19] Six subjects who would not have been predicted to purchase the information did so. All six had only given to CARE in prior decisions.

made use of it.[20] Among the 36 subjects who purchased information and did record their allocation decisions, many reallocated money into the charity with the highest matching rate but again fewer than half of them substituted perfectly, consistent with behavior in earlier decisions, even though the efficiency costs were much larger in this case. Of course, the bigger inefficiency in this case comes from subjects who chose not to purchase information. These subjects forfeited matching funds ranging from 30-150% of the value of their unmatched gifts, with the median donor sacrificing matching funds exactly equal to the value of her unmatched gift, a truly staggering sum.

Did subjects simply not care about the potential to substitute into the charity with the highest matching rate? Given the evidence from prior decisions, this seems quite plausible and the data are broadly consistent with this possibility. Referring back to substitution patterns in decisions 2 and 3, 10 (71%) of the 14 subjects who had perfectly substituted purchased information whereas only 40% of those who weakly substituted did so.

Another possible explanation for why so few subjects purchased the information is that they simply did not understand the cost of not doing so. In order to gauge whether or not subjects could decide "correctly" in the absence of warm glow, a subset of subjects were asked to repeat decision 8 for personal profit.[21] Rather than having the option to give to three charities with different matching rates, subjects were asked to divide the amount of their donation between three generic investment options labeled simply A, B, and C, one of which had a \$3 rate of return with the other two earning \$1.50 per dollar invested. Again, information regarding the assignment of rates of return to investment options was available for purchase at a cost of \$5 (deducted from the investment). This exercise indicates that indeed some subjects may not have understood the consequences of their decision not to purchase information since among the 137 subjects for whom it would have been personally profitable, 40% opted not to purchase information.[22] That said, 24 subjects (29% of those who inefficiently chose not purchase information regarding charities' matching rates) did purchase information in the investment decision, indicating that they had the ability to

[20]The prevalence of missing data in this decision is likely due to the logistics of how the information purchase was handled in the experiment. Information about which charity would receive the \$3 matching rate was only provided after everyone had made their purchase decisions. After this delay, many of the subjects who had purchased information did not make the effort to record their allocations in response to knowledge about the matching rates.

[21]This question was added to the experiment after the first 5 sessions had been held.

[22]This is almost definitely an upper bound on the fraction of subjects who would make such a poor financial decision in real life. Some of these subjects probably did not read the question carefully, since the experiment was almost over and the wording closely resembled that of the prior question about charities. Another possibility is that this behavior is more evidence of the $\frac{1}{n}$ heuristic.

assess whether or not the cost of the information was worthwhile in terms of the ability to earn higher rates of return. This discrepancy is illustrated in Figure 3 as the vertical distance between the percent of subjects who purchased information about the personal investment and those who purchased information about the charities' matching rates.

Figure 3 also suggests a concerning pattern in how subjects seem to have approached the information decisions. We would expect the percent of subjects who purchase information to be increasing in the size of their gifts, since the cost of the information was fixed but the financial benefit of the information is greatest for subjects whose gifts are larger; this is also true for the personal investment decision. The fact that these curves are so flat might be an indication that subjects thought only about the fixed cost of information and not how it related to the higher payouts they could earn, either for the charities or themselves.

I conclude that the results of the information purchase decision are not necessarily meaningful for all subjects, since some might not have been able to accurately judge the value of information in terms of the potential to capture higher matching rates. Nonetheless, roughly one-third of subjects who were capable of assessing the value of information inefficiently chose not to purchase, with major consequences in terms of forfeited matching funds. Since information actually reduces risk in this situation, these inefficient choices must result from warm glow. Puzzlingly, there seems to be much higher willingness to pay for a delivery guarantee that eliminates downside risk over matching rates than there is for information that enables subjects to lock in a higher matching rate when there is risk over which charity will be most valuable. Loss aversion and framing effects might help explain some of the difference, but this paradox warrants further investigation in future work.

4.5 Classifying preference types

Panel A of Table 1 synthesizes the data on substitution and responses to risk in order to classify subjects' preference types. The 14 subjects who perfectly substituted when matching rates changed require no further explanation. Among those who did not perfectly substitute, two-thirds (60 subjects) appear to have done so because of warm glow since they did not exhibit risk aversion.[23] For the 29 subjects who showed signs of risk aversion, no conclusions can be drawn regarding the balance of altruism and warm glow in their utility functions, placing a lower bound on the the importance of warm glow in explaining choices in the

[23] Of the nine subjects who initially divided their gifts unequally between the three charities, six showed evidence of risk aversion, leaving only three who might be mistakenly classified as warm glow even if they were actually extremely risk averse. Clearly, these few subjects would not meaningfully change the conclusion that warm glow explains more of the subjects who imperfectly substituted than risk aversion does.

experiment. Bootstrapped confidence intervals for the proportion of subjects of each type (excluding those who maintained constant allocations) indicate that significantly more subjects who inefficiently split their gifts between multiple development charities did so out of warm glow rather than risk aversion. Comparing the inefficiency that resulted from the two potential motivations for imperfect substitution, subjects motivated by warm glow contributed disproportionately to the aggregate inefficiencies, sacrificing, on average, a third of the value of their unmatched gift relative to the fifth that risk averse donors left on the table in the form of unclaimed matching funds. This difference is statistically significant at more than 99% confidence.

5 Results from Professional Subjects

One of the strengths of this paper is that the subjects participating in the experiment are all real-world donors, in contrast to most laboratory studies of altruism which rely on professional subjects, the majority of whom are undergraduate economics and psychology students. However, one potential problem with the novel subject pool I recruited is that they were generally unfamiliar with this sort of study. As a result they might not have fully understood that their decisions could have real monetary consequences or might have been confused about the details of the different decisions. As a robustness check, in this section I present evidence from an additional 191 professional subjects who participated in the experiment and were unlikely to have been confused by the protocol.[24] In general, the experimental choices made by the two subject pools are remarkably similar, despite major demographic differences.[25]

The only modification to the experimental protocol described in Section 3 and Appendix II was that the outside option for the professional subjects was to keep the gift for oneself, rather than keep it for one's club.[26] Unsurprisingly, professional subjects showed less generosity toward the development charities relative to the real-world donors' allocations, as evidenced by the differences between the upper left panel of Figure 4 and Panel A of Figure 1. It is also interesting to note that the professional subjects showed much more heterogeneity

[24] Professional subjects were recruited through the XLab at U.C. Berkeley.

[25] The average age among the professional subjects was 21 and only 43% were male. I did not collect data on race from the real-world donors, but based on my observations at the club meetings, the vast majority were Caucasian, with only a few African, Hispanic, or Asian Americans. Among the professional subjects 70% identified themselves as Asian American.

[26] In accordance with XLab policies, each subject was paid a $10 participation fee for the 30-minute experiment in addition to the randomly chosen decision for one subject that was paid in real money.

in the size of the gift that they gave to the development charities relative to the real-world donors, perhaps indicating that the professional subjects took a more sophisticated approach to solving the allocation problem. Nonetheless, the professional subjects were just as likely as the real-world donors to divide their gifts across all three development charities in the baseline decision with equal matching rates as shown in the upper right panel of Figure 4.

The professional subjects' responses to changes in matching rates were also quite similar to those of the real-world donors, as shown in the lower left panel of Figure 4.[27] Although some professional subjects did perfectly substitute into whichever of their preferred charities had the highest matching rate, weak substitution was again much more common. As with the real-world donors, there were also several professional subjects who chose a less efficient allocation, moving a higher share of their gift into a charity that was less valuable.[28] Like the real-world donors, the professional subjects' substitution patterns led them to forfeit large sums of matching rates, equal to roughly 30% of the value of un-matched gifts.

A slightly higher share of the professional subjects who did not perfectly substitute showed evidence of risk aversion based on the decision to spread their gifts more equally among preferred charities when exogenous risk over matching rates was introduced. As with the real-world subjects, this test was uninformative for many subjects, several of whom were identified as risk averse based on their decision to purchase a delivery guarantee. Moreover, there is only weak evidence that delivery guarantees promoted more efficient allocations, since a number of professional subjects who purchased the guarantees actually chose less efficient allocations in that decision and half of the professional subjects who purchased the guarantees did not change their allocations despite having eliminated the risk; only a third of the professional subjects who purchased the guarantees reallocated in a more efficient manner. Ultimately, combining results from all decision rounds involving risk over matching rates, slightly more than half of the professional subjects who did not perfectly substitute are explained by warm glow, with the remainder showing evidence of risk aversion as shown in Panel A of Table 5.

Finally, the lower right panel of Figure 4 shows that the professional subjects were no more willing to pay for information than their real-world donor counterparts. Only half of the professional subjects who could have increased the social benefit of their gifts chose to

[27] Again, there were very few (6) subjects whose choices violated WARP.

[28] Perhaps the most important difference between the real-world donors and the professional subjects is the fact that a much smaller share of the latter maintained constant allocations even when the matching rates changed (see the bottom row of Panel A in Tables 1 and 5). This finding suggests that some of the real-world donors who adopted this behavior might have done so because of confusion or laziness, strengthening the argument for excluding such subjects from the analysis.

purchase the information that would have allowed them to do so. Furthermore, although a surprisingly high fraction (32%) of the professional subjects were not able to correctly assess the value of the information even when they stood to personally profit, half of the professional subjects who were capable of assessing the value of the information nevertheless chose not to purchase information about charities' matching rates, consistent with the prevalence of warm glow.

In summary, real-world donors and professional subjects made remarkably similar choices throughout the experiment, mediating concerns that the real-world donors might not have understood the experimental protocol. At the same time, part of the explanation for the consistency of the results between the two subject pools likely lies in the fact that 75% of the professional subjects said they had given either time or money to at least one charity in the past year; in this regard the difference between the two subject pools is not so much that fewer professional subjects are donors, but that the professional subjects' gifts are more likely to be in terms of time rather than money.

6 Donor Control

While this paper's focus is on substitution between charities, substitution between programs run by a single charity is an increasingly important issue, with more and more charities offering donors the option to earmark their gifts for specific projects, a trend I refer to as "donor control". As long as the charity has sufficient unrestricted funds in order to implement its desired budget, restricted gifts are simply non-binding constraints and there will be no efficiency costs of introducing donor control.[29] In contrast, when donors' attention is focused on one specific issue, as is the case in the wake of a well-publicized natural disaster, many small restricted gifts can add up to binding constraints for charities. As an example, after the Indian Ocean tsunami the international medical charity Medecins Sans Frontieres (MSF) estimated it would need 25 million euros to address the medical needs of victims. Within a week, MSF had received 105 million euros, at which point its leadership decided to stop accepting restricted donations, asking instead that donors make gifts to MSF's general fund to be used in places of need all over the world. In addition, MSF called donors back to ask for permission to de-restrict their gifts; less than 1% of donors asked for a refund rather than agreeing to let MSF reallocate their gifts.

[29] Of course, this also raises the question of why donors would respond to donor control, when their restricted gifts do not actually have any effect on how the charity will allocate its resources. I suggest that warm glow can explain this potential puzzle.

Natural disasters are not the only instances in which even small restricted gifts become binding constraints. In the past several years a number of websites that aim to link donors with projects of their choice have become increasingly popular. Operating without any unrestricted funds, every gift becomes a binding constraint in these cases. For example, the online micro-lending website Kiva.org has experienced blistering growth since it first opened operations in October 2005, facilitating over $43 million in micro-loans by connecting lenders with borrowers on a person-to-person basis, giving "donors" complete control over where their money goes. DonorsChoose.org and GlobalGiving.com deal with projects rather than individual beneficiaries but also follow the model of letting donors choose exactly which programs they want their funds to support.

Donor control has the potential to increase total giving by enabling donors to find the projects that they care most about, in essence resolving agency problems between donors and charities. At the same time, however, there are important efficiency consequences of this new paradigm that deserve special attention. Donors who gave to MSF after the tsunami wanted to help, but they did not realize that their gifts could be more productive if spend on needs other than those of the crisis victims. Lenders who pick borrowers based on their profiles on Kiva's website might not be as discerning as a loan officer at the local micro-finance institution would be. By understanding why donors respond to donor control, it might be possible to use it as a tool to promote generosity without also risking inefficient resource allocation.

6.1 Conceptual framework

In keeping with the conceptual framework developed for substitution between charities, we can understand substitution between programs within a single charity in terms of the altruistic and warm glow utility that restricted gifts generate for a donor.[30] In the case of donor control, it is necessary to differentiate between restricted gifts that place binding constraints on the charity and those that the charity can offset with its unrestricted funds. Donor control can only solve agency problems between an altruistic donor and a charity if her restricted gifts are binding. Otherwise, the rational donor will realize that her restricted gifts have no net effect on the charity's budget and donor control will be irrelevant.

[30]See Duncan (2004) for an alternate model of "impact" philanthropy in which a donor cares about the change in production of a public good rather than its level. As a consequence, such a donor might prefer to target her gifts to a specific purpose (such as child sponsorship) not because she experiences more warm glow or altruistic utility, but rather because she perceives that her gift has more of an impact when restricted to a particular purpose.

If donor control increases the marginal warm glow that a donor experiences from giving, perhaps because some projects (such as vaccines) are more pleasant to think about than others (such as latrines), then restricted gifts might be relevant regardless of whether or not they place binding constraints on the charity. I will say that such a donor experiences "differential" warm glow as a result of restricting her donations. A donor whose preferences are both altruistic and warm glow might prefer to place restrictions on her gifts (so that she can experience differential warm glow) but simultaneously release the charity from these restrictions (so that the charity can maximize production of the public good, trusting it to know best how to do so).

The experiment I use to investigate how and why donors respond to donor control allows me to differentiate donors who have agency problems with the charity (and therefore want their restricted gifts to be binding constraints) from donors who respond to donor control primarily because of differential warm glow (and therefore do not care if their gifts are binding constraints).

6.2 The experiment

In addition to the decisions described in Section 3, the experimental protocol contained three questions that investigated how much and why gifts increase when a charity introduces donor control. Subjects were asked to divide a $100 gift from the author's grant between their club and an un-named "study charity" that was described as follows[31]:

- The study charity's goal is to improve children's health in rural areas of the world's poorest countries.
- Because latrines serve many children and are more costly than vaccines, the study charity plans to spend 75% of all donations improving sanitation by digging pit latrines and the other 25% providing free vaccines.[32]

Gifts to the study charity were matched at 50¢ per dollar.

Next, subjects were again asked to divide the $100 gift between their club and the study charity (with the same matching rate), but were given the option to give directly to latrines

[31] For the professional subjects, the outside option was to keep the money for themselves.

[32] In actuality, gifts allocated to the study charity for latrines were given to WaterAid America and gifts allocated for vaccines were given to the GAVI Alliance (Global Alliance for Vaccines and Immunization) according to the stated budget rule of 75% latrines / 25% vaccines. The identity of the "study charity" was revealed to subjects at the conclusion of the experiment when payouts were issued.

and/or vaccines in addition to being able to contribute to the charity's general budget. Subjects were informed that all of the study charity's other funding had already been committed to its latrine and vaccine programs, with the implication that additional gifts would be binding, though this point was not emphasized in particular. Subjects who placed restrictions on their gifts were also asked if they would give the study charity permission to re-direct the donation in the event that it could be more useful if spent on the other program. Finally, in order to break cases in which the subject might have been indifferent between restricted and unrestricted gifts, the donor control question was repeated, but this time subjects were informed that there was an administrative cost of 10¢ per dollar associated with gifts to specific programs. The experimental protocol clearly specified that unrestricted gifts would be matched at 50¢ per dollar whereas the matching rate for gifts specific to latrines or vaccines was 50¢ - 10¢ cost.

Subjects who made restricted gifts as a means of resolving an agency problem with the charity would not be expected to grant the charity permission to reallocate their gifts. On the other hand, subjects who made restricted gifts because such gifts increased their warm glow utility relative to general donations might not actually care if their restricted gifts place binding constraints on the charity, and might even prefer that the charity spend the money in the most efficient way if their preferences are also somewhat altruistic. By allowing the charity permission to reallocate their restricted gifts, warm glow subjects could potentially experience higher warm glow utility without having to sacrifice altruistic utility that would come from maximizing production of the public good. Panel B of Table 1 summarizes the method for classifying subjects' preferences on the basis of their responses in the experiment.

6.3 Results

6.3.1 Effects on intensive and extensive margins of giving

Figure 5 documents increases on both the extensive and intensive margins of giving when subjects were given control over how their gifts would be spent. As in the other part of the experiment, the professional subjects were generally less generous to the study charity with an average gift of $40 compared to the real-world donors' average gift of $60 (p-value for the difference in means is less than 0.01; p-value for a Kolmogorov-Smirnov equality of distributions test is also less than 0.01). That said, after accounting for the level difference between the two groups, their responses to the introduction of donor control are remarkably similar, so I integrate the two subject pools in the results below. Of the 32 subjects who

kept the full $100 gift for their club in the baseline decision without donor control, nine (28%) showed generosity towards the charity when donor control was introduced. Strangely, three subjects who gave some of the gift to the charity in the baseline decision then kept the full $100 when donor control was introduced, so the net effect on the external margin was an increase of six donors. While this is a small sample on which to test for effects on the extensive margin, converting close to 30% of non-donors into donors is a major success.

Effects on the intensive margin are also notable, with 72 (20%) of the 354 subjects who had given to the charity in the baseline decision increasing their gifts when donor control was introduced. These subjects gave between $1 and $70 more than they had in the baseline decision, with an average increase of $27 among the real-world donors and $13 among the professional subjects.[33] Among all subjects who gave to the charity in the baseline decision, the average gift increased significantly with the introduction of donor control (p-value less than 0.01). Total gifts to the charity increased 8% between the baseline and the decision in which donor control was introduced.

An important robustness check is to confirm that subjects who gave more after the introduction of donor control actually did so because they wanted to exercise the privilege of telling the charity how to spend the money. As expected, all but 11 (14%) of the 81 subjects who increased their gifts after the introduction of donor control (on either the extensive or intensive margin) made restricted gifts to the charity. Perhaps the option of donor control also carries informational value about the charity that is meaningful for donors regardless of whether or not they place restrictions on their gifts, leading these 11 donors to increase their gifts even though they did not have preferences over how the charity spent the money.

This experiment tracked only the immediate response to the introduction of donor control, which seems to have been quite profitable for the charity, but there are also some indications that offering donors control over how their gifts are spent might have broader benefits for the charity than can be measured simply in terms of the size of gifts. Among the 282 subjects who gave to the charity in the baseline decision but did not increase their gifts when donor control was introduced, 169 (60%) nonetheless made restricted gifts. To the extent that donors feel more goodwill towards the charity when they are offered control over their gifts, they may be more likely to continue to support the charity if donation decisions are dynamic and are influenced by prior decisions.

[33]It should be noted that subjects' ability to increase their gifts in response to the introduction of donor control was truncated on the upper end by the $100 cap on how much they could give to the charities through the experiment. Without this upper bound, the effects of donor control on gift size could have been even larger.

6.3.2 Permission for the charity to reallocate restricted gifts

Why is donor control so popular? By allowing subjects to specify whether or not they wanted their restricted gifts to be binding for the charity, I can determine whether donor control was being used as a means of resolving agency problems or if it appealed to subjects because of differential warm glow. Panel B of Tables 1 and 5 combines the data on restricted gifts and permission for the charity to reallocate in order to classify subjects according to their preference types. Among the 202 subjects who exercised donor control and indicated whether or not they wanted their restricted gifts to be binding for the charity, one-third were using it in order to resolve an agency problem, since they made binding restricted gifts and refused the charity permission to reallocate. The other two-thirds of the subjects who responded to donor control were doing so because of the warm glow that they derived from restricting their gifts, since they made non-binding restricted gifts to the charity.

The one-third/two-thirds split between subjects motivated by agency problems and warm glow is based on the conservative assumption that subjects who exercised donor control but who gave to both programs actually wanted to shift the proportion of spending devoted to those two programs. If instead such subjects simply miscalculated or were not aware of the charity's stated allocation rule between the two programs, the prevalence of warm glow would be much higher. In fact, 80% of subjects who exercised donor control actually gave to both programs, and over one-third of subjects who exercised donor control also made an unrestricted gift to the charity. Only 57 (24%) of the 237 subjects who exercised donor control and 18 (22%) of the 81 subjects who increased their gifts in response to donor control expressed a clear preference for one or the other of the charity's programs by giving their entire gift in restricted form to that particular program. Even among these subjects, almost half agreed to give the charity permission to spend the gift on the other program.[34]

The decision to make a restricted gift that is not binding for the charity is also consistent with a voting model, in which donors signal their preferences to the charity but do not seek to immediately change the allocation of the charity's resources. In a sense, such a response is like a weak form of an agency problem. I do not explicitly deal with that interpretation in this paper, but in future work it would be interesting to formally test whether donors use restricted gifts in order to vote their preferences, for example by comparing donations in response to a solicitation plus a survey of budget priorities versus in response to a solitication that offers donor control.

[34] Unsurprisingly, given the emphasis on vaccines rather than sanitation in the media and the service clubs' program priorities, 80% of subjects who gave to only one program earmarked their gift for vaccines.

Regardless of the interpretation, it is clear that many subjects who responded to donor control did not necessarily want their restricted gifts to be binding constraints for the charity. This finding implies that charities could potentially avoid inefficient binding budget constraints without sacrificing the increased generosity generated by donor control simply by asking donors for permission to reallocate their restricted gifts.

6.3.3 Costly donor control

The results discussed so far have all been under the assumption that donor control is costless. While this very well might be the working assumption of most real-world donors who make restricted gifts, the experiment could have overstated the importance of warm glow by counting subjects who were indifferent between unrestricted and restricted gifts as responding to donor control. In order to be more conservative, the experimental protocol also included a donor control allocation decision in which restricted gifts carried an administrative cost in terms of a lower matching rate. Remarkably, I find that even when donor control is costly to exercise, a sizable proportion of subjects place restrictions but then give the charity permission to reallocate their gifts, effectively negating their restrictions, and an even larger fraction of subjects who make restricted gifts show no clear preference for one of the programs in particular. Specifically, one-third of the 347 subjects who made a gift to the charity placed costly restrictions on it but only half of the subjects who made restricted gifts wanted them to be binding. Only 26 (21%) of the 121 subjects who exercised costly donor control showed a clear preference for one of the two programs and even among those, 12 (10%) agreed to give the charity permission to spend the gift on the other program if it would be more useful that way. Clearly, for the 59 subjects (15% of the total sample) who forfeited matching funds on restricted gifts that they then de-restricted, the illusion of donor control is worthwhile, even when it has no consequences other than to reduce the value of the gift, proving that some donors are willing to pay for warm glow.

7 Conclusions

In summary, warm glow explains roughly two-thirds of the inefficiencies that arise from donors' choices about which charities and which programs to support. While charities might be able to increase contributions by exploiting this aspect of their donors' utility functions, there can be important efficiency consequences of doing so, since by definition such donors do not make their allocation decisions solely on the basis of the social benefit of their gifts.

This was clear in both parts of the experiment, which documented that 1) very few subjects perfectly substituted to whichever of their preferred charities had the highest social benefit and 2) some subjects sacrificed matching funds in order to make restricted gifts, even if they did not ultimately care whether or not the restrictions they placed on their gifts were binding for the charity. If this experiment is indeed indicative of how donors would respond to changes in the social benefit of their gifts in the real world, the results bode ill for efficiency of resource allocation across charities and across programs within a single charity. Specifically, when one of the charities they already supported became exogenously more productive in the experiment, subjects continued to give to lesser-valued charities, forfeiting social surplus (matching funds) equal to 25% of the value of their gifts.

The experiment also sheds light on one potential reason why there are so few rigorous evaluations of development projects in the real world. Just as half of the subjects in the experiment were unwilling to pay for information that could have increased the social benefit of their gifts, rigorous evaluations might be rare in part because they are costly to charities and largely ignored by donors. Moreover, the fact that those who made larger gifts were no more likely to purchase information suggests that subjects focused on the fixed cost of acquiring information rather than the benefits of knowing which programs are most effective.

There is mixed evidence on the importance of risk aversion in subjects' decisions about how many charities to support. Some subjects did diversify their charitable portfolios when exogenous risk over matching rates was introduced, but the discrepancy between the popularity of guarantees against down-side risk and the low willingness to pay for information that could have allowed subjects to lock in a higher matching rate presents a puzzle. Loss aversion and simple marketing (use of the word "guarantee") could help to explain this seeming inconsistency, but the issue warrants further research. It might also suggest a potential strategy for charities that want to fundraise for rigorous evaluations: frame the cost of such information as protection against risk. On the other hand, drawing attention to the inherent risk associated with donations might actually be counter-productive.

On the question of why donors prefer to make restricted gifts rather than general contributions to a charity, agency problems seem to be a relevant factor for a third of the subjects who responded to donor control in the experiment. It is less clear whether or not donors who make restricted gifts in the real world think of these gifts as binding constraints for the charity, although to a certain degree this ambiguity is also present in the experimental results since there was not an explicit test of subjects' perceptions of the fungibility of restricted gifts. That said, for the two-thirds of subjects who responded to donor control in

the experiment and gave the charity permission to reallocate their restricted gifts, we can be confident that they were accurately identified as being motivated by warm glow.

Having established that warm glow, risk aversion, and agency problems influence donors' decisions about their charitable giving, we can begin to design and test policies that could help mitigate the inefficiencies that result from these aspects of donors' preferences. In the case of donor control, the solution to warm glow inefficiency is simple: ask donors for permission to reallocate gifts at the same time that they earmark their donations for certain projects. This way donor control can still be used to solve agency problems but the charity can avoid inefficiencies caused by warm glow donors. The policy implications concerning risk aversion are less straightforward. Many subjects in the experiment were willing to pay for an intervention that eliminated experimental risk over matching rates but such guarantees are infeasible in the real world. Information from charity evaluators such as CharityNavigator.org, the Better Business Bureau's Wise Giving Alliance, and the revised IRS form 990 (filed by tax-exempt organizations) could all be viewed as forms of "insurance", but there is no way to entirely avoid the uncertainty surrounding the social benefit of charitable donations and it is not at all clear that donors would equate the cost of time that it takes to gather such information with the simple cost of the guarantee in the experiment. Finally, it will be difficult to design policies that can mitigate warm glow inefficiencies across charities without also having unintended consequences for the level of giving, since warm glow can affect both the size of the pie as well as how it is divided. To the extent that warm glow donors can not be trusted to discriminate against charities that engage in fraudulent activities, government regulation is even more important to ensure that only charities which exceed a minimum quality threshold are allowed to operate.

This paper has highlighted several ways in which the classical model of a purely altruistic donor who contributes to a single public good fails to predict real-world donors' behavior. Incorporating risk aversion, agency problems, and especially warm glow into our understanding of how donors choose which charities to support makes it possible to identify interventions that can promote efficient resource allocation.

Figure 1: First Stage Budget Decisions and Initial Allocations Across Charities

Panel A

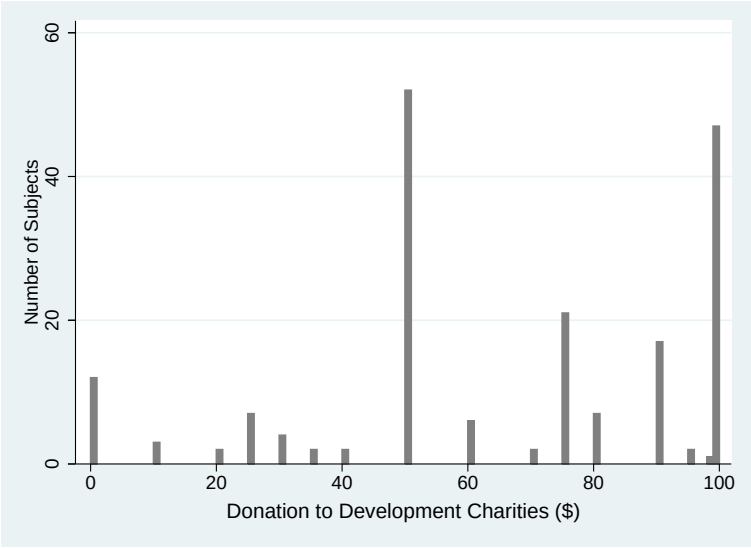

Panel B

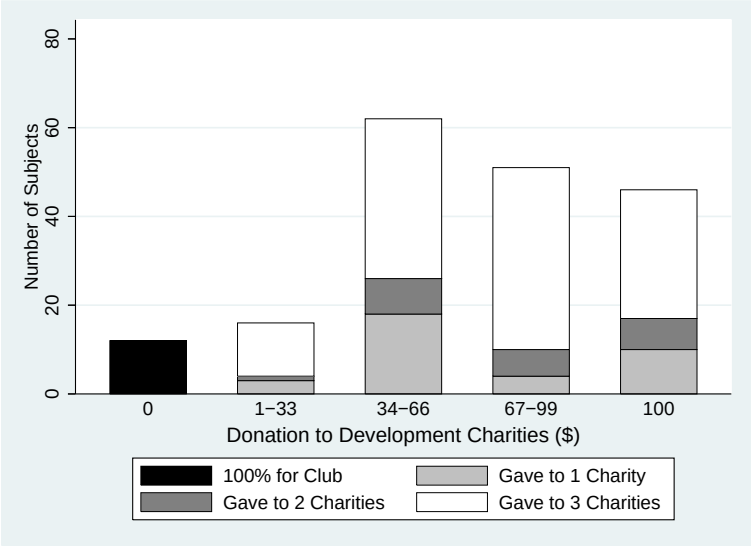

Notes: Subjects' responses when asked to divide $100 between their club and a set of three development charities (CARE, Mercy Corps, and Oxfam America), under the conditions that 1) in the subsequent decisions they would be able to choose which charity(ies) of the three in particular to give the money to and 2) gifts to the development charities would be matched at a rate of 50¢ per dollar. The number of charities receiving gifts from a subject is calculated based on responses to decision 1, in which all three development charities had the same matching rate. Eighteen of the 205 subjects skipped this section of the experiment.

Figure 2: Substitution Across Charities in Response to Changes in Matching Rates

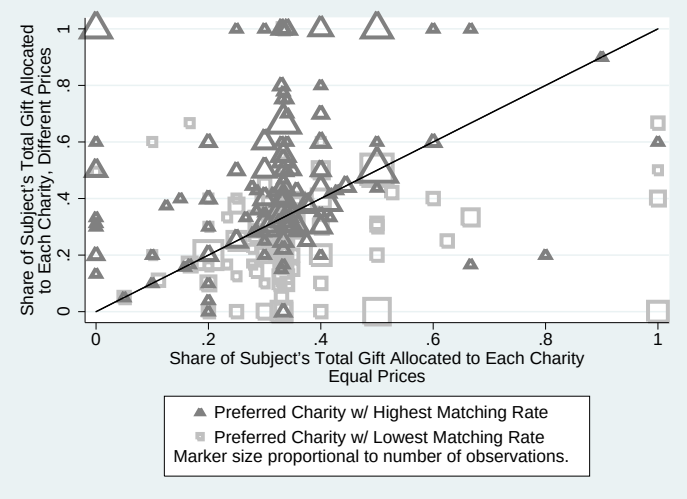

Notes: Data on allocations from decisions 1 (x axis) and 2-3 (y axis) . Graph excludes observations in which the subject never gave any of the gift to the charity in question or always gave the entire gift to the charity in question. See Section 4.2 for a description of the substitution patterns depicted in this figure.

Figure 3: Decisions to Purchase Information about Pairings of Matching Rates to Charities

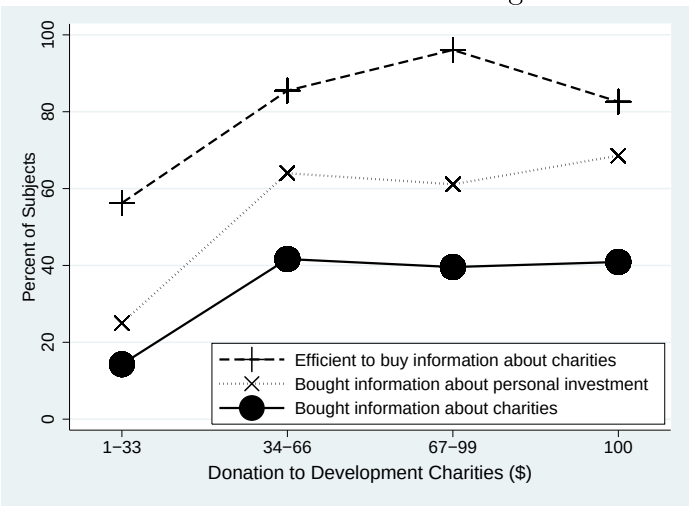

Notes: Subjects' responses when asked if they wanted to spend $5 of their gift in order to find out which of the three charities would have a $3 matching rate, knowing that the other two would each receive matching rates of $1.50. "Efficient to buy information about charities" refers to subjects whose gift was at least $20 and gave to all three charities at some point during the experiment or whose gift was at least $35 and gave to two of the charities at some point during the experiment. See Section 4.4 for a description of the personal investment robustness check.

Figure 4: Results from Professional Subjects

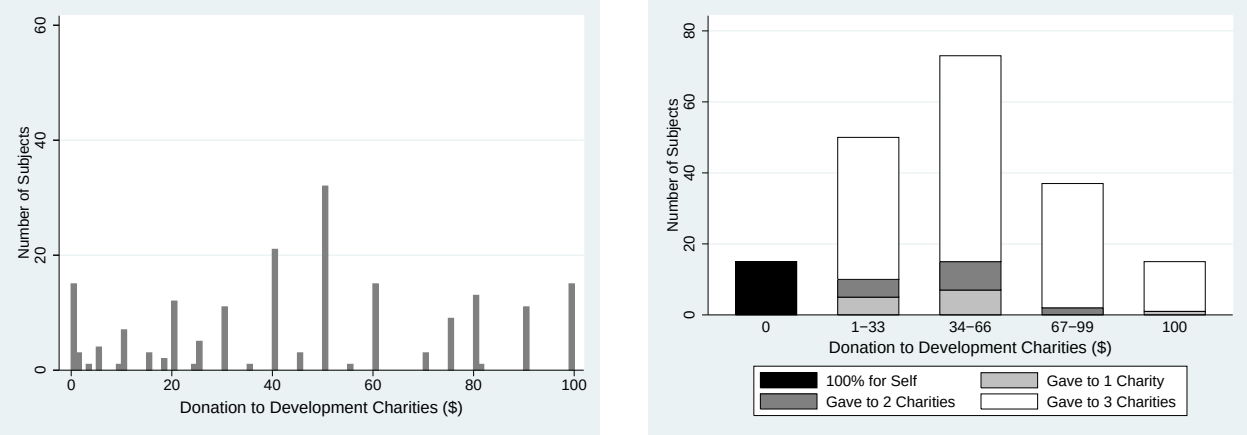

Notes: Subjects' responses when asked to divide $100 between oneself and a set of three development charities (CARE, Mercy Corps, and Oxfam America). See Figure 1 for further details.

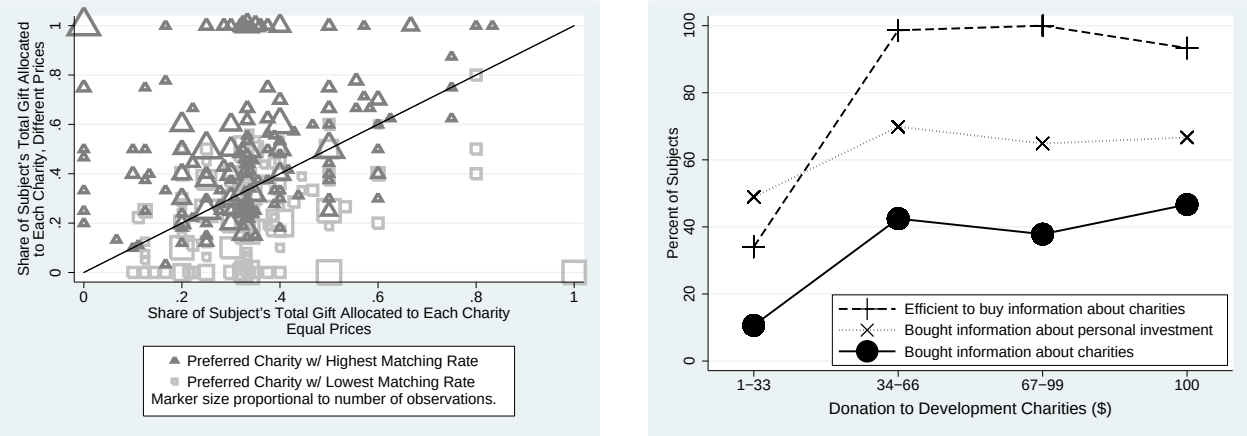

Notes: See Figures 2 and 3 for details on the data and interpretations.

Figure 5: Effect of Introducing Donor Control on Generosity

Panel A: Real-world Donors

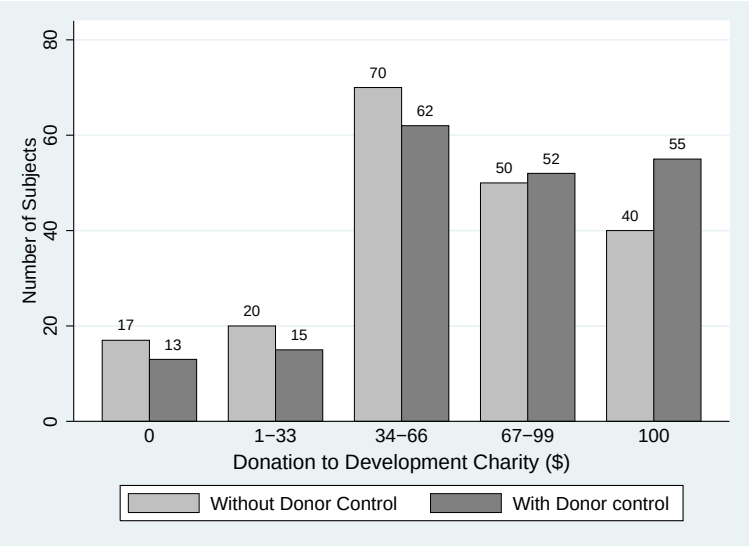

Notes: Subjects' responses when asked to divide $100 between their club and an un-named development charity that digs latrines and distributes vaccines. Eight of the 205 subjects had insufficient data necessary to make the comparison.

Panel B: Professional Subjects

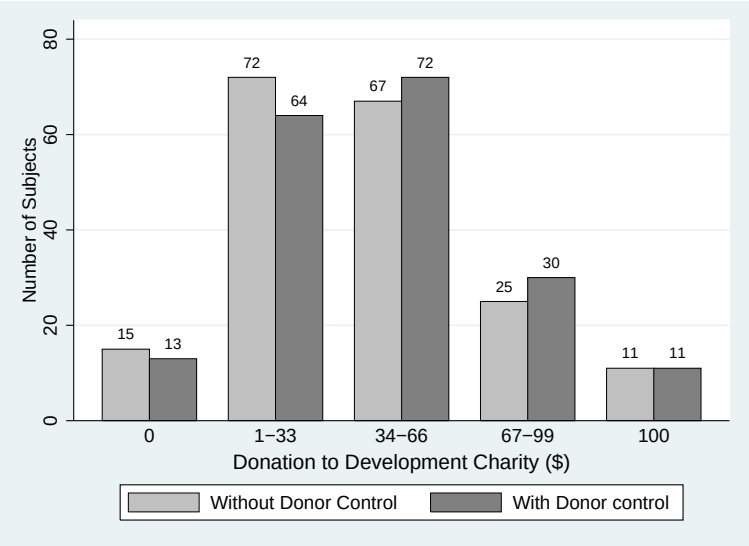

Notes: Subjects' responses when asked to divide $100 between oneself and an un-named development charity that digs latrines and distributes vaccines. One of the 191 subjects had insufficient data necessary to make the comparison.

Table 1: Classifying Preference Types

Panel A: Simultaneous Giving to Multiple Charities

	Not Risk Averse ($\rho \leq 1$)	Risk averse ($\rho > 1$)
Pure altruist ($\alpha = 1$)	perfectly substitutes n=14 [8-20% of sample]	weakly substitutes, diversifies in response to risk n=29 [20-37% of sample]
Both altruist and warm glow ($\alpha < 1$)	weakly substitutes, does not diversify in response to risk n=60 [48-68% of sample]	
Pure warm glow ($\alpha = 0$)	maintains constant allocations throughout n=44	

Notes: Refer to Section 2 for an interpretation of α and ρ. "Perfectly (weakly) substitutes" means that the subject gave all (more) of her total gift to the one of her preferred charities with the highest matching rate. "Diversifies in response to risk" refers to subjects who spread their total gift more equally among their preferred charities when exogenous risk over matching funds was introduced; subjects for whom that test was uninformative are counted as responding to risk if they purchased a delivery guarantee. Number of subjects in each cell synthesizes data in Tables 3 and 4. Twenty-one subjects were missing data and 35 gave their entire gift either to their club or a single one of the development charities. Bootstrapped 95% confidence intervals for the proportion of subjects of a given type are shown in brackets, excluding subjects who gave constant allocations throughout.

Panel B: Donor Control

	No agency problems	Agency problems
Pure altruist ($\alpha = 1$)	unrestricted gifts, non-binding n=60 [31-46% of sample]	restricted gifts, binding n=38 [18-31% of sample]
Both altruist and warm glow ($\alpha < 1$)	restricted gifts, non-binding n=55 [28-44% of sample]	
Pure warm glow ($\alpha = 0$)	restricted gifts, agency problems irrelevant indistinguishable from ($\alpha < 1$)	

Notes: Refer to Section 2 for an interpretation of α. "(Non-)Binding" refers to the subject's choice to (give) refuse the charity permission to reallocate restricted gifts. Number of subjects in each cell synthesizes data discussed in Section 6.3. Nine subjects were missing data on gift allocations, 13 subjects kept the full \$100 for their clubs, and 30 subjects gave binding restricted gifts but did not answer the question regarding permission for the charity to reallocate their gifts. Bootstrapped 95% confidence intervals for the proportion of subjects of a given type are shown in brackets.

Table 2: Correlates of Generosity to Development Charities

	All	CARE	Mercy Corps	Oxfam America
	(1)	(2)	(3)	(4)
Male	-.72	-1.08	-.02	-1.03
	(0.6)	(1.48)	(1.03)	(0.96)
Age	-.02**	0.003	-.02	-.03
	(0.01)	(0.03)	(0.02)	(0.02)
Male*Age	0.01	0.006	0.007	0.02
	(0.01)	(0.03)	(0.02)	(0.02)
Previously heard of charity	0.34*	-.19	0.39	0.44
	(0.18)	(0.41)	(0.32)	(0.3)
Previously given to charity	0.15	-.26		
	(0.39)	(0.47)		
Mercy Corps	-.43*			
	(0.22)			
Oxfam America	-.60***			
	(0.22)			
Const.	2.21***	2.01	1.66**	1.83**
	(0.52)	(1.34)	(0.77)	(0.8)
Obs.	370	123	124	117
Log likelihood	-163.08	-32.86	-60	-64.89

Notes: Marginal effects from a probit regression where the dependent variable is a binary indicator equal to 1 if a subject allocated any of her gift to the development charity in question. Data from the first decision of the experiment. Column one includes all three charities (the omitted category is CARE). Prior history of giving to the charity is dropped from columns 3 and 4 because it perfectly predicted gifts in the experiment. The number of observations are less than the number of subjects participating in the experiment due to missing data from the questionnaires, included as explanatory variables in these regressions. Standard errors reported in parentheses, significant at * 90%, ** 95%, and *** 99% confidence.

Table 3: Substitution in Response to Changes in Matching Rates

	Implied values of α & ρ	Summarizing results from decisions 2-3 (unequal matching rates) relative to decision 1 (equal matching rates)
Perfectly substituted	$\alpha = 1$, $\rho = 1$	14
Weakly substituted	$\alpha < 1$ or $\rho < 1$	62
Allocated less efficiently	$\alpha < 1$	23
Never changed allocation	$\alpha = 0$	44
Mixed evidence		6
Preference for 1 charity		23
Gave only to club		12
Missing data		21
Total number of subjects		205

Notes: In decision 2 matching rates were 50¢, \$1, and \$1.50 (order randomly assigned to the three charities); in decision 3 matching rates were 50¢, \$1.25, and \$1 (order paired with the random assignment in the prior decision). Refer to Section 2 for an interpretation of the parameters referenced in the table. Subjects are categorized as having perfectly substituted if in both decisions they moved their entire gift into the one of their preferred charities that had the highest matching rate. Weak substitution occurred when subjects moved some but not all of their gift into a preferred charity with a higher matching rate. Subjects who gave more of their gift to a charity whose matching rate was lower than that of other preferred charities are referred to as having allocated less efficiently. Subjects who perfectly substituted once and allocated less efficiently once are counted as having "mixed evidence" on substitution.

Table 4: Reallocations in Response to Changes in Risk Among Subjects who Gave to Multiple Charities

Panel A

Reallocations across preferred charities:	decision 4 (risk) relative to decision 1 (no risk) (equal matching rates)	decision 5 (risk) relative to decision 2 (no risk) (unequal matching rates)
Divided gift more equally	9	19
Divided gift less equally	9	12
No change in how gift was divided	14	36
Gift initially divided equally	56	20
Missing data necessary for comparison	6	7
Total number of subjects (did not perfectly substitute)	94	94

Panel B

	Implied value of ρ	Summarizing results from decisions 4-5 (risk) relative to decisions 1-2 (no risk)
Evidence of risk aversion	$\rho > 1$	18
Evidence of risk loving	$0 < \rho < 1$	12
Mixed evidence on risk preferences		7
No change in response to risk	$\rho = 1$	11
Uninformative test		41
Missing data necessary for comparison		5
Total number of subjects (did not perfectly substitute)		94

Notes: Subjects are counted as risk averse if they reallocated their gifts across preferred charities more equally in at least one of the decisions and did not consolidate gifts in either decision, risk loving is defined as the reverse, and mixed evidence on risk preferences refers to subjects who were risk averse in one decision and risk loving in the other. Subjects categorized as "No change in response to risk" were those who could have allocated more equally in decision 4 when matching rates were equal but did not and also made no changes in their allocations in decision 5 when matching rates were not equal; those who had no capacity to reallocate more equally in decision 4 are counted as having an "Uninformative test" provided that they did not change their allocations in decision 5.

Table 5: Classifying Preference Types - Professional Subjects

Panel A: Simultaneous Giving to Multiple Charities

	Not Risk Averse ($\rho \leq 1$)	Risk averse ($\rho > 1$)
Pure altruist ($\alpha = 1$)	perfectly substitutes n=22 [8-20% of sample]	weakly substitutes, diversifies in response to risk n=58 [30-45% of sample]
Both altruist and warm glow ($\alpha < 1$)	weakly substitutes, does not diversify in response to risk n=75 [41-56% of sample]	
Pure warm glow ($\alpha = 0$)	maintains constant allocations throughout n=14	

Notes: Refer to Section 2 for an interpretation of α and ρ. "Perfectly (weakly) substitutes" means that the subject gave all (more) of her total gift to the one of her preferred charities with the highest matching rate. "Diversifies in response to risk" refers to subjects who spread their total gift more equally among their preferred charities when exogenous risk over matching funds was introduced; subjects for whom that test was uninformative are counted as responding to risk if they purchased a delivery guarantee. One subject was missing data and 21 gave their entire gift either to their club or a single one of the development charities. Bootstrapped 95% confidence intervals for the proportion of subjects of a given type are shown in brackets, excluding subjects who gave constant allocations throughout.

Panel B: Donor Control

	No agency problems	Agency problems
Pure altruist ($\alpha = 1$)	unrestricted gifts, non-binding n=63 [30-45% of sample]	restricted gifts, binding n=30 [12-23% of sample]
Both altruist and warm glow ($\alpha < 1$)	restricted gifts, non-binding n=79 [38-53% of sample]	
Pure warm glow ($\alpha = 0$)	restricted gifts, agency problems irrelevant indistinguishable from ($\alpha < 1$)	

Notes: Refer to Section 2 for an interpretation of α. "(Non-)Binding" refers to the subject's choice to (give) refuse the charity permission to reallocate restricted gifts. One subject was missing data on gift allocations, 13 subjects kept the full \$100 for themselves, and 5 subjects gave binding restricted gifts but did not answer the question regarding permission for the charity to reallocate their gifts. Bootstrapped 95% confidence intervals for the proportion of subjects of a given type are shown in brackets.

Appendix I - A Parametrized Model

Begin with the assumptions and notation from Section 2. Let the donor's problem be

$$\max_{g_1,\cdots,g_n} \underbrace{\alpha\mathbb{E}\left[\left(\sum_{j=1}^{n} m_j g_j\right)^{1/\rho}\right]}_{\text{altruism}} + (1-\alpha)\cdot\underbrace{\left(\sum_{j=1}^{n}\delta_j f(g_j)\right)}_{\text{warm glow}}$$

$$\text{s.t.}: \sum_{j=1}^{n} g_j \leq D \text{ and } \; g_j \geq 0 \text{ for } j = 1,\cdots,n$$

where the warm glow function $f(g_j)$ converts the utility of giving into comparable units as the altruistic utility of public good production without loss of generality, $\delta_j \geq 0$ allow her to experience different degrees of warm glow from different charities, and $f'(g) > 0$ & $f''(g) < 0$ so that there is diminishing marginal utility of warm glow from each charity she supports.

Clearly, a risk neutral ($\rho = 1$) purely altruistic ($\alpha = 1$) donor maximizes her utility by giving her entire gift to the charity(ies) with the highest marginal public benefit (highest value of m_j). In general, if either $\rho > 1$ (so that the donor is risk averse) or if $\alpha < 1$ and $\delta_j > 0$ for at least two charities (so that the donor experiences warm glow from multiple charities), the concavity of her utility function could lead her to make gifts to multiple charities even if the marginal public benefit of these gifts is not equal.

To see why risk aversion ($\rho > 1$) can lead a donor to split her gift across charities with unequal marginal social benefits, consider a simple example. A purely altruistic donor with $\rho = 2$ is deciding how to divide a \$100 gift between charities j and k. She believes $m_j = 1$ or 0 each with probability .5 and that $m_k = 0$ when $m_j = 1$ and $m_k = \frac{1}{4}$ otherwise. Obviously, a risk neutral donor would give the full \$100 to charity j since it has a higher marginal social benefit in expectation ($\mathbb{E}[m_j] = \frac{1}{2} > \mathbb{E}[m_k] = \frac{1}{8}$). The risk averse donor, on the other hand, is choosing the amount g_j that maximizes her expected utility which is equal to $\frac{1}{2}\sqrt{g_j} + \frac{1}{2}\sqrt{\frac{1}{4}(100 - g_j)}$ where the budget constraint has already been substituted into her utility function. The maximum expected utility she can achieve is 5.6 utils by giving charity j only \$80 and charity k the other \$20, relative to the 5 utils of expected utility that she would get from giving the full \$100 to charity j. Loosely speaking, when deciding whether or not to add a charity to her portfolio, a risk averse donor will consider the effect on both the mean and the covariance of possible outcomes in terms of public good production whereas the risk neutral donor considers only the mean.

For the case of warm glow ($\alpha < 1$), begin by assuming that $\rho = 1$ so that any splitting that arises is due to the warm glow and not risk aversion. Then the first order constraints from the problem above set the following allocation rule when neither non-negativity constraint is binding:

$$\alpha[m_j - m_k] = (1 - \alpha)[\delta_k f'(g_k) - \delta_j f'(g_j)]$$

If $\alpha = 0$ or $m_j = m_k$, the donor divides the gift so as to equate the marginal warm glow from the two charities, giving more to the one with the larger value of δ but at least a token gift to the other charity since the marginal warm glow of such a small gift is very large.[35] If $\alpha > 0$ and $m_j \neq m_k$, the donor considers both the marginal public benefit and the marginal warm glow of her gifts. Of course corner solutions are still possible depending on the exact parameters of the utility function, but in general, if the charity with the higher marginal public benefit also has the higher δ, she will substitute more towards that charity than she would if she were purely warm glow; on the other hand, if the marginal public benefit is higher for the charity with the lower marginal warm glow, she will substitute less toward the socially desirable charity than she would if she were purely altruistic.

Appendix II - Logistics of the Experiment

The experiment was conducted as a pen and paper survey during regularly scheduled club meetings at which the author was the guest speaker.[36] Each subject was given a packet of materials that included an Informed Consent form, an index card with the subject's identification number, and the experimental protocol which was already labeled with the corresponding identification number.[37] Instructions prior to the survey, given verbally as

[35] This may seem implausible, but could be avoided by introducing a small fixed transaction cost to the budget constraint. This would not change the results in any meaningful way.

[36] Participating clubs and dates (all in 2008) were: Walnut Creek Sunrise Rotary (n=8) and Berkeley Kiwanis (n=14) on May 20, Richmond Kiwanis (n=10) on May 21, Berkeley Lions (n=15) on May 22, East Oakland Rotary (n=10) on May 28, San Ramon Kiwanis (n=22) on June 19, El Cerrito Rotary (n=14) on June 26, San Leandro Kiwanis on August 26 (n=14), Walnut Creek Kiwanis on September 2 (n=11), Danville Lions on September 3 (n=7), El Sobrante Rotary on September 16 (n=14), Montclair Lions on September 17 (n=10), Moraga Rotary on September 23 (n=22), Albany Lions on September 25 (n=8), Berkeley Rotaract on October 14 (n=11), and San Pablo Rotary on October 15 (n=15). Rotaract is the Rotary Club's college affiliate. "n" refers to number of subjects; the number of members present at a given meeting was generally higher because not everyone was willing and able to participate. Three sessions of the experiment, each with approximately 30 professional subjects, were run at the U.C. Berkeley XLab on October 2 and 16, 2008.

[37] Copies of the experimental protocol and email / phone script used to recruit the service clubs are available from the author upon request. Professional subjects knew nothing about the nature of the experiment other than the time and location of the session they registered for.

well as printed on the protocol, emphasized that all decisions would be confidential and that one decision for one participant would be chosen to be paid out in real money at the end of the experiment. The author then described each decision verbally as subjects progressed through the experiment.[38] The information purchases were implemented as follows: 1) subjects first recorded whether or not they wanted to purchase information, 2) subjects who did not purchase information recorded their allocations between the three charities (investments), 3) after step 2 was completed for both charities and investments, the author announced which charity (investment) would have the high matching rate, and 4) subjects who purchased information recorded their allocations between the three charities (investments).[39] When everyone was finished, a 12-sided die was rolled to determine which of the decisions would be paid and the identification cards were collected and one was selected at random for payout. The chosen decision and subject's identification number were announced in order to assure subjects that payments would be issued as promised, although only the chosen subject would ultimately be able to verify the payments since no one else knew the identity of the chosen subject or how s/he had allocated the $100 in the chosen decision. The entire process, from signing the consent forms to collecting the protocols, generally took about 20 minutes. At club meetings the author then gave a short 10-minute talk on the motivation for the research project, the way the data would help to answer the research questions, and results from a pilot study conducted at U.C. Berkeley using undergraduates as subjects. As specified by the chosen subject, the author then wrote checks to the host club and/or participating charities which were placed in open but stamped and addressed envelopes so that the chosen subject could confirm the payout amounts before mailing in the donations. These envelopes were then put in a larger manila envelope which also included a de-briefing thank you note from the author and a copy of the Informed Consent form for their personal records. There were identical larger envelopes for subjects whose decisions were not chosen for payout in order to maintain confidentiality; the envelopes were labeled with only the identification numbers and each subject claimed his or her own. Payments to professional subjects were administered by XLab staff in the absence of the author.

[38] Professional subjects worked through the experiment at their own pace without the verbal descriptions of each decision.

[39] During the experimental sessions with professional subjects, the information purchase decisions and allocations based on information were recorded on two separate pages of the experimental protocol. The author collected the purchase decisions and allocations without information before the information was announced, so that there was no way for subjects who had not purchased the information to benefit from it. Logistically, this was not possible at the experimental sessions conducted with the real-world donors, though there is no evidence that these subjects attempted to "steal" information in the absence of the extra safeguard.

It should be noted that despite emphasizing provisions to protect subjects' confidentiality at the beginning of the experiment and immediately prior to announcing which decisions would be paid, real-world donors whose decisions were chosen for payout frequently identified themselves at the conclusion of the experiment. This may be because it was not clear to them how they could receive their payouts without revealing their identities, or it may simply reflect their indifference to protecting the anonymity of their decisions. Since there was no such thing as a "selfish" decision as the experiment was designed, it is likely that subjects felt they had nothing to hide. On a related note, not all chosen subjects actually claimed the manila envelope with the donation checks that were issued based on their decisions. Perhaps the sums involved were too small as to be consequential in their view, with other demands on their time and attention as the meeting was wrapping up, but it is also possible that they never completely understood that their decisions would be converted into real money. That said, several of the chosen subjects who identified themselves to the author at the end of the survey later said that they would "just trust [the author]" to put the checks in the mail. There were no such abnormalities regarding confidentiality or claim of payments at experimental sessions with professional subjects.

Appendix III - Charity Mission Statements

As listed on CharityNavigator.org, a major independent charity evaluator:

Founded in 1945, **CARE** fights root causes of poverty in the world's poorest communities. CARE places special focus on working alongside poor women because, equipped with the proper resources, women have the power to help whole families and entire communities escape poverty. In 70 countries, women are at the heart of CARE's community-based efforts to improve education, prevent the spread of HIV, increase access to water and sanitation, expand economic opportunity and protect natural resources. Each year, CARE helps tens of millions of people around the world effect real, positive changes in their lives.

Mercy Corps exists to alleviate suffering, poverty and oppression by helping people build secure, productive and just communities. Since 1979, Mercy Corps has provided $1.3 billion in assistance to people in 100 nations. Mercy Corps pursues its mission through: emergency relief services that assist people afflicted by conflict or disaster; sustainable economic development that integrates agriculture, health, housing and infrastructure, economic development, education and environment, and local management; and civil society initiatives that promote citizen participation, accountability, conflict management and the rule of law.

Founded in 1970, **Oxfam America**, an affiliate of Oxfam International, works to end global poverty through saving lives, strengthening communities, and campaigning for change. Oxfam America works on the scene, helping people gain the hope, skills, and direction to create a new future. We are also active in the global arena, addressing social injustice through our advocacy, public education, and emergency assistance programs. While Oxfam employs a variety of strategies to achieve our mission, the goal in all our endeavors is the same: to enable poor people to exercise their right to manage their own lives. The specific issues we work on include making a living, natural resources, peace and security, equality for women, indigenous and minority rights, and global trade.

References

Andreoni, J., "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *Journal of Political Economy*, 1989, *97(6)*, 1447–1458.

__ , "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," *Economic Journal*, 1990, *100(401)*, 464–477.

__ , "An Experimental Test of the Public-Goods Crowding-Out Hypothesis," *The American Economic Review*, 1993, *83(5)*, 1317–1327.

__ **and A. Payne**, "Crowding Out Both Sides of the Philanthropy Market: Evidence from a Panel of Charities," *Working Paper*, 2008.

__ **and J. Miller**, "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism," *Econometrica*, 2002, *70(2)*, 737–753.

Benartzi, S. and R. Thaler, "Naive Diversification Strategies in Defined Contribution Saving Plans," *The American Economic Review*, 2001, *91(1)*, 79–98.

Camerer, C., *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, 2003.

Charness, G. and M. Rabin, "Understanding Social Preferences with Simple Tests," *The Quarterly Journal of Economics*, 2002, *117(3)*, 817–869.

Crumpler, H. and P. Grossman, "An Experimental Test of Warm Glow Giving," *Journal of Public Economics*, 2008, *92*, 1011–1021.

Darby, M. and E. Karni, "Free Competition and the Optimal Amount of Fraud," *Journal of Law and Economics*, 1973, *16(1)*, 67–88.

Desai, R. and H. Kharas, "The California Consensus: Can Private Aid End Global Poverty?," *Survival*, 2008, *50(4)*, 155–168.

Duflo, E., "Scaling Up and Evaluation," *Annual World Bank Conference on Development Economics 2004: Accelerating Development*, 2004, *1(30228)*.

Duncan, B., "A Theory of Impact Philanthropy," *Journal of Public Economics*, 2004, *88*, 2159–2180.

Eckel, C. and P. Grossman, "Rebate versus Matching: Does How We Subsidize Charitable Contributions Matter?," *Journal of Public Economics*, 2003, *87*, 681–701.

Fisman, R., S. Kariv, and D. Markovits, "Individual Preferences for Giving," *American Economic Review*, 2007, *97(5)*, 1858–1876.

Fong, C., "Evidence from an Experiment on Charity to Welfare Recipients: Reciprocity, Altruism and the Empathic Responsiveness Hypothesis," *The Economic Journal*, 2007, *117(July)*, 1008–1024.

Giving USA Foundation, "Giving USA 2007: The Annual Report on Philanthropy for the Year 2006," researched and written by the Center on Philanthropy at Indiana University Indianapolis, IN 2007.

__ , "Giving USA 2008: The Annual Report on Philanthropy for the Year 2007," researched and written by the Center on Philanthropy at Indiana University Indianapolis, IN 2008.

Greenlee, J., M. Fischer, T. Gordon, and E. Keating, "An Investigation of Fraud in Nonprofit Organizations: Occurrences and Deterrents," *Nonprofit and Voluntary Sector Quarterly*, 2007, *36(4)*, 676–694.

Hoeffler, A. and V. Outram, "Need, Merit or Self-Interest – What Determines the Allocation of Aid?," *Center for the Study of African Economies at the University of Oxford Working Paper Series*, 2008, *2008-19*.

Jacobsson, F., M. Johannesson, and L. Borgquist, "Is Altruism Paternalistic?," *The Economic Journal*, 2007, *117(April)*, 761–781.

Karlan, D. and J. List, "Does Price Matter in Charitable Giving?: Evidence from a Large-Scale Natural Field Experiment," *American Economic Review*, 2007, *97(5)*, 1774–1793.

Kerlin, J. and S. Thanasombat, "The International Charitable Nonprofit Subsector: Scope, Size, and Revenue," Urban Institute Policy Brief September 2006.

Landry, C., A. Lange, J. List, M. Price, and N. Rupp, "Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment," *Quarterly Journal of Economics*, 2006, *121(2)*, 747–782.

List, J. and D. Lucking-Reiley, "The Effects of Seed Money and Refunds on Charitable Giving: Experimental Evidence from a University Capital Campaign," *Journal of*

Political Economy, 2002, *110*, 215–233.

Meier, S., "Do Subsidies Increase Charitable Giving in the Long-Run? Matching Donations in a Field Experiment," *Boston Federal Reserve Board Working Paper*, 2007.

Reinstein, David, "Empirical Studies on Consumer Demand and Philanthropy." PhD dissertation, University of California, Berkeley 2006.

Ribar, D. and M. Wilhelm, "Altruistic and Joy-of-Giving Motivations in Charitable Behavior," *Journal of Political Economy*, 2002, *110(2)*, 425–457.

Rollins, K., "The Index of Global Philanthropy," Hudson Institute 2007.

Rondeau, D. and J. List, "Matching and Challenge Gifts to Charity: Evidence from Laboratory and Natural Field Experiments," *Experimental Economics*, 2008, *11*, 253–267.

Savedoff, W. and R. Levine, "Learning from Development: the Case for an International Council to Catalyze Independent Impact Evaluations of Social Sector Interventions," *Center for Global Development Brief*, 2006, *May.*

Siegel, Sarah, "The Effect of Choice on Charitable Contributions: A Field Experiment." PhD dissertation, Massachusetts Institute of Technology 2006.

The Center on Philanthropy at Indiana University, "American Express Charitable Gift Survey," November 2007.