Giving gifts to groups: How altruism depends on the number of recipients [☆]

James Andreoni

University of California, San Diego

Abstract

When a single gift goes to a group of recipients, how does giving depend on the size of the group? This question is important for understanding charitable giving and fund-raising, public goods provision, family altruism, and more. If we think of the gift as giving up a dollar to create a social surplus, then we want to know how the number of recipients of that surplus affects its value to the giver. In other words, how congestible is altruism? This paper builds a theoretical framework for this question and begins to answer it with a controlled experiment. The finding is that for most subjects altruism is congestible. For the average subject, a gift that results in one person receiving x is equivalent to one in which n people receive $x/n^{0.68}$ each.
© 2007 Elsevier B.V. All rights reserved.

1. Introduction

How much would you give up if each dollar you forfeited benefited another person by $10? What if instead your dollar gave $5 each to two others, or $2 each to five others, or $1 each to ten others? In each case the private costs and social benefits are the same, but would we view these as the same situation? Similarly, what if each dollar you gave up benefited one other by $1 each, or benefited 2 others by $1 each, or 3 others, and so on? If the same private cost generates benefits for more people, do you spend more or less on giving?

These questions are important for understanding gift giving when there are groups of recipients. Consider two public goods, one at a local level and one at a national level. Suppose the local good

---

provides high value to a few people, the national good provides low value to many, and that the total value of these is identical. Which would an altruist prefer? Another example would be charitable solicitations from both the local chapter and the national organization of a charity. A particular donation might affect more people at a higher level. Would donors give more or less as a result?

Suppose instead the gift is to a food bank. There is a finite amount of food, and more recipients means each gets less. Do you give more or less as the number of recipients increases? Another example is a restaurant staff that pools its tips. A given tip will now only go partly to your waiter, and partly to the others. Does that make people more or less generous? Or consider a parent thinking of buying a new television for the family. Will she spend more, all else equal, if the household has 4 members rather than 3?

These questions are also relevant for experimental research. If a generous act in an experiment, such as giving to a public good or striking a fair bargain, is shared with more people at the same cost, do people value it more or less?

Economic theory can give us some guidance in answering these questions. Consider the first counterfactual above: The same social surplus is divided among more individuals. If people care only about the *total* amount received, then increasing the number of others, n, will have no effect. If, at the opposite extreme, they care about the *average* amount received, then increasing n increases the effective price of giving, which means they may reduce giving.

Consider the second counterfactual: Increasing the group size expands the surplus, but keeps the average benefits the same. Increasing n here means that any dollar given up will generate more benefits, indicating people should be more generous as n increases—an income effect. On the other hand, the same social benefit can be attained at lower cost. This means people may give less away—a substitution effect. Which effect dominates will depend on preferences of individuals.

A unifying approach to this is to frame the issue in terms of congestion, as in the literature on congestible public goods. The congestion here, however, is in the hearts of the altruists, rather than in the technology of the public good. Think first of the total social surplus generated for others by the dollars forfeited. Call this total surplus Π_o. Then let $\pi_o = \Pi_o/n$ be the average surplus for others. If the individual giver views the gift, g, as being the *total* social surplus, independent of n, then this is a good that is not congested at all and $g_0 = \Pi_o = n\pi_o$. On the other extreme, if the giver thinks of the gift as the *average* social surplus, then the good is completely congested, $g_1 = \Pi_o/n = \pi_o$. As with congestible public goods, actual congestion may be better described as a combination of these two, $g = g_0^b g_1^{1-b} = n^b \pi_o$, with $b \in [0, 1]$. This paper will be about identifying b.

There have been no field studies and only a few experimental studies of altruism and the effects of group size. Comparing groups of size two and four in public goods provision, Goeree, Holt and Laury (2002) found no clear effect. Isaac and Walker (1988) and Isaac, Walker and Williams (1994) compared public goods games with groups of size 4, 10, 40 and 100, with varying marginal returns to giving. They found no effect of group size when the price was low (0.33), but when the price of giving was high (2.33) groups of 4 were less generous than groups of 10—the income effect dominates the substitution effect—but the effect largely vanished when comparing groups of 40 and 100. The experiments of Kahneman and Knetsch (1992) find similar effects. When subjects were asked their willingness to pay for a local public good, followed by asking about a regional or national public good, giving rose.[1]

---

[1] Others study group size and giving to discrete or step level public goods. These are quite different from what is in mind here, as these experiments usually require a minimum total contribution or number of contributors in order to move the provision of a public good from zero to a single positive value. See Bagnoli and McKee (1991), Croson and Marks (1998), for instance.

Although the main interest in studying congestion in giving is, obviously, real world giving, we begin the study of it here with a controlled laboratory experiment. The effects we find could help shape the field studies necessary to demonstrate the validity of the laboratory findings.

The experiment reported in this paper offers subjects a series of 24 decisions which vary the number of recipients and the "price" of generosity. The experiment builds on earlier work by Andreoni and Miller (2002) who used similar techniques to uncover preferences for altruism, but who kept the number of recipients fixed at one.

We will use three approaches to uncover preferences. First is a non-parametric revealed preference approach that holds n fixed. Next is a semi-parametric approach. This begins by hypothesizing a parametric "giving" function that enters utility, and then by using non-parametric analysis of utility to find the best parameters on the giving functions. That is, we find the parameter in the giving function that results in no violations (or the smallest violations) of revealed preference. Third is a fully parametric approach to the estimation of utility functions. With each approach we describe and catalogue the variety of preferences exhibited by our subjects.

The findings are, first, that there is a remarkable degree of consistency within subjects, and surprisingly few violations of the axioms of revealed preference. Second, while there are many examples of individuals who behave as though $b=0$ or $b=1$, the majority of subjects are best described by a b between 0 and 1. Combining the entire sample and estimating a single representative utility function, we estimate the best aggregate b to be $b=0.68$. This means that, whether looking at the problem non-parametrically, semi-parametrically or fully parametrically, or whether it is individual or aggregate analysis, we must conclude that altruism is a congestible good, but the congestion is partial—doubling the number of recipients increases but does not double the value of altruism to the giver.

The paper is organized as follows. The next section will present the formal framework of congestible altruism. Section 3 will present the experimental design. We present the general results in Section 4, followed by non-parametric and semi-parametric analysis of the data using revealed preference techniques in Section 5. Sections 6 and 7 turn to parametric analysis, presenting estimates of utility for each individual subject as well as for a representative altruist. Section 8 concludes.

2. Model

Let π_i be the marginal payoff to individual i and let π_o be the additional payoff to each of the others as a consequence of i's choice. Again, let n be the number of others affected. Then the maintained hypothesis is that individuals have utility $u_i = u_i(\pi_i, \pi_o; n)$, where preferences are convex in π_i and π_o, and (weakly) increasing in n. It is also natural to expect that u is increasing in both π_i and π_o, as found by Andreoni and Miller (2002).[2]

Suppose utility is separable on (π_o, n). Then there is some function $g(\pi_o, n)$ such that utility can be written as $u_i(\pi_i, \pi_o, n) = u_i(\pi_i, g(\pi_o, n))$. We can think of g abstractly as an individual's "giving" function, that is, the particular aggregation of π_o that this person cares about.[3] The question for

---

[2] See also Fisman, Kariv and Markovits (in press). With downward sloping budgets it is impossible to identify preferences where π_o may be a "bad" at the margin. However, offering some upward sloping budgets, Andreoni and Miller (2002) found that there was a significant minority of subjects for whom the indifference curve may actually "bend back" for some extreme inequality. Andreoni, Castillo and Petrie (2003) confirm this for strategic games as well. Since we are concerned with voluntary altruism here, we do not concern ourselves with such "jealous" or "resentful" preferences, but this is an interesting area for future research.

[3] The main restriction imposed by the assumption of separability is that how people aggregate giving is independent of their own level of consumption, that is, the marginal rate of substitution between π_o and n is independent of x. In our experiments, a subject's income varies only slightly, so this is unlikely to be a costly assumption. Whether it is a valid assumption for field data is an interesting question.

research is to identify these aggregations, $g(\pi_o, n)$, as well as learn how utility changes as the value of this aggregation changes.[4] Thinking first about $g(\pi_o, n)$, at one extreme we could have $g(\pi_o, n) = \pi_o$, that is, people care only about the average received. At the other extreme we could have $g(\pi_o, n) = n\pi_o$, so people care about the total amount received. What about something intermediate? In the literature on congestion in public goods, there is a history of estimating "crowding" with a parameter, b, and a function, $g(\pi_o, n) = n^b \pi_o$. Thus for $b = 0$ or $b = 1$ we have the extreme cases, while for $b \in [0, 1]$ we allow intermediate cases.[5]

Consider a person facing a linear budget over payoffs. For simplicity we will always treat π_i as the numeraire, and use p as the price of π_o. Thus, write the optimization problem as

$$\max_{\pi_i, \pi_o} u_i(\pi_i, n^b \pi_o)$$
$$\text{s.t. } \pi_i + p\pi_o = m$$

It is convenient to think of the person as choosing $g = g(\pi_o, n)$ rather than π_o. Thus, substitute $\pi_o = g/n^b$ and let the price be p/n^b. Then we can rewrite this optimization as

$$\max_{\pi_i, g} u_i(\pi_i, g)$$
$$\text{s.t. } \pi_i + \frac{p}{n^b} g = m \tag{1}$$

Let $p' = p/n^b$. Solving Eq. (1) we get a demand function $g(p', m)$. As usual, assume $\partial g / \partial p' < 0$. Then we can measure total giving as $p\pi_o = p'g$. How does this change with n? Taking the derivative we see

$$\frac{\partial p\pi_0}{\partial n} = \left(p' \frac{\partial g}{\partial p'} + g \right) \frac{\partial p'}{\partial n} \lessgtr 0.$$

We can now see the income and substitution effects of n on giving, π_o. As n increases, the price of g falls but the total amount of benefit rises, and the effect on total spending is ambiguous.

With this as our setup, we use the axioms of revealed preference to identify whether, for a given b, preferences are consistent with maximizing behavior. We do this by applying the Generalize Axiom, GARP. We say a bundle (π_{ia}, g_a) is *directly revealed preferred* to a bundle (π_{ib}, g_b) if $\pi_{ia} + p'_a g_a \geq \pi_{ib} + p'_a g_b$. The relation is *strict* if the inequality is strict. We then say a bundle (π_{ia}, g_a) is *revealed preferred* to a bundle (π_{ib}, g_b) if there is a chain of directly revealed preferred relations that links the two. That is, the revealed preferred relation is the transitive closure of the directly revealed preferred relation. Then define GARP this way (Varian, 1982):

Definition 1. Generalized Axiom of Revealed Preference (GARP): If (π_{ia}, g_a) is revealed preferred to (π_{ib}, g_b) then (π_{ib}, g_b) is never strictly directly revealed preferred to (π_{ia}, g_a).

---

[4] Note that we are generalizing the model of "warm-glow giving" (Andreoni, 1989, 1990) to mean that utility of giving depends on the number of recipients. With n constant (or as we see later, n large) the distinction is irrelevant. The analysis could be generalized to a strategic environment where altruism also depends on the contributions of others, but the fundamental insights of this model would be unaffected.

[5] Suppose, for instance, that $g(\pi_o, n)$ combines the two extreme notions in a Cobb–Douglas production function for giving, so $g(\pi_o, n) = (n\pi_o)^b \pi_o^{1-b} = n^b \pi_o$. On the genesis of the literature on congestible public goods, see Bergstrom and Goodman (1973), for instance. For a general treatment of public goods and group size, see Bergstrom, Blume and Varian (1986) and Andreoni (1988a).

As Varian (1982) shows, if the data is consistent with GARP, then there exists a utility function that could have generated the data. Checking GARP will be at the heart of our analysis.

2.1. Interpreting the price

How do we think about price p in this setting? Recall, we constructed π_o as the value per recipient. Thus, if giving is to a pure public good, the average value can simply be measured as the cost to the giver. In this case, absent any subsidies or other distortions, we can assume $p=1$.

What if instead the gift is a pure private good, as in the example of pooling of tips? If π_o is the tip received by each of the waitstaff, then total spending on tips is $n\pi_o$. This means $p=n$.

It is possible, as well, that the gift provided is itself a congestible public good. Perhaps the gift provides for a seminar series in the economics department. If seminars are crowded the pleasure of the experience is decreased for all. Following the literature on congestion in public goods, we would say the total surplus is $\Pi=n^c\pi_o$, where $0 \le c \le 1$. In order to generate a given π_o in average value, therefore, a person will need to spend $n^c\pi_o$ dollars. Thus, $p=n^c$, meaning that techno-logical congestion has the effect of raising the price of altruism.

In the experiments below we enforce giving as a pure public good, but we artificially vary the price in order to allow us to identify preferences. When applying these results to the real world, however, we need to be thoughtful about what is determining the price.

3. Experimental design

The restatement of the problem in Eq. (1) indicates that we will need variation in n, m, p and p/n in order to identify both the utility function and the parameter b. With this in mind, we designed a task with 24 decisions. The decisions were presented to the subjects on a computer screen like that shown in Fig. 1. Subjects filled in the blanks in sentences like this: "Divide 60 tokens: *Hold* _ _ _ @ 20 cents, and *Pass* _ _ _ @ 10 cents, to each of 2 other people". The computer automatically made sure that choices were on the budget constraint. It also calculated
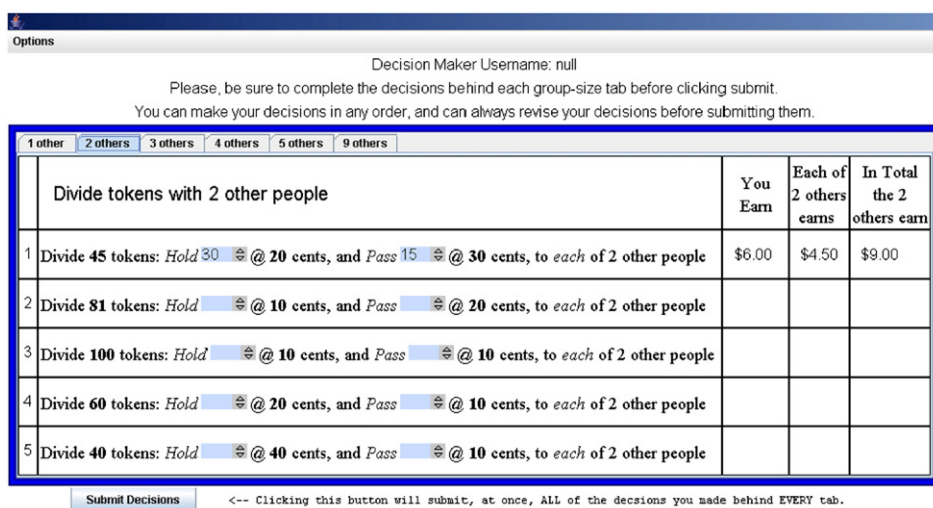

Fig. 1. Sample decision screen.

the earnings for the subjects in the study, and reported them in the columns on the right, as seen in Fig. 1. It calculated earnings in two ways. First is average earnings ("Each of n others earns") and second is total earnings ("In Total the n others earn"), as can be seen in the figure.

The decisions were presented on pages organized by n. Subjects, however, were free to tab back and forth across pages and could fill in blanks in any order. They could not submit their decisions until all blanks were filled. To check for possible order effects we presented the game in two orders, forward, as shown in the figure, and reverse, where the tabs and the decisions on each tab were all in reverse order.

For any choice, we can calculate p and m as follows. The task is to divide M tokens with a hold value of H cents and a pass value of P cents. Let X be the amount held and Y be the amount passed, adhering to the token constraint $X+Y=M$. Then $\pi_i=HX$ and $\pi_o=PY$. Substitute these into the token constraint and rearrange to get

$$\pi_i + \frac{H}{P}\pi_o = HM.$$

Finally, substituting $g(\pi_o, n)=n^b\pi_o$, and defining $p=H/P$ and $m=HM$, we rewrite this as the budget constraint

$$\pi_i + \frac{p}{n^b}g = m,$$

as in Eq. (1) above. Thus, m has the interpretation of the earnings by player i if he passes zero and keeps M, and p is the cost to person i of increasing the payoff of each other person.

The full set of experimental parameters is shown in Table 1. There are six possible values for the number of others, $n \in \{1, 2, 3, 4, 5, 9\}$. Likewise p takes on 10 possible values, and p/n takes on 12. Care was taken to assure that there would be several sets of budgets that would be identical on p and m but differ on p/n (budgets 4 and 12, 5 and 9, 10 and 21), and that others would be the same on p/n and m but differ on p (budgets 6 and 12, and 11 and 22). This will help with identifying b apart from utility parameters. We also assured that all budgets crossed within a given n.

While subjects made 24 choices, they were only paid for one choice selected at random.[6] Each choice was equally likely to be chosen. Subjects were told this prior to their decisions.

We were also interested in the possible effect of reciprocal altruism. For instance, is there a difference in giving when it can be reciprocated, as in a public good, and giving when it is explicitly non-reciprocal? We explored this with two matching conditions. Under *reciprocal matching* people are explicitly put into groups and the same decision is carried out for each member of the group. We tell subjects that those who receive their pass amounts are also in a position to pass to them.[7] We weighted the probabilities of choosing different groups so that each decision is equally likely to be chosen. We compare this to *circular matching*, which only allows for unilateral altruism.[8] After choices were made, the computer program randomly placed subjects in a "circle".

---

[6] Hey and Lee (2005) present evidence to suggest that randomly selecting one of a subject's choices is successful in getting the subject to treat them as separate independent events, rather than as a portfolio of random events.

[7] For the reciprocal matching the instructions read: "You will also be randomly selected to be a recipient of others' Pass allocations. The recipients of your Pass portion will be those who will be passing to you. That is, any passed tokens you receive will be from the same other participants who may receive tokens passed by you. Moreover, these other participants will have the same decision chosen by the computer that you do. For example, suppose the computer selects this decision for you and 4 others: Divide 80 tokens: *Hold* _ _ _ _ @ 20 cents, and *Pass* _ _ _ _ @ 30 cents to each of 4 other people. Then you will be in a group of 5 people. Each of the 5 has the same decision chosen, and each is potentially passing tokens to the same other members in that group".

[8] For the circular matching the instructions read: "You will also be randomly selected to be a recipient of others' Pass allocations. The recipients of your Pass portion will not be the same participants who will be passing to you. That is, you will not receive any passed tokens from those other participants who may receive tokens passed by you".

Table 1
Experimental parameters and average passed

Budget	n	M	Hold value	Pass value	p/n	p	m	Circular tokens	Reciprocal tokens
			Experimental parameters [a]					Average passed [b]	
1	1	40	10	40	0.25	0.25	400	14.6	15.9
2		48	10	30	0.33	0.33	480	17.3	18.6
3		60	10	20	0.50	0.50	600	17.1	22.9
4		90	10	10	1	1	900	19.9	24.2
5		60	20	10	2	2	1200	12.1	14.3
6	2	45	20	30	0.33	0.67	900	15.1	18.1
7		81	10	20	0.25	0.50	810	27.7	32.2
8		100	10	10	0.50	1	1000	27.5	32.2
9		60	20	10	1	2	1200	12.8	16.7
10		40	40	10	2	4	1600	7.0	9.3
11	3	40	10	30	0.11	0.33	400	15.8	16.6
12		90	10	10	0.33	1	900	29.9	31.4
13		40	30	20	0.50	1.5	1200	10.9	12.1
14		52	30	10	1	3	1560	10.8	12.6
15	4	80	10	10	4	0.25	800	28.7	28.0
16		72	20	10	2	0.50	1440	22.3	23.5
17		42	10	20	0.5	0.125	420	17.3	17.9
18		50	40	10	4	1	2000	11.4	12.3
19	5	60	10	10	0.20	1	600	20.7	21.3
20		51	20	10	0.40	2	1020	15.5	17.5
21		40	40	10	0.80	4	1600	10.3	11.2
22	9	40	10	10	0.11	1	400	16.6	13.5
23		30	20	10	0.22	2	600	10.6	10.0
24		20	40	10	0.44	4	800	6.3	6.6

[a] M is reported in tokens. Hold value, Pass value and m are reported in cents.
[b] Average across all subjects, pooled across the two orders.

For each subject i a choice among the 24 made was selected at random, and the n people to the right of i were the recipients. Hence, this game assures subjects that none of the people who receive their generosity will have an option to give anything back, and this is explained to all the subjects.

We ran 6 sessions of 20 subjects each, for a total of 120 subjects in this study. Three sessions were run with circular matching and three with reciprocal matching. Within each matching type, two sessions were run with the forward presentation and one with the reverse presentation. Each session lasted less than 1 h, and subjects earned an average of $14.23. Complete instructions are available from the author.[9]

We should note that this design differs from prior public goods experiments. Earlier experiments, typified by Isaac and Walker (1988), usually examine a single group size and a single price during any session, plus they usually have repeated interactions either with the same group or with "strangers" rematching (Andreoni, 1988b, 1995). This experiment, by contrast, asks for a series of decisions without any feedback, and compares choices over a number of different prices and group sizes.

4. Overall results

Table 1 shows the average amount passed on each budget for both the circular and reciprocal treatments. These results are also shown in Figs. 2 and 3. Both figures show the budgets for (π_i, g).
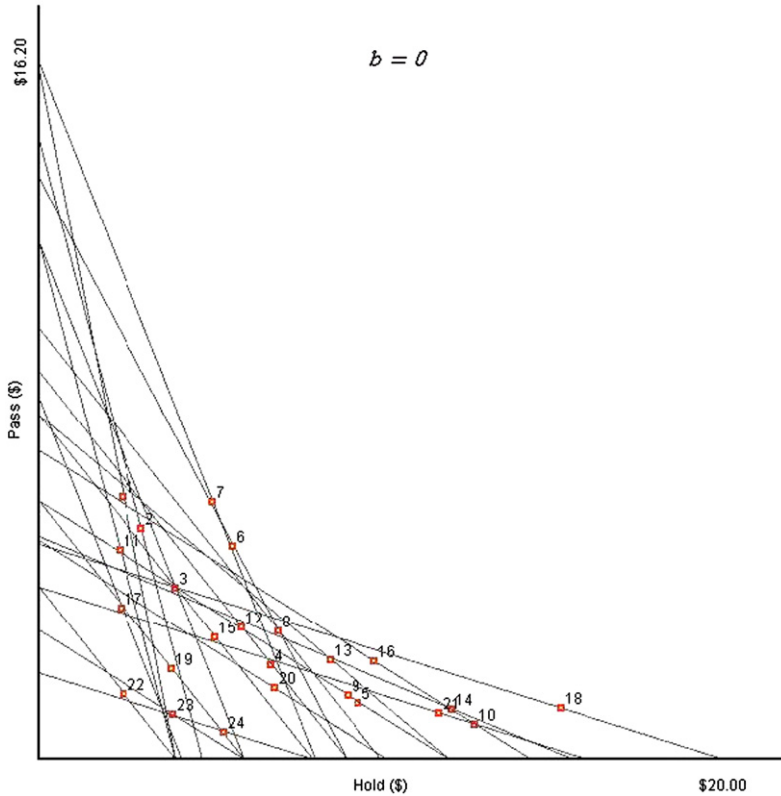
---

[9] Go to http://econ.ucsd.edu/~jandreon/.

Fig. 2. Budgets assuming $g(\pi_o, n) = \pi_o$.

Fig. 2 shows the average choice under the assumption of $b=0$, while Fig. 3 assumes $b=1$. As can be seen, both show sensitivity to price and income in the predicted way.[10]

Next ask whether there were effects of either the order of decisions or the matching of subjects. We conducted Mann–Whitney tests on the distributions of average hold amounts across subjects. Keeping the matching constant, we found no difference for order for either the circular or reciprocal matching, or for the whole sample combined. Combining all 120 subjects, for instance, we find $z=1.22$ ($\alpha \leq 0.22$).[11] This gives us confidence that the data is not biased by the presentation of the decision task. Comparing the matching schemes, we again employ Mann–Whitney tests and, perhaps surprisingly, we again found no significant difference between matching rules. Comparing all subjects in circular matching to those reciprocal matching, we find $z=0.33$ ($\alpha \leq .74$).[12] This is evidence that there is no real effect of reciprocal versus unilateral altruism. These two findings indicate that we can safely pool all 120 subjects into a single sample. We will nonetheless report

---

[10] Although aggregate data need not satisfy GARP, there are no violations of revealed preferences for $b=0$, while for the $b=1$ aggregation there is one violation involving budgets 17 and 22. These two budgets are nearly identical when $b=1$ and the choices nearly overlap at the intersection.

[11] For reciprocal matching, $z=0.679$ ($\alpha \leq 0.49$) and for circular matching $z=0.88$ ($\alpha \leq 0.38$).

[12] This also holds when we look at the forward and reverse orders separately. For forward order, $z=0.48$. ($\alpha \leq .63$) and for reverse order $z=0.13$ ($\alpha \leq 0.89$).
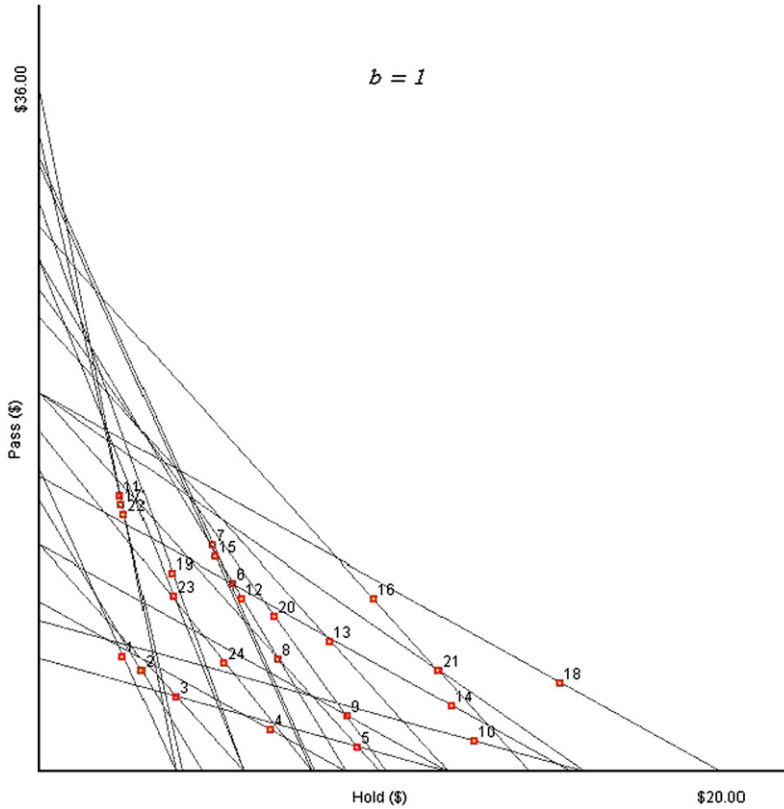
Fig. 3. Budgets assuming $g(\pi_o, n) = n\pi_o$.

results separated by condition, when convenient. As we will see, however, the behavior in all four conditions is remarkably similar throughout the analysis.

4.1. The power of the GARP tests

There are several methods for looking at the power of revealed preference tests.[13] The easiest are the Bronars index (Bronars, 1987) and a bootstrapping index (Andreoni and Harbaugh, 2006). Bronars specifies as the alternative hypothesis that choices are drawn randomly from a uniform distribution on each budget. The bootstrapping index also draws randomly on each budget, but instead uses the actual distribution of choices made on each budget by the pooled set of subjects. After creating synthetic subjects, GARP tests are conducted. The process is then iterated a sufficient number of times.

Applying these power tests to our data requires that we also specify a giving function, $g(\pi_o, n)$. Choosing $g(\pi_o, n) = n^b \pi_o$, we calculated both power indices for 41 values of b between 0 and 1. In each case, using 10,000 random iterations, over 99.6% of all randomly generated data had violations of GARP, for both Bronars and bootstrapping indices.[14] Given these methods of randomly generating

---

[13] See Andreoni and Harbaugh (2006) for a complete discussion of the power of revealed preference tests.

[14] This is true whether boostrapping was conditional on circular matching, reciprocal matching, or pooled.

alternative organizations of the data, these two tests indicate that we have designed a powerful test of consistent preferences.[15]

5. Non-parametric analysis

This section begins by asking whether choices are consistent with GARP. We first hold group size constant and check GARP conditional on n. We then combine all the data and try to uncover the optimal b in $g(\pi_o, n) = n^b \pi_o$ for each subject.

5.1. Checking GARP within group size

The advantage of checking GARP within a group size is that we need no assumptions on how preferences depend on n. Table 2 shows the number of violations of GARP conditional on each n and the matching scheme. Within each n only a small number of subjects had any violations of GARP. Of those violations, only a fraction were "significant" as measured by the Afriat critical cost efficiency index, $\mathcal{A}_i$. This index is a measure of goodness of fit. It indicates how much a budget has to be "relaxed" to avoid violations of GARP. The index has the interpretation that a person could have purchased a revealed preferred bundle at a fraction $\mathcal{A}_i$ of what was spent.[16] If $\mathcal{A}_i = 1$ there are no violations, and the lower $\mathcal{A}_i$ the worse the violation. The standard proposed by Varian for a significant violation of GARP is $\mathcal{A}_i < 0.95$. Of the 30 violations listed in Table 2, 12 are significant by this definition, which come from only nine individual subjects. This means 92% of the subjects have choices that are consistent based on these tests. This is a very high degree of rationalizability.

5.2. Checking GARP across group sizes

We next turn to evaluating GARP when we pool across group sizes. Whether data is consistent with GARP will depend on the definition of the giving function $g(\pi_o, n) = n^b \pi_o$. Here we ask, does there exist a $b \in [0, 1]$ for an individual i such that there are no violations of GARP?

The answer to this question is yes for most subjects. For both the circular and reciprocal matching, exactly 83 subjects, 69%, have choices that did not violate GARP a single time for at least one value of b. Those for whom the answer is no are summarized in Tables 3 and 4. Here we report the value of b that resulted in the highest value of the Afriat index $\mathcal{A}_i$ for all the subjects that always violate GARP. Of the 37 subjects with violations, 16 have indices above the $\mathcal{A}_i \geq 0.95$ threshold, and 20 below. By this measure, 16% of subjects have severe violations of GARP, and 84% are rationalizable.[17] Within the context of revealed preference experiments, this is a relatively high degree of agreement with the model, especially for such a large number of budget constraints.[18]

---

[15] We also conducted Bronars (1987) power over each subset of budgets with the same n. This will necessarily reveal lower power since there are fewer budgets. For $n = 1, 2, 3, 4, 5, 9$ the Bronars power finds the fraction of iterations with GARP violations is, respectively, 0.7, 0.60, 0.50, 0.36, 0.30, 0.36. This was over 10,000 iteration each. To get a perspective on how high these numbers are, note that with two budgets the *maximum* value a Bronars index can take is 0.25.

[16] See Afriat (1972) and Varian (1982) for discussions of this index.

[17] We also conducted further tests of differences between circular and reciprocal matching and again found no significant differences. Comparing the b's reported in Tables 3 and 4 with a Mann–Whitney test, $z = 0.127$. Comparing the $\mathcal{A}_i$'s reported, $z = 0.427$. Finally, comparing the number of GARP violations, $z = 0.174$.

[18] See Andreoni and Harbaugh (2006) for a discussion of recent experiments using revealed preference tests. An intresting recent comparison is Fisman, Kariv and Markovits (in press) who provide subjects with 50 choices of dividing a pie with one other subject. They find only 10.2% of subjects had no violations of GARP, and 41% were above the $\mathcal{A}_i \geq 0.95$ threshold.

Table 2
Violations of GARP, holding number of others n constant

	Circular matching		Reciprocal matching	
n	Number with GARP violations	$\mathcal{A}_i$ Afriat efficiency[a]	Number with GARP violations	$\mathcal{A}_i$ Afriat efficiency
1	2	0.73, **0.96**	2	0.84, **0.97**
2	5	**0.97, 0.97, 0.99, 0.99, 0.99**	4	0.84, 0.88, **0.95, 0.98**
3	1	0.87	4	0.84, **0.96, 0.97, 0.99**
4	1	**0.95**	2	0.83, **0.96**
5	3	0.75, **0.95, 0.98**	1	**0.96**
9	2	0.88, **0.99**	3	0.88, 0.88, 0.88
Total violations	14		16	
No. of unique subjects	10		11	

[a] Bold indicates $\mathcal{A}_i \geq 0.95$.

This leads to the question, what values of b lead to consistent behavior for those subjects without violations of GARP? Unfortunately, there is typically no single b that leads to $\mathcal{A}_i = 1$, but often many b's can be rationalized. To summarize the information across subjects, we evaluated each subject for violations of GARP for 41 values of b between 0 and 1. If the subject had no violations for *a b*, we called that b "admissible" for that subject.[19] That is, the b could rationalize the data. We then add up how often a value of b was admissible across subjects for the pool of subjects for whom at least one b leads to $\mathcal{A}_i = 1$. This is summarized in Fig. 4.

In Fig. 4 we see the vast heterogeneity across subjects, but also the weakness of this semi-parametric approach in identifying the best values of b. There are too many subjects that have a wide array of admissible values of b. Making a more precise statement about the distribution of "optimal" b's will require that we turn to parametric analysis.

6. Estimating utility

This section begins by asking whether the altruism shown here is easily identifiable as one of the standard social welfare functions. We then go on to estimate parametric utility functions for each subject.

6.1. Identifying the pure types

Andreoni and Miller (2002) identified that a substantial fraction of subjects have preference that conform to one of three "pure types" when $n = 1$. These are perfectly selfish, $u(\pi_i, g) = \pi_i$, utilitarian or perfect substitutes, $u(\pi_i, g) = \pi_i + g$, and Rawlsian or Leontief, $u(\pi_i, g) = \min\{\pi_i, g\}$. We ask here if, for each subject, there is some b such that choices match one of these types.

---

[19] This implies, of course, that a subject can have many admissible values of b. A subject who always passes 0, for instance, is fit by any b.

Table 3
Values of b, given $g(\pi_o, n)=n^b\pi_o$, that maximize the Afriat efficiency for circular matching

Subject	$\mathcal{A}_i$ Maximizing b	Afriat efficiency [a] $\mathcal{A}_i$	Number of GARP violations
38	0	**0.988**	1
42	0	0.856	136
58	0	**0.994**	1
80	0	0.700	15
90	0.075	0.914	24
43	0.175	**0.975**	3
32	0.250	0.889	119
35	0.275	**0.970**	16
22	0.300	**0.975**	16
46	0.325	0.893	20
89	0.400	0.680	92
96	0.400	0.904	28
85	0.425	0.889	36
39	0.625	0.733	115
88	0.775	0.942	28
31	0.900	**0.998**	1
83	1	**0.990**	7
98	1	0.556	253
Subjects with $\mathcal{A}_i = 1$:		42	
Subjects with $\mathcal{A}_i<1$:		18	
Subjects with $\mathcal{A}_i<0.95$:		11	

[a] Bold indicates $A_i \geq 0.95$.

To answer this, we calculated Euclidean distance, in tokens, to one of the pure choices for each of 41 b's between 0 and 1. Table 5 shows that 15 subject fit one of these exactly, 11 of which were perfectly selfish. Of the three perfect substitutes, one exists for $b=0$, and the other two are for b of 0.65 and 0.68. If we relax our criterion to so that the Euclidean distance is 6 or less, we include 3 more (nearly) selfish players. If we relax the definition to a distance of 10, we find four more Leontief players. These 4 have optimal b's of 0.02, 0.08, 0.10, and 0.28. Looking more closely at the data, these are subjects who are best described as having Leontief preferences $u=\min\{\pi_i, \pi_o\}$ (i.e. $b=0$), except for one or two "mistakes" that deviate from this.

6.2. Parametric estimation

Following Andreoni and Miller (2002), we will estimate a CES utility function.[20] One reason for doing this is that the three functions listed above are all CES functions, hence we can think of all subjects as having preferences from the same family of utility functions. A second reason is that the CES is reasonably flexible and has very few parameters to estimate. With 24 observations we found estimation of functions with more parameters, such as the translog, to be untenable.
Consider the utility function

$$u(\pi_i, g) = [a\pi_i^\rho + (1-a)g^\rho]^{1/\rho}, \qquad (2)$$

Table 4
Values of b, given $g(\pi_o, n)=n^b\pi_o$, that maximize the Afriat efficiency for reciprocal matching

Subject	$\mathcal{A}_i$ Maximizing b	Afriat efficiency[a] $\mathcal{A}_i$	Number of GARP violations
3	0	**0.963**	3
5	0	0.833	78
8	0	0.750	120
10	0	0.875	59
12	0	0.833	28
15	0	**0.988**	1
117	0	0.933	6
102	0.125	**0.975**	37
6	0.400	0.960	5
112	0.400	**0.992**	7
62	0.650	**0.996**	3
108	0.675	0.875	11
119	0.675	**0.996**	3
101	0.850	0.797	210
75	0.900	**0.981**	22
106	0.900	**0.950**	66
73	1	0.783	21
78	1	**0.951**	45
Subjects with $\mathcal{A}_i = 1$:		42	
Subjects with $\mathcal{A}_i < 1$:		18	
Subjects with $\mathcal{A}_i < 0.95$:		9	

[a] Bold indicates $A_i \geq 0.95$.

where $-\infty \leq \rho \leq 1$ and $\sigma = 1/(\rho-1) < 0$ is the elasticity of substitution, and $0 \leq a \leq 1$. Assuming $g = n^b\pi_o$ and choosing π_i as the numeraire, then $p_g = p/n^b$. Maximizing Eq. (2) subject to the budget $\pi_i + p_g g = m$ yields the demand function

$$g^* = \frac{p_g^{r-1}}{p_g^r + A} m \qquad (3)$$

where $r = -\rho/(1-\rho)$ and $A = [a/(1-a)]^{1/(1-\rho)}$.

Because $0 \leq g^* \leq m/p_g$, we estimate Eq. (3) with a two-limit Tobit, censored at both ends of the budget constraint. The techniques are the same as those used by Andreoni and Miller (2002), and the same technical appendix applies.

We estimated utility functions for 109 of the subjects, excluding those identified as perfectly selfish in panel a of Table 5. We report the results of all 109 regressions in an appendix (available from the author). Table 6 gives an overview of the results.

Table 6 divides the estimates based on parameters b, a, and σ. A σ closer to 0 means the indifference curves are increasingly L-shaped, they become increasingly flat as σ approaches $-\infty$, and $\sigma = -1$ is Cobb–Douglas utility.

This table reveals a vast heterogeneity across subjects. Of the 109 utility functions estimated, 25% are best described with $b=0$ and fully congested altruism. Another 17% have $b=1$ and purely uncongested altruism. The remaining 58% have preferences best described as partially congested altruism, $0 < b < 1$.

We can also characterize our sample as favoring themselves more than others. That is, there are very few subject with a's that favor the other subject. Looking at the individual regressions, 12 have estimated $a < 0.5$, and only four of these estimates are significantly lower than 0.5.
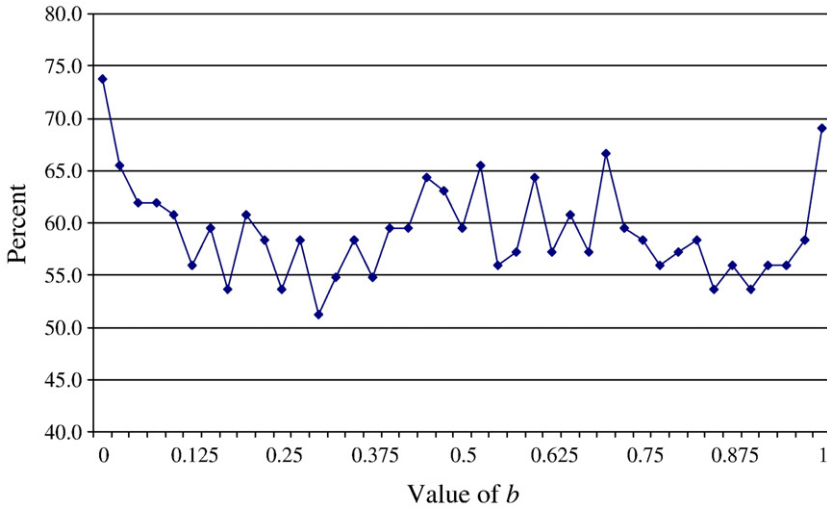
Fig. 4. Percent of subjects for whom b is admissible, for subjects with $\mathcal{A}_i = 1$.

Finally, subjects also differ greatly on the curvature of their preferences, ranging from very sharp indifference curves with $-0.5 < \sigma$ (23% of subjects) to very flat indifference curves with $\sigma < -4$ (24% of subjects).

On the one hand, the results of this section are gratifying—they illustrate what we know from personal experience, that is, there is vast heterogeneity on altruism. Such differences in tastes needs to be kept in mind, for instance, when designing institutions for giving or calculating the incidence of policy. On the other hand, the heterogeneity makes it difficult to characterize the

Table 5
Numbers of subjects as pure types with $u(\pi_i, g(\pi_o, n))$ and $g(\pi_o, n) = n^b \pi_o$

	$b=0$	$0<b<1$	$b=1$	Total	
a. Euclidean distance = 0					
Selfish [a]	–	–	–	11	
Substitutes	1	2	0	3	
Leontief	1	0	0	1	
Total				15	(12.5%)
b. Euclidean distance ≤ 6					
Selfish	–	–	–	14	
Substitutes	1	2	0	3	
Leontief	1	0	0	1	
Total				18	(15%)
c. Euclidean Distance ≤ 10					
Selfish	–	–	–	16	
Substitutes	1	2	0	3	
Leontief	1	4	0	5	
Total				24	(20%)

[a] Any $b \in [0, 1]$ will rationalize selfish choices.

Table 6
Distribution of utility parameters for $u=[a\pi_i^\rho+(1-a)(n^b\pi_o)^\rho]^{1/\rho}$, where $\sigma=1/(\rho-1)$ is the elasticity of substitution

a. $b=0$

	$0<a<0.4$	$0.4\leq a<0.6$	$0.6\leq a<0.8$	$0.8\leq a<1$	Total
$-0.5<\sigma$	0	0	3	7	
$-1.5\leq\sigma<-0.5$	0	0	9	2	
$-4\leq\sigma<-1.5$	0	4	1	0	
$\sigma<-4$	0	1	0	0	27

b. $0<b<.5$

	$0<a<0.4$	$0.4\leq a<0.6$	$0.6\leq a<0.8$	$0.8\leq a<1$	
$-0.5<\sigma$	4	2	0	2	
$-1.5\leq\sigma<-0.5$	2	1	2	1	
$-4\leq\sigma<-1.5$	1	6	3	0	
$\sigma<-4$	1	0	3	5	33

c. $0.5\leq b<1$

	$0<a<0.4$	$0.4\leq a<0.6$	$0.6\leq a<0.8$	$0.8\leq a<1$	
$-0.5<\sigma$	2	1	2	0	
$-1.5\leq\sigma<-0.5$	0	2	1	0	
$-4\leq\sigma<-1.5$	0	4	4	3	
$\sigma<-4$	0	2	10	0	31

d. $b=1$

	$0<a<0.4$	$0.4\leq a<0.6$	$0.6\leq a<0.8$	$0.8\leq a<1$	
$-0.5<\sigma$	0	0	0	2	
$-1.5\leq\sigma<-0.5$	0	1	4	3	
$-4\leq\sigma<-1.5$	0	2	2	0	
$\sigma<-4$	0	2	1	1	18
Total	10	28	45	26	109

central tendency of the society as a whole from individual analysis. To say something about the "average" person, we must estimate a representative utility function.

7. A representative altruist

If we were to characterize the population in aggregate, what b best describes the "representative" altruist? To answer this we pooled all 120 subjects and each of their 24 decisions into a single regression to estimate an aggregate CES utility function. The results are reported in Table 7. All parameters are precisely estimated—b, a and σ are all significantly different from both 0 and 1 beyond the $\alpha\leq.001$ level. Starting with σ, we see that the elasticity of substitution is -1.70, indicating indifference curves that are flatter than Cobb Douglas, but still quite curved. The parameter $a=0.70$ indicates that subjects on average give their own payoff about twice the weight they give to the aggregate of the others. Finally, $b=0.68$ indicates that the representative altruist experiences significant congestion of altruism.

How do we interpret $b=0.68$? Consider Al and his identical twin brother Bob, with the same income and utility function as estimated above. Al is giving up a dollar, and each of two others is gaining a dollar. Bob, by contrast, is giving up a dollar and a single other person is getting $1.59. Al and Bob have the same utility, that is, $2^b=1.59$.

Table 7
Estimates of a CES utility $u=[a\pi_i^\rho+(1-a)(n^b\pi_o)^\rho]^{1/\rho}$ for a representative altruist

Coefficient (standard error)	
b	0.676
	(0.088)
a	0.705
	(0.010)
$\sigma=1/(\rho-1)$	−1.704
	(0.066)
SE	0.432
N	2880

Here is another way to look at it. Suppose Cindy is spending a dollar which creates one dollar benefit for each of 4 others. The value of this gift is $4^b=2.55$. Suppose the benefit expands to one dollar to each of 5 others. This value of this gift is $5^b=2.97$. That is, adding the extra one dollar of social benefits only increased the private value of the gift to the giver by 0.42. Repeating this exercise moving from 9 to 10, the private benefit of the extra dollar of social benefit falls to 0.32. Although it extrapolates well outside of our sample, the estimated b implies that the 101st dollar of social benefit is only worth 0.15 to the giver, and the 1001st dollar is only worth 0.07 to the giver. For extremely large populations—such as when giving to a pure public good with many consumers—people are essentially indifferent to the number of recipients.

With the utility function estimated above, we can ask how contributions to a public good change with the size of the group. Fig. 5 shows the demand for the average payoffs of others, π_o, under the assumption that giving is to a pure public good, and under the assumption that giving is to a pure private good as in the example of pooling of tips described in Sections 1 and 2 above. Since the levels of giving depend on m, we normalized gifts to equal one at $n=1$ (note, demands are identical for both public and private goods at $n=1$). Fig. 5 shows that for both public and private goods, the demand for average payoffs of others is declining. For public goods, the decline is coming because the substitution effect is dominating the income effect, that is, because one can create the same social value for a
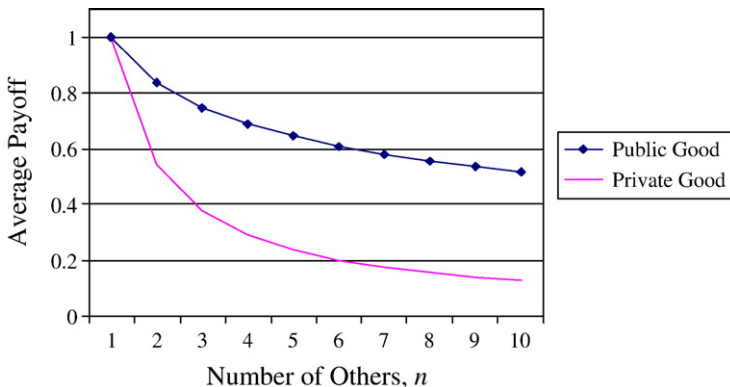

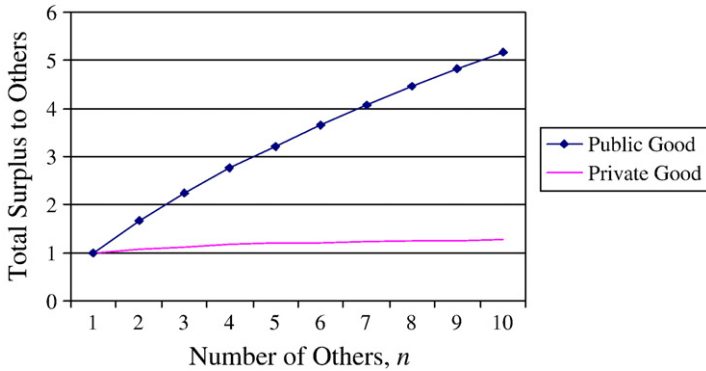Fig. 5. Demand for average payoff to others.

Fig. 6. Demand for total social surplus of others.

smaller gift people are giving less.[21] For private goods this effect is compounded by the fact that a higher n raises the price of average payoffs.

Perhaps the more important question is what happens to the total social surplus as n grows. Do gifts fall so much that the surplus is not increased? Fig. 6 shows that for both public and private goods the total social surplus is rising with n. For public goods, we definitely see the role of congestion of altruism. If altruism were not congested, the slope of the top line in Fig. 6 would be one—adding additional consumers would not diminish the gifts made. Because there is congestion, this slope is less than one.

The line in Fig. 6 for private goods also tells an interesting story. Because congestion is not complete, total giving actually rises here—spreading out the food to more food bank recipients lowers the total utility of the gift. As a result, the marginal utility of a any gift rises, so giving must increase. By the same token, tip pooling should, all else equal, raise total tips.

One can also notice that total spending on giving for public goods (the top line in Fig. 5) is decreasing for n. This would suggest that if you asked for willingness to give to a "national" public good, and for a "local" public good that is nested within the national good, then giving should be lower at the national level.

8. Conclusion

This paper explores preferences for giving gifts to groups. How does the value of the gift to the giver depend on the number of recipients? Is altruism in the hearts of givers a congestible good?

We address this question by assuming that givers care about an aggregate benefit of the gift, where the aggregator can reflect congestion, that is $g = n^b \pi_o$ and $0 \leq b \leq 1$. Only if $b = 1$ is altruism uncongested.

Using experiments that vary the number of recipients and the price of giving, we found that 70% of subjects made all 24 choices without a single violation of GARP, and fewer than 17% had violations with Afriat indices below 0.95. We are then able to identify individual preferences for giving, and to find the b for each subject that best fits the choices.

---

[21] The line for public goods in Fig. 5 also tells us how a giver's spending on gifts changes with n. While one needs to add a grain of salt when generalizing outside of the experimental parameters and context, these preferences would also lead us to say that for large groups, giving to public goods will be largely unaffected by n. When n goes from 100 to 101, for instance, spending on giving falls by only 0.4%, and from 1000 to 1001 it falls by 0.04%.

We found that only 17% of our subjects have $b=1$ and thus have uncongested altruism. At the opposite extreme, 25% are fully congested with $b=0$. For the remaining 58% we estimate b's between 0 and 1. The single b that best describes the central tendency in the data is $b=0.68$. That is, as groups grow, altruism of the givers is congested and the value of the gift to the giver does not grow proportionately with the social value of the public good.

The most direct and immediate consequence of these results is for experiments. In situations where we think altruism, fairness, and inequality aversion play a role, the preferences for allocations should depend on the number of people playing the game, even if they are treated symmetrically.

To the extent we can generalize these findings outside the lab, they have other interesting implications. First, as groups grow, average giving will decline. This is not, however, for the usual strategic reasons of increased free riding, but follows simply from the fact that congested altruism changes the value of the social surplus to the giver. This means that, for instance, when presented the option of giving to the economics department or to the university as a whole, the donor may give more to the economics department. This may be one reason why university fund-raisers often direct givers to individual departments, research consortia, or research labs when soliciting major gifts.

Second, when donations buy private goods for a set of people, an individual's total giving rises slightly with n. This means that waiters who pool tips shouldn't worry about reducing altruism so much that total tips will fall. A new niece or nephew means that the new child's siblings will get smaller birthday presents from their uncle, but the total gifts to the family will not change much.

Finally, for purposes of tax policy, this suggests that giving to charities that have many consumers will not be greatly affected by numbers of givers. The congestion of altruism means that the marginal utility of giving will decline as the number of consumers increases, eventually making giving insensitive to n, and tax subsides will not have a differential incidence across charities of different sizes.

This laboratory study is, obviously, only the first step in understanding how people give gifts to groups. The most interesting applications of these findings are to real world giving, which makes the appeal to external validity all the more important. Future field studies will be needed to show whether the results found for the small numbers and special circumstances of the lab can be extrapolated to the real world.

References

Afriat, S., 1972. Efficiency estimates of productions functions. International Economic Review 13 (3), 568–598 (October).

Andreoni, James, 1988a. Privately provided public goods in a large economy: the limits of altruism. Journal of Public Economics 35, 57–73 (February).

Andreoni, James, 1988b. Why free ride? Strategies and learning in public goods experiments. Journal of Public Economics 37 (3), 291–304 (December).

Andreoni, James, 1989. Giving with impure altruism: applications to charity and Ricardian equivalence. Journal of Political Economy 97, 1447–1458.

Andreoni, James, 1990. Impure altruism and donations to public goods: a theory of warm-glow giving. Economic Journal 100, 464–477.

Andreoni, James, 1995. Cooperation in public goods experiments: kindness or confusion? American Economic Review 85 (4), 891–904 (September).

Andreoni, James, Harbaugh, William, 2006. Power indices for revealed preference tests. Working Paper, UCSD.

Andreoni, James, Miller, John H., 2002. Giving according to GARP: an experimental test of the consistency of preferences for altruism. Econometrica 70 (2), 737–753 (March).

Andreoni, James, Castillo, Marco, Petrie, Ragan, 2003. What do bargainers' Preferences look like? Exploring a convex ultimatum game. American Economic Review 93 (3), 893–902 (June).

Bagnoli, Mark, McKee, Michael, 1991. Voluntary contribution games: efficient private provision of public goods. Economic Inquiry 29 (2), 351–366 (April).

Bergstrom, Theodore C., Goodman, Robert P., 1973. Private demands for public goods. American Economic Review 63 (3), 280–296 (June).

Bergstrom, Theodore C., Blume, Lawrence, Varian, Hal R., 1986. On the private provision of public goods. Journal of Public Economics 29, 25–49.

Bronars, Stephen G., 1987. The power of nonparametric tests of preference maximization. Econometrica 55 (3), 693–698 (May).

Cox, James, Friedman, Daniel, Gjerstad, Steven, 2007. A tractable model of reciprocity and fairness. Games and Economic Behavior 59 (1), 17–45 (April).

Croson, Rachel, Marks, Melanie, 1998. Alternative rebate rules in the provision of a threshold public good: an experimental investigation. Journal of Public Economics 67 (2), 195–220 (February).

Fisman, Raymond, Shachar Kariv, and Daniel Markovits, in press. Individual Preferences for Giving. Working paper, American Economic Review.

Goeree, Jacob K., Holt, Charles A., Laury, Susan K., 2002. Private costs and public benefits: unraveling the effects of altruism and noisy behavior. Journal of Public Economics 83 (2), 255–276 (February).

Hey, John D., Lee, Jinkwon, 2005. Do subjects separate (or are they sophisticated)? Experimental Economics 8 (3), 233–265 (September).

Isaac, Mark R., Walker, James M., 1988. Group size effects in public goods provision: the voluntary contributions mechanism. Quarterly Journal of Economics 103 (1), 179–199 (February).

Isaac, Mark R., Walker, James M., Williams, Arlington W., 1994. Group size and the voluntary provision of public goods: experimental evidence utilizing large groups. Journal of Public Economics 54 (1), 1–36 (May).

Kahneman, Daniel, Knetsch, Jack L., 1992. Valuing public goods: the purchase of moral satisfaction. Journal of Environmental Economics and Management 22 (1), 57–50 (January).

Varian, Hal R., 1982. The nonparametric approach to demand analysis. Econometrica 50 (4), 945–973 (July).