Individual Preferences for Giving*

Raymond Fisman[†] Columbia University
Shachar Kariv[‡] UC Berkeley
Daniel Markovits[§] Yale University

March 28, 2007

Abstract

We utilize graphical representations of Dictator Games which generate rich individual-level data. Our baseline experiment employs budget sets over feasible payoff-pairs. We test these data for consistency with utility maximization, and we recover the underlying preferences

*The results reported here were previously distributed in three different papers titled "Individual Preferences for Giving," "Distinguishing Social Preferences from Preferences for Altruism" and "Pareto Damaging Behaviors." This research was supported by the Experimental Social Science Laboratory (X-Lab) at the University of California, Berkeley. We thank Jim Andreoni and Gary Charness for detailed comments and suggestions. We are grateful to Bruce Ackerman, Ian Ayres, Colin Camerer, Ken Chay, Stefano DellaVigna, Liran Einav, Douglas Gale, Steve Goldman, Alvin Klevorick, Botond Koszegi, John List, Tom Palfrey, Ben Polak, Jim Powell, Matthew Rabin, Al Roth, Ariel Rubinstein, Alan Schwartz, Andrew Schotter, Chris Shannon, Hal Varian, and Bill Zame for helpful discussions. This paper has also benefited from suggestions by the participants of seminars at Caltech, Institute for Advanced Studies, NYU, Yale SOM, UC Berkeley, UC San Diego, and the AEA 2006 annual meeting in Boston. Syngjoo Choi and Benjamin Schneer provided excellent research assistance. In addition, we would like to thank Brenda Naputi and Lawrence Sweet from the X-Lab for their valuable assistance, and Roi Zemmer for writing the experimental computer program. For financial support, Fisman thanks the Columbia University Graduate School of Business; Kariv acknowledges UC Berkeley COR grant; Markovits thanks Yale Law School and Deans Anthony Kronman and Harold Koh. Kariv is grateful to the hospitality of the Institute for Advances Studies School of Social Science.

[†]Graduate School of Business, Columbia University, Uris 823, New York, NY 10027 (E-mail: rf250@columbia.edu, URL: http://www-1.gsb.columbia.edu/faculty/rfisman/).

[‡]Department of Economics, University of California, Berkeley, 549 Evans Hall # 3880, Berkeley, CA 94720 (E-mail: kariv@berkeley.edu, URL: http://socrates.berkeley.edu/~kariv/).

[§]Yale Law School, P.O. Box 208215, New Haven, CT 06520. (E-mail: daniel.markovits@yale.edu, URL: http://www.law.yale.edu/outside/html/faculty/ntuser93/profile.htm)

for giving (tradeoffs between own payoffs and the payoffs of others). Two further experiments augment the analysis. An extensive elaboration employs three-person budget sets to distinguish preferences for giving from social preferences (tradeoffs between the payoffs of others). And an intensive elaboration employs step-shaped sets to distinguish between behaviors that are compatible with well-behaved preferences and those that are compatible only with not well-behaved cases.

JEL Classification Numbers: C79, C91, D64.

We study *individual* preferences for giving. Our experiments employ a graphical interface that allows subjects to see geometric representations of choice sets on a computer screen and to make decisions through a simple point-and-click response. The rich data generated by this design facilitate statistical analysis at the level of the individual subject with no need to pool data or assume homogeneity across subjects.

Our first experiment employs a modified dictator game, developed by James Andreoni and John H. Miller (2002), that varies the endowments and the prices of giving, so that a person $self$ faces a menu of *budget sets* over his own payoff and the payoff of *other*. We begin our analysis by testing for consistency with utility maximization using revealed preference axioms. The broad range of budget sets that our experiment employs leads to high power tests of consistency. We find that most subjects exhibit behavior that appears to be *almost* optimizing so that the violations are minor enough to ignore for the purpose of constructing appropriate utility functions. We then move to estimate constant elasticity of substitution (CES) demand functions for giving at the individual level. The parameter estimates vary dramatically across subjects, implying that preferences for giving are very heterogeneous, ranging from perfect substitutes to Leontief. However, we do find that subjects display a pronounced (although far from monolithic) emphasis on increasing aggregate payoffs (the elasticity of substitution between the payoffs to persons $self$ and *other* is smaller than -1) rather than reducing differences in payoffs (the elasticity of substitution is greater than -1).

While *preferences for giving* govern the tradeoffs that $self$ makes between his payoffs and the payoffs of $others$ (all persons except $self$), *social preferences* govern the tradeoffs $self$ makes among the payoffs to $others$. Although these two types of *distributional preferences* often operate together, as when we decide both how much to give to charity and how to allocate our donations across causes, they remain conceptually distinct.[1] Certainly there is no *a priori* reason to insist that preferences for giving and social

[1] The terms "distributional preferences" and "social preferences" have been used inter-

preferences have the same (or even a similar) form. In order to distinguish preferences for giving from social preferences and to compare these two classes of distributional preferences, we use *three-person* dictator games that vary the prices of giving, so that $self$ faces a menu of budget sets over his own payoff and the payoffs of two $others$.

With the three-person data, we extend the conclusions of the two-person experiment that preferences for giving are highly heterogeneous. We also compare preferences for giving and social preferences and find (although with a few interesting exceptions) that subjects employ a unified approach to efficiency-equity tradeoffs across both realms. Thus, although there is considerable heterogeneity in preferences for giving and social preferences *across* subjects, there is a strong positive association between preferences for giving and social preferences *within* subjects.

According to Sidney N. Afriat's (1967) theorem, our analysis over linear budget sets necessarily treats preferences as *well-behaved* (continuous, increasing and concave), since price and quantity data do not allow us to distinguish between decisions that are compatible with a well-behaved utility function and those that are compatible only with less well-behaved cases. This is crucial, since several prominent theories of distributional preferences, namely difference aversion models, posit a utility function that is not well-behaved, and the aim of our analysis is to identify "correct" individual-level utility functions. We therefore turn to a version of our experimental design in which each subject faces a menu of *step-shaped sets* (as illustrated in Figure 1 below) representing the feasible monetary payoffs to person $self$ and one $other$.

The step-shaped set enables us to differentiate among various prototypical preferences – competitive, self-interested, *lexself* (lexicographic for $self$ over *other*), difference averse, and social welfare. Most importantly, the non-convexity and sharp nonlinearity of the step-shaped constraint means that $self$ always faces choices with an extreme price of giving. In this context, either $self$ or *other* must be made monetarily strictly worse off in order to create greater equality or greater inequality. With the step-shaped data, we extend the conclusion of the linear budget set experiments that preferences for giving vary widely across subjects, ranging from competitive to selfish to lexself (the single commonest form) to difference averse to social welfare. Moreover, some of our difference averse subjects (and some other subjects also) systematically display behaviors that are consistent only

changeably in the literature. Nevertheless, the distinctions that we draw are straightforward and (as our analysis reveals) capture important differences.

with not well-behaved preferences. Finally, several of our subjects display a balance of selfishness and difference aversion that leads them to make allocations that cannot be accommodated by the canonical models of distributional preferences encapsulated in Gary Charness and Matthew Rabin (2002).

Our paper contributes to a large and growing body of work on distributional preferences, including George E. Loewenstein, Leigh Thompson and Max H. Bazerman (1989), Gary E. Bolton (1991), Rabin (1993), David K. Levine (1998), Ernst Fehr and Klaus M. Schmidt (1999), Bolton and Axel Ockenfels (1998, 2000), Charness and Rabin (2002), and Andreoni and Miller (2002) among others (Colin F. Camerer (2003) provides a comprehensive discussion). First, we extend the analysis Andreoni and Miller (2002) by collecting richer data about preferences for giving. Second, we present an *extensive* elaboration that uses three-person budget sets to distinguish preferences for giving from social preferences. Third, we present an *intensive* elaboration that employs step-shaped sets to provide further tests of the structure of preferences for giving.[2]

The rest of the paper is organized as follows. Section 1 describes the experimental design and procedures. Sections 2 and 3 provide the results from the budget set experiments and Section 4 from the step-shaped experiment. Section 5 unifies the results and contains some concluding remarks. All individual-level estimates and technical digressions are relegated to appendices.

1 Design and Procedures

In this section, we define a number of concepts and terms that will be used throughout the paper and describe the theory on which the experimental design is based as well as the design itself.

1.1 Two-person budget sets

We denote persons $self$ and *other* by s and o, respectively, and the associated monetary payoffs by π_s and π_o. The set of feasible payoff pairs $\pi = (\pi_s, \pi_o)$ may take many forms. Yet in a typical dictator experiment,

[2] Syngjoo Choi, Ray Fisman, Douglas M. Gale, and Shachar Kariv (2007) employ a similar experimental methodology to study decisions under uncertainty. While the papers share a similar experimental methodology that allows for the collection of many observations per subject, they address very different questions and produce very different behaviors.

subject $self$ divides his *endowment* m between $self$ and an *anonymous other* such that $\pi_s + \pi_o = m$. This framework restricts the set of feasible payoff pairs to the line with a slope of -1, so that the problem faced by $self$ is simply allocating a fixed total income between $self$ and $other$. The simplest and perhaps most important generalization of the dictator game, developed by Andreoni and Miller (2002), maintains the assumption of linearity but allows an endowment to be spent on π_s and π_o at fixed *price* levels p_s and p_o such that $p_s\pi_s + p_o\pi_o \leq m$. This configuration creates *budget sets* over π_s and π_o where p_o/p_s is the *relative price of giving*.

Initially, we wish to examine whether the observed individual-level data could have been generated by a subject maximizing a utility function $U_s = u_s(\pi_s, \pi_o)$ that captures the possibility of giving. If a utility function $u_s(\pi_s, \pi_o)$ that the choices maximize exists, then the techniques of demand analysis may be brought to bear on modeling and predicting behavior governed by these preferences. The crucial test for this is provided by the Generalized Axiom of Revealed Preference (GARP) which requires that if π is revealed preferred to π' then π' is not *strictly* directly revealed preferred to π. GARP is tied to utility representation through the following theorem, which was first proved by Afriat (1967). This statement of the theorem follows Hal R. Varian (1982):

Afriat's Theorem The following conditions are equivalent: (i) The data satisfy GARP. (ii) There exists a non-satiated utility function that rationalizes the data. (iii) There exists a continuous, increasing, concave, non-satiated utility function that rationalizes the data.

1.2 Three-person budget sets

We next investigate choices made by $self$ that have consequences for her own payoff and the payoffs of two anonymous $others$. This is of particular interest insofar as it facilitates the analysis of the two types of *distributional preferences* – *preferences for giving* ($self$ versus $others$) and *social preferences* (*other* versus *other*). With a slight abuse of notation, we denote $others$ by $o = \{A, B\}$ and the associated monetary payoffs and corresponding prices by $\pi_o = (\pi_A, \pi_B)$ and $p_o = (p_A, p_B)$. This configuration creates budget sets over π_s and $\pi_o = (\pi_A, \pi_B)$.

A common assumption used in demand analysis that allows for a clear demarcation between preferences for giving and social preferences is *independence* which entails that if π_o is preferred to π'_o for some π_s, then π_o is preferred to π'_o for all π_s. That is, the preferences of $self$ over the payoffs of *others* are independent of her *self*-interestedness. If this independence property is satisfied, then the utility function $u_s(\pi_s, \pi_o)$ is (weakly) *separable* in the sense that we can find a *subutility* function $w_s(\pi_A, \pi_B)$ and a *macro* function $v_s(\pi_s, w_s(\pi_o))$ with v_s strictly increasing in w_s such that $U_s \equiv v_s(\pi_s, w_s(\pi_o))$. This formulation makes it possible to represent distributional preferences in a particularly convenient manner, because the macro utility function $v_s(\pi_s, w_s(\pi_o))$ represents preferences for giving, whereas the subutility function $w_s(\pi_A, \pi_B)$ represents social preferences.[3] Moreover, separability makes convenient (if restrictive) assumptions on the form of the utility function, which yield empirically testable restrictions on the relationship between preferences for giving and social preferences.

1.3 The step-shaped set

The equivalence of (*i*) and (*ii*) in Afriat's theorem establishes GARP as a direct test for whether the data from our budget set experiments may be rationalized by a utility function, and the equivalence of (*ii*) and (*iii*) tells us that when a rationalizing utility function exists, it may be chosen to be *well-behaved* (continuous, increasing and concave). This last connection entails that when a rationalizing utility function exists, price and quantity data do not allow us to reject the hypothesis that it is well-behaved. The intuitive reason for this is that choices subject to linear budget constraints will never be made at points where the underlying utility function is not well-behaved. Hence, satisfying GARP entails only that choices are consistent with the utility maximization model, whereas the further implication of consistency with a well-behaved utility function is a consequence of the specification of the linear budget constraint.

Given these limitations, we analyze preferences for giving more intensively by studying decisions over *step-shaped sets*. This enables us to distinguish effectively between choices that are compatible with a well-behaved utility function and those that are compatible only with less well-behaved cases. Figure 1 illustrates the step-shaped set in our experiment. In this case, there are only two *socially* optimal allocations: $\pi^s = (\pi^s_s, \pi^s_o)$ maximizes the payoff for *self*; and $\pi^o = (\pi^o_s, \pi^o_o)$ maximizes the payoff for *other*. It should be noted that π^s and π^o cannot be ranked. Thus, the step-shaped set can also be interpreted as presenting subjects with monetarily incompa-

[3] Edi Karni and Zvi Safra (2000) introduce an axiomatic model of choice among random social allocation procedures. Their utility representation is also decomposed in a similar way, and they also provide conditions under which the representation is *additively* separable.

rable *binary* choices like those commonly employed in experiments of distributional preferences, with the added possibility of *free-disposal*. Accordingly, while the step-shaped set follows prior literature in using binary dictator games, it does not "force" subjects into discrete choices and thus permits reducing or increasing differences in payoffs.

[Figure 1 here]

Certain choices within a step-shaped constraint may readily be associated with various prototypical distributional preferences. To aid us in developing these associations, we first define an allocation as $self$- (*other-*) *damaging* if and only if $self$- (*other-*) *monetary* improvements can be made. Figure 1 also depicts the subsets of the step-shaped constraint associated with each type of damaging behavior. The horizontal subsets

$$\Pi^1 = \{\pi : \pi_s = \pi_s^s, 0 < \pi_o < \pi_o^s) \text{ and } \Pi^3 = \{\pi : \pi_s = \pi_s^o, \pi_o^s < \pi_o < \pi_o^o)$$

involve *other*-damaging behavior (that disposes payoffs of *other*), whereas the vertical subsets

$$\Pi^2 = \{\pi : \pi_o = \pi_o^s, \pi_s^o < \pi_s < \pi_s^s) \text{ and } \Pi^4 = \{\pi : \pi_o = \pi_o^o, 0 < \pi_s < \pi_s^o)$$

involve $self$-damaging behavior (that disposes payoffs of $self$).

We further distinguish *inequality-decreasing* from *inequality-increasing* $self$- and *other*-damaging behavior. Whether $self$- or *other*-damaging behavior is inequality increasing or decreasing depends on π^s and π^o, and on π^e, which is the *unique* equal $\pi_s^e = \pi_o^e$ allocation on the step-shaped constraint. More precisely, a $self$- or *other*-damaging allocation π is inequality-decreasing if $|\pi, \pi^e| < |\pi^i, \pi^e|$ where $\pi^i > \pi$ for some person $i = o, s$ and inequality-increasing otherwise.[4] That is, allocation π is inequality-decreasing (increasing) if it is closer to (further from) π^e relative to either π^s or π^o. Indeed, in contrast to choices made on linear budget sets, reducing or increasing differences in payoffs involves $self$- or *other*-damaging behavior.

By separating decisions that damage $self$ and that damage *other* and distinguishing between decisions that are inequality-increasing and inequality-decreasing, we can differentiate among various prototypical distributional

[4] Notice that $\pi^d = (\pi_s^o, \pi_o^s)$ is the only allocation on the step-shaped constraint that involves both $self$- and *other*-damaging behavior. We shall say that π^d is inequality-decreasing if $|\pi^d, \pi^e| < |\pi^i, \pi^e|$ for all $i = o, s$ and apply an analogous characterization of the circumstance in which π^d is inequality-increasing.

preferences: (i) *competitive* preferences, where utility increases in the difference $\pi_s - \pi_o$, are consistent only with the competitive allocation $\pi^c = (\pi^s_s, 0)$; (ii) *narrow self-interest* or *selfish* preferences, where utility depends only on π_s, are consistent with any allocation π where $\pi_s = \pi^s_s$; (iii) *difference aversion* preferences, where utility increases in π_s and decreases in the difference $\pi_s - \pi_o$, are generally consistent with the allocations π^s and π^e if $\pi^e_s = \pi^o_s$; (iv) *social welfare* preferences, where utility increases in both π_s and π_o, are consistent only with π^s and π^o; (v) *lexself* preferences, where utility is lexicographic for π_s over π_o, are consistent with π^s only. These definitions are inspired by the model of Charness and Rabin (2002), which embeds several canonical models of distributional preferences as special cases. We refer the interested reader to Appendix I for more details.

Notice that within the linear budget set, competitive, selfish and lexself preferences are all consistent with only the "selfish" allocation $\pi = (m/p_s, 0)$, so that tests of behavior that employ only such sets cannot distinguish among these preferences completely. Hence, the step-shaped set differs from the linear budget set in two ways. First, it does not allow for incremental efficient sacrifices that decrease inequality and therefore provides a challenging test of difference aversion. Second, it also permits distributional preferences that increase inequality such as selfishness and competitiveness. Thus, the step-shaped sets "span" a range of prototypical preferences, enabling a more refined classification of behaviors than was possible based solely on linear budget sets.

1.4 Experimental procedures

The subjects in the experiments were recruited from all undergraduate classes and staff at UC Berkeley. The procedures used in the three versions of the experiment were identical, with the exception that the sets of feasible monetary payoff choices were different. The treatment was held constant throughout a given experimental session, and each subject participated in only one session. Each session consisted of 50 independent decision-problems. In each decision problem, each subject was asked to allocate tokens between himself and an anonymous subject(s), where the anonymous subject(s) was chosen at random from the group of subjects in the experiment. Each choice involved choosing a point on a two- (three-) dimensional graph representing the set of possible payoff allocations $\pi = (\pi_s, \pi_o)$.

In the two- and three-person budget set versions (subjects ID 1-76 and ID 135-199, respectively), each decision problem started by having the computer select a budget set randomly from the set of budget sets that intersect with at least one of the axes at 50 or more tokens, but with no intercept exceeding 100 tokens. In the step-shaped version of the experiment (subjects ID 77-134), each decision problem started by having the computer select a set randomly from the set $\{\pi : \pi \leq \pi^s\} \cup \{\pi : \pi \leq \pi^o\}$ where $10 \leq \pi^s, \pi^o \leq 100$ and $\pi^o_s < \pi^s_s$ and $\pi^s_o < \pi^o_o$. The sets selected for each subject in different decision problems were independent of each other and of the sets selected for any of the other subjects in their decision problems. In the two-person versions of the experiment, choices were not restricted to allocations on the constraints so that subjects could freely dispose of payoffs. In the three-person version, choices were restricted to allocations on the budget constraint, which made the computer program easier to use. The computer program dialog window is shown in the experimental instructions that are reproduced in Appendix II.

At the end of the experiment, the experimental program randomly selected one decision round to carry out for the purpose of generating payoffs. In the two-person versions, each subject received the tokens that he held in this round (π_s) and the subject with whom he was matched received the tokens that he passed (π_o). Thus, as in Andreoni and Miller (2002), each subject received two groups of tokens, one based on his own decision to hold tokens and one based on the decision of another random subject to pass tokens. In the three-person version, each subject received the tokens that he held in this round (π_s) and the subjects with whom he was matched received the tokens that he passed (π_A and π_B). Thus, each subject received three groups of tokens, one based on his own decision to hold tokens and two based on the decisions of two other random subjects to pass tokens. The computer program ensured that the same two subjects were not paired twice as $self$-*other* and *other*-$self$.

2 Two-person budget sets

2.1 Data description

We begin with an overview of some basic features of the experimental data. Figure 2 depicts the distribution of the expenditure on tokens given to *other* as a fraction of total expenditure $p_o\pi_o/(p_s\pi_s + p_o\pi_o)$. We present the distribution for *all* allocations as well as the distributions by three price ratio terciles: intermediate prices of around 1 ($0.70 \leq p_o/p_s \leq 1.43$), steep prices ($p_o/p_s > 1.43$) and symmetric flat prices ($p_o/p_s < 0.70$). For the full sample there is a local mode at the midpoint of 0.5 (note that we divide the bottom decile in half because of the very striking decline within this decile).

The number of allocations then decreases as we move to the left, before increasing rapidly to selfish allocations of 0.05 or less of the total expenditure on tokens for *other*, which account for 40.5 percent of all allocations. This masks some heterogeneity by price. For the middle tercile, the pattern is somewhat more pronounced, while for the flat tercile, there is no peak at the midpoint. Not surprisingly, the distribution is generally further to the left for steeper-sloped budgets. The distributions of the tokens given to *other* as a fraction of the sum of the tokens kept and given $\pi_o/(\pi_s+\pi_o)$ show similar patterns, though they are somewhat more skewed to the left.

[Figure 2 here]

Compared with studies of split-the-pie dictator games, the mode at the midpoint is relatively less pronounced and the distribution is much smoother, even for the intermediate tercile allocations. Over all prices, our subjects gave to *other* about 19 percent of the tokens, accounting for 21 percent of total expenditure, which is very similar to typical mean allocations of about 20 percent in the studies reported in Camerer (2003). Hence, although the behaviors of our subjects vary widely at the aggregate level, important features of the experimental data are very similar to the data that come out of previous studies.

The aggregate distributions tell us little about the particular allocations chosen by individual subjects. Of our 76 subjects, 20 (26.3 percent) behaved perfectly selfishly. Only two (2.6 percent) subjects allocated all their tokens to $self$ if $p_s < p_o$ and to *other* if $p_s > p_o$ implying utilitarian preferences, and two (2.6 percent) subjects made nearly equal expenditure on $self$ and *other* indicating Rawlsian preferences.[5] We also find many intermediate cases, but these are difficult to see directly from the data due to the fact that both p and m shift in each new allocation.

2.2 Testing rationality

Before turning to GARP violations, we note initially that half of our subjects have no violations of *budget balancedness* ($p_s\pi_s + p_o\pi_o < m$) even with a narrow one token confidence interval.[6] If we allow for a five token

[5] By comparison, Andreoni and Miller (2002) report that 40 subjects (22.7 percent) behaved perfectly selfishly, 25 subjects (14.2 percent) could fit with utilitarian preferences, and 11 subjects (6.2 percent) were consistent with Rawlsian preferences.

[6] We allow for small mistakes resulting from the slight imprecision of subjects' handling of the mouse. Thus, the subsequent results allow for a narrow confidence interval of one token (for any π and $\pi' \neq \pi$ if $|\pi, \pi'| \leq 1$ then π and π' are treated as the same allocation).

confidence interval, 64 subjects (84.2 percent) have no violations of budget balancedness.[7] We next assess how nearly the data comply with GARP by calculating Afriat's (1972) *Critical Cost Efficiency Index* (CCEI) which measures the amount by which each budget constraint must be adjusted in order to remove all violations of GARP. Hence, the CCEI is bounded between zero and one and can be interpreted as measuring the upper bound of the fraction of his wealth that person $self$ is 'wasting' by making inconsistent choices. The closer the CCEI is to one, the smaller the perturbation of the budget constraints required to remove all violations and thus the closer the data are to satisfying GARP. Appendix III provides details on testing for consistency with GARP and other indices that have been proposed for this purpose by Varian (1991) and Martijn Houtman and J. A. H. Maks (1985).

Next, we generate a benchmark level of consistency with which we may compare our CCEI scores. As in Andreoni and Miller (2002), we use the test designed by Stephen G. Bronars (1987) that employs the choices of a hypothetical subject who randomizes uniformly among all allocations on each budget line as a benchmark. Figure 3 shows the distribution of CCEI scores generated by a sample of 25,000 hypothetical subjects and the actual distribution. It makes plain that the significant majority of our subjects came much nearer to consistency with utility maximization than random choosers and that their CCEI scores were only slightly worse than the score of one of the perfect utility maximizers.[8] We therefore conclude that most subjects exhibit behavior that appears to be *almost* optimizing in the sense that their choices nearly satisfy GARP, so that the violations are minor enough to ignore for the purposes of recovering preferences or constructing appropriate utility functions.

[Figure 3 here]

2.3 Econometric specification

Our subjects' CCEI scores are sufficiently near one to justify treating the data as utility-generated, and Afriat's theorem tells us that the underlying

[7] A few subjects required large confidence intervals to remove *all* budget balancedness violations, but these subjects also have many GARP violations even if the choices that violate budget balancedness are removed.

[8] By comparison, Andreoni and Miller (2002) report that only 18 of their 176 subjects (10.2 percent) violated GARP, and of those only 3 had CCEI scores below the 0.95 threshold. This is as expected, as our subjects were given a larger and richer menu of budget sets, which provides more opportunities to violate GARP.

utility function $u_s(\pi_s, \pi_o)$ that rationalizes the data can be chosen to be well-behaved. Like Andreoni and Miller (2002), we further assume that $u_s(\pi_s, \pi_o)$ is a member of the constant elasticity of substitution (CES) family given by

$$U_s = [\alpha(\pi_s)^\rho + (1-\alpha)(\pi_o)^\rho]^{1/\rho}$$

where α represents the relative weight on the payoff for $self$, ρ represents the curvature of the indifference curves, and $\sigma = 1/(\rho - 1)$ is the (constant) elasticity of substitution. When $\alpha = 1/2$, $U_s \to \pi_s + \pi_o$ (the purely utilitarian case) as $\rho \to 1$, and $U_S \to \min\{\pi_s, \pi_o\}$ (the Rawlsian case) as $\rho \to -\infty$. As $\rho \to 0$, the indifference curves approach those of a Cobb-Douglas function, which implies that the expenditures on tokens kept and given are equal to fractions α and $1-\alpha$ of the endowment m, respectively. Further, if $\rho > 0$ ($\rho < 0$) a fall in the relative price of giving p_o/p_s lowers (raises) the expenditure on tokens given to *other* as a fraction of total expenditure. Thus, any $\rho > 0$ ($\sigma < -1$) indicates distributional preferences weighted towards increasing total payoffs, whereas any $\rho < 0$ ($-1 < \sigma < 0$) indicates distributional preferences weighted towards reducing differences in payoffs.

The CES demand function is given by

$$\pi_s(p_s, p_o, m) = \left[\frac{g}{\left(\frac{p_o}{p_s}\right)^r + g}\right]\frac{m}{p_s}$$

where $r = -\rho/(1-\rho)$ and $g = [\alpha/(1-\alpha)]^{1/(1-\rho)}$. This generates the following individual-level econometric specification for each subject n:

$$\frac{p_{s,n}\pi^t_{s,n}}{m^t_n} = \frac{g_n}{(\frac{p^t_{o,n}}{p^t_{s,n}})^{r_n} + g_n} + \epsilon^t_n$$

where $t = 1, ..., 50$ and ϵ^t_n is assumed to be distributed normally with mean zero and variance σ^2_n. We generate estimates of $\hat{g}_n$ and $\hat{r}_n$ using non-linear tobit maximum likelihood, and use this to infer the values of the underlying CES parameters $\hat{\alpha}_n$ and $\hat{\rho}_n$ (we generate virtually identical parameter values using non-linear least squares). Before proceeding to the estimations, we omit the 11 subjects (26.3 percent) with CCEI scores below 0.80, as the choices of subjects with CCEI scores not sufficiently close to one cannot be utility-generated. We also screen out 20 subjects (14.5 percent) with uniformly selfish allocations (average $p_s\pi_s/m \geq 0.95$) whose preferences are easily identifiable. This leaves a total of 45 subjects (59.2 percent) for whom we need to recover the underlying preferences by estimating the CES model. Appendix IV presents, subject by subject, the results of the estimations.

2.4 Preferences for giving

Of the 45 subjects with consistent, non-selfish preferences, two subjects (4.4 percent) have perfect substitutes preferences ($\hat{\rho} \approx 1$), five subjects (11.1 percent) exhibit Cobb-Douglas preferences ($\hat{\rho} \approx 0$), and two subjects (4.4 percent) exhibit Leontief preferences ($\hat{\rho}$-values far below 0). More interestingly, there are many subjects with intermediate values of $\hat{\rho}$: 22 subjects (48.9 percent) show a preference for increasing total payoffs ($0.1 \leq \hat{\rho} \leq 0.9$). The 14 other subjects (31.1 percent) show a preference for reducing differences in payoffs ($-0.9 \leq \hat{\rho} \leq -0.1$). Therefore, like Charness and Rabin (2002), our results lean overall toward a social welfare conception of preferences.[9] To economize on space and to facilitate comparison across the two- and three-person budget set experiments, we will present the estimation results in the form of figures together with those of the three-person experiment.

3 Three-person budget sets

3.1 Data description

We next provide an overview of some important features of the three-person experimental data, which we summarize by reporting the distribution of allocations in a number of ways. Figure 4 depicts the distribution of the expenditure on tokens given to *others* as a fraction of total expenditure $p_o\pi_o/(p_s\pi_s + p_o\pi_o)$, and compares it with the analogous distribution in the two-person experiment. The distributions are quite similar, although, perhaps as expected, in the three-person case, subjects gave more than half of the tokens to *others* with greater frequency than in the two-person case.

[Figure 4 here]

Interestingly, only seven subjects (10.8 percent) in the three-person experiment spent, on average, more than half of their endowment on tokens given to *others*. We consider this to be surprisingly low, although no subjects in the two-person experiment spent more than half of their endowment on *others* on average. Overall, subjects gave approximately 26 percent of the tokens to *others* accounting for 25 percent of total expenditure, which

[9] Charness and Rabin (2005) extend the Charness-Rabin model, adding non-distributional parameters. They estimate population means from data on sequential two-person games and find significant effects for both distributional and non-distributional parameters.

is only marginally higher than the 19 percent and 21 percent, respectively, in the two-person experiment. Thus, the addition of a second *other* fell far short of generating a proportional increase in the overall level of giving.[10]

To investigate how $self$ trades off the payoff of person A against that of person B, Figure 5 depicts the distribution of the expenditure on tokens given to person A as a fraction of total expenditure on tokens given to $others$, $p_A\pi_A/(p_A\pi_A+p_B\pi_B)$. After screening the data for selfish allocations of 0.05 or less of total expenditure on tokens for $others$, which account for 50.2 percent of all allocations, we present the distribution based on the full sample, as well as distributions with the sample divided into three relative price terciles: intermediate relative prices of around 1 ($0.70 \leq p_A/p_B \leq 1.43$), steep prices ($p_A/p_B > 1.43$) and symmetric flat prices ($p_A/p_B < 0.70$). For the full sample, the distribution is nearly symmetric around the midpoint, indicating that $others$ are treated identically on average. For the distributions by tercile, the distribution for the steep tercile is bimodal with local modes at $0.95 - 1$ and $0.35 - 0.45$. For the flat tercile, the pattern is the mirror image. Thus, subjects respond symmetrically to changes in the relative price p_A/p_B. This is a natural result of the anonymity of $others$.

[Figure 5 here]

3.2 Econometric specification

Our subjects' CCEI scores are again sufficiently near one (see Appendix III) to justify treating the data as utility-generated. In order to recover the underlying distributional preferences and to assess any possible relationship between preferences for giving and social preferences, we assume a separable utility function, which may be expressed in terms of a subutility function $w_s(\pi_A, \pi_B)$ (*other* versus *other*) and macro utility function $v_s(\pi_s, w_s(\pi_o))$ ($self$ versus $others$). Additionally, we assume that the subutility function and the macro function are members of the CES family.

We therefore write:

$$U_s = [\alpha(\pi_s)^\rho + (1-\alpha)[\alpha'(\pi_A)^{\rho'} + (1-\alpha')(\pi_B)^{\rho'}]^{\rho/\rho'})]^{1/\rho}$$

where α (α') represents the relative weight on $self$ versus $others$ (*other* versus *other*) and ρ (ρ') expresses the curvature of the indifference curves for giving (social indifference curves). Clearly, when $\alpha = 1/3$ and $\alpha' = 1/2$,

[10]It is worthy of note that this suggests that π_o is a function only of the prices p_o and the total expenditure on $others$. The price p_s is relevant only insofar as it affects the total expenditure on $others$, as entailed by separability.

$U_s \to \pi_S + \pi_A + \pi_B$ (the purely utilitarian case) as $\rho, \rho' \to 1$, and $U_s \to \min\{\pi_S, \pi_A, \pi_B\}$ (the Rawlsian case) as $\rho, \rho' \to -\infty$. As $\rho, \rho' \to 0$, the indifference curves approach those of a Cobb-Douglas function. Further, any $0 < \rho, \rho' \leq 1$ indicate distributional preference weighted towards increasing total payoffs, whereas any $\rho, \rho' < 0$ indicate distributional preference weighted towards reducing differences in payoffs.

We use a two-stage estimation (first estimating parameters for the subutility function, and then using these parameter estimates in our estimation for the macro utility function) that is a direct generalization of the econometric specification in the two-person case. We refer the interested reader to Appendix V for precise details on the estimation. Before proceeding to the estimations, we omit the eight subjects (12.3 percent) with a CCEI score below 0.80, as their choices are not sufficiently consistent to be considered utility-generated. We also screen subjects with readily identifiable preferences. These include 24 subjects with uniformly selfish allocations (average $p_s\pi_s/m \geq 0.95$), as well as three pure utilitarians (ID 139, 154 and 199) and one pure Rawlsian (ID 158).[11] This leaves a set of 29 subjects (44.6 percent) for whom we need to recover the underlying distributional preferences by estimating the CES model. Appendix V also presents, by subject, the results of the estimations. Throughout this section, whenever we list the number and percentages of subjects with particular properties, we will be considering the 33 subjects with consistent non-selfish preferences. These are the 29 subjects listed in Appendix V plus the four subjects whose choices correspond precisely to utilitarian or Rawlsian distributional preferences.

3.3 Preferences for giving

The estimates of the two relevant parameters for the macro function $v_s(\pi_s, w_s(\pi_o))$, α and ρ, reflect preferences for giving ($self$ versus $others$). As a preview, Figure 6 shows a scatterplot of $\hat{a}_n$ and $\hat{\rho}_n$, and compares the estimated parameters with the analogous parameters for the two-person experiment (to facilitate presentation of the data, subjects ID 3, 46, 55, 73, 158 and 179 are excluded because they have very negative $\hat{\rho}$-values). Note that in both the two- and three-person experiments there is considerable heterogeneity in both parameters, $\hat{a}_n$ and $\hat{\rho}_n$. Perhaps not surprisingly, $\hat{a}_n > 1/2$ for all n in the two-person case, whereas in the three-person case $\hat{a}_n > 1/3$ for all n.

[11] One subject (ID 199) perfectly implemented utilitarian social preferences and implemented utilitarian preferences for giving with slight imperfections. Throughout this section, we will also classify this subject as utilitarian.

[Figure 6 here]

Of the 33 subjects with consistent, non-selfish preferences, eight subjects (24.2 percent) have preferences for giving that are easily identifiable: four subjects (12.1 percent) have perfect substitutes preferences for giving ($\hat{\rho} \approx 1$), three subjects (9.1 percent) exhibit Leontief preferences ($\hat{\rho}$-values far below 0) and one subject exhibits Cobb-Douglas preferences ($\hat{\rho} \approx 0$). There are additionally many subjects with intermediate values of $\hat{\rho}$: 18 subjects (54.5 percent) show a preference for increasing total payoffs of $self$ and $others$ ($0.1 \leq \hat{\rho} \leq 0.9$) and seven subjects (21.2 percent) show a preference for reducing differences in payoffs between $self$ and $others$ ($-0.9 \leq \hat{\rho} \leq -0.1$). Figure 7 presents the distribution of $\hat{\rho}_n$ for the 33 subjects with consistent, non-selfish preferences, rounded to a single decimal and compares it with the analogous distribution in the two-person experiment. The distributions are very similar and skewed to the right so that, as in the two-person experiment, our results lean overall toward a social welfare conception of preferences for giving.

[Figure 7 here]

3.4 Social preferences

The estimated parameters for the subutility function $w_s(\pi_A, \pi_B)$, α' and ρ', reflect social preferences (*other* versus *other*). We cannot reject the hypothesis that $\hat{\alpha}'_n = 1/2$ for all but four subjects at the 95 percent significance level (24 subjects (72.7 percent) have $0.45 \leq \hat{\alpha}' \leq 0.55$, and this increases to a total of 31 subjects (93.9 percent) if we consider $0.4 \leq \hat{\alpha}' \leq 0.6$). This provides strong support for the inference that subjects do not have any bias towards a particular person, A or B. Figure 8 presents the distribution of $\hat{\rho}'_n$, which parameterizes attitudes towards the efficiency-equity tradeoff concerning $others$, rounded to a single decimal. Of the 33 subjects with consistent, non-selfish preferences, 14 subjects (42.4 percent) have social preferences that are easily identifiable: five subjects (15.2 percent) have perfect substitutes social preferences ($\hat{\rho}' \approx 1$), three subjects (9.1 percent) exhibit Cobb-Douglas social preferences ($\hat{\rho}' \approx 0$), and six subjects (17.2 percent) exhibit extreme aversion to inequality for Leontief social preferences ($\hat{\rho}'$-values far below 0). Since $others$ are treated symmetrically by $self$, we conclude that both utilitarian and Rawlsian social preferences are well represented among our subjects. Moreover, 17 subjects (51.5 percent) show a preference for increasing the total payoffs of $others$ ($0.1 \leq \hat{\rho}' \leq 0.9$) while only two subjects (6.1 percent) show aversion to inequality between *others* ($-0.9 \leq \hat{\rho}' \leq -0.1$). We thus conclude that a significant majority of subjects are concerned with increasing the aggregate payoffs of *others* rather than reducing differences in payoffs between *others*.

[Figure 8 here]

3.5 Preferences for giving versus social preferences

Finally, we make within-subject comparisons of the estimated CES parameter of the macro utility function $\hat{\rho}$ (preferences for giving) and the parameter of the subutility function $\hat{\rho}'$ (social preferences). Figure 9 shows a scatterplot of $\hat{\rho}_n$ and $\hat{\rho}'_n$ (subjects ID 148, 158, 161, 177, 179, 191, and 197 are omitted because they have very negative values of $\hat{\rho}_n$ or $\hat{\rho}'_n$). The data are concentrated in the upper right quadrant ($0 < \hat{\rho}_n, \hat{\rho}'_n \leq 1$). Of the 33 subjects with consistent, non-selfish preferences, 21 subjects (63.6 percent) have positive values for both $\hat{\rho}_n$ and $\hat{\rho}'_n$, so that for a majority of subjects, both preferences for giving and social preferences emphasize increasing aggregate payoffs rather than reducing differences in payoffs. Two of the remaining subjects on the graph and six of the seven subjects omitted from the graph because of low $\hat{\rho}_n$ or $\hat{\rho}'_n$ values are located in the lower left quadrant ($\hat{\rho}_n, \hat{\rho}'_n < 0$). Hence, a total of eight subjects (24.2 percent) emphasize reducing difference in payoffs for both social preferences and preferences for giving.

[Figure 9 here]

Interestingly, four subjects exhibit opposite tradeoffs between efficiency and equity in their social preferences and preferences for giving. Two subjects (ID 157 and 193), who fall in the lower right quadrant ($0 < \hat{\rho}_n \leq 1$ and $\hat{\rho}'_n < 0$), show a preference for increasing total payoffs of *self* and *others* while reducing differences in payoffs between *others*. In contrast, two subjects (ID 148 who is omitted from the graph because of a low $\hat{\rho}_n$-value and ID 185) who fall in the top left quadrant ($\hat{\rho}_n < 0$ and $0 < \hat{\rho}'_n \leq 1$) show a preference for reducing differences in payoffs between *self* and *others* while increasing total payoffs of *others*. However, in only two of these four cases are both $\hat{\rho}_n$ and $\hat{\rho}'_n$ significantly different from zero. In conclusion, although we find considerable heterogeneity of attitudes towards the efficiency-equity tradeoff *across* subjects, there is a strong association between preferences for giving and social preferences *within* subjects. Thus, at least with respect to preferences concerning efficiency versus equity subjects apply the same distributive principles universally to $self$ versus $others$, and among anonymous $others$.

4 Step-shaped sets

Since some canonical models of distributional preferences posit not well-behaved preferences, and given that such preferences cannot be detected using choices on linear budget sets, we next turn to step-shaped sets. Appendix VI shows the distribution of decisions aggregated to the subject level by summarizing the number of decisions corresponding to each subset of the step-shaped constraint depicted in Figure 1 above, with an additional column that lists the number of equal allocations π^e. Whenever possible, in Appendix VI, we also adhere to the preference classifications described in the model of Charness and Rabin (2002).

As a preliminary step, we examine the extent to which subjects damage both $self$ and $other$ by choosing *strictly* interior allocations. With the narrow one token confidence interval, of the $58 \times 50 = 2900$ allocations, only 186 allocations (6.4 percent) were not on the step-shaped constraint. Of these, 156 allocations (83.9 percent) are concentrated in five subjects (8.6 percent), with the remaining 30 spread among the 53 other subjects. We do not observe any patterns in these five subjects' choices that could effectively distinguish them from random allocations. Thus, their behavioral rules are not clear and there is no taxonomy that allows us to classify their behaviors unambiguously. We therefore omit the five subjects (ID 89, 92, 93, 116, and 117) with many interior allocations (17, 50, 20, 37, and 32 respectively) from the analyses below, leaving a total of 53 subjects (91.4 percent). We also screen out the 30 interior allocations distributed among the remaining subjects.[12]

Of the 53 subjects listed in Appendix VI, 43 (81.1 percent) have cleanly classifiable preferences. Of those, 26 subjects (49.0 percent) have lexself preferences ($\pi = \pi^s$), three subjects (5.7 percent) have competitive preferences ($\pi = \pi^c$), seven subjects (13.2 percent) exhibit selfish preferences ($\pi_s = \pi_s^s$ and $0 \leq \pi_o \leq \pi_o^s$), and seven subjects (13.2 percent) exhibit social welfare preferences ($\pi = \pi^s$ or $\pi = \pi^o$). Of the 10 remaining subjects, nine (17.0 percent) have intermediate preferences that incorporate elements of preferences for $self$, concerns for $other$, and difference aversion. The remaining subject exhibits preferences that incorporate both difference

[12] Appendix VI also lists, by subject, the number of interior allocations, and the average distance of these allocations from the constraint.

aversion and social welfare preferences. Thus, we find a large fraction of subjects that incorporate elements of less well-behaved preferences which could not have been detected given choices on linear budget sets, as implied by Afriat's Theorem. We next provide a more refined analysis and discussion of the behaviors of each preference type by examining the characteristics of their individual decisions beyond their broad classification in the various subsets of the step-shaped constraint.

lexself preferences We begin with the 26 subjects (49.0 percent) whose choices correspond to lexself preferences. Of these, 15 subjects choose π^s in all 50 decision-rounds. For all but one of the remaining subjects that we classify as lexself, at least 49 allocations are within two tokens of π^s. To be sure, always choosing π^s could potentially be consistent also with social welfare or difference aversion preferences. However, given the rich menu of step-shaped sets faced by each subject, social welfare preferences that generate $\pi = \pi^s$ for all allocations would require a great "weight" on $self$.

To illustrate this point, we use payoff calculations to measure the relative *surplus* of π^s and π^o defined by $(\pi^s_s - \pi^o_s)/(\pi^o_o - \pi^s_o)$. That is, the relative surplus depicts the surplus for $self$ $\pi^s_s - \pi^o_s$ (the difference between the payoffs for $self$ at π^s and π^o) as a fraction of the surplus for *other* $\pi^o_o - \pi^s_o$. The lower bound on the relative surplus varies by subject but it is uniformly low and ranges empirically from 0.07 to 0.33. Accordingly, these subjects chose π^s even when π^o was relatively inexpensive, so that if these subjects did indeed have social welfare preferences, the weight on *other* would be sufficiently low so that, for practical purposes, preferences could be approximated as being lexicographic for $self$ over *other*.

The allocations of these 26 subjects are also difficult to reconcile with difference averse preferences, since according to the Charness-Rabin model this would imply choosing $\pi = \pi^e$ when $\pi^e \in \Pi^1$. Of these, 15 subjects faced sets in which $\pi^e \in \Pi^1$, and the equal allocation π^e was never chosen. Further, while the allocations of the subjects that always choose $\pi = \pi^s$ are also consistent with perfectly selfish preferences, selfishness suggests no systematic pattern in the choice of π_o, whereas we always observe $\pi_o = \pi^s_o$. Thus, any explanation for the behavior of these subjects that relies upon social welfare, difference aversion, and selfishness seems inadequate.

Social welfare We next analyze the behavior of the seven subjects (13.2 percent) whose choices correspond to social welfare preferences. Of these, five subjects choose either π^s or π^o in all 50 decision-rounds, with the remaining two subjects making all but four allocations each within two tokens of π_s or π_o. To probe further the validity of our classification, we consider whether the relative surplus $(\pi_s^s - \pi_s^o)/(\pi_o^o - \pi_o^s)$ is significantly different when subjects choose π^s relative to when they choose π^o. Table 1 summarizes the means, standard deviations, and number of observations for each of these seven subjects, according to whether π^s or π^o was chosen. For each of these subjects, relative surplus is higher when π^s is chosen, and this difference is significant at the 1 percent level in all cases. This further bolsters the validity of our classification of these subjects as having social welfare preferences.'

[Table 1 here]

Difference aversion Next, we turn to the ten subjects (17.0 percent) that exhibit *self*-damaging behavior. These subjects display a balance of selfishness and difference aversion that leads them to make *self*-damaging allocations on Π^2 that cannot be accommodated by the canonical models of distributional preferences encapsulated in Charness and Rabin (2002). Of these, at least four subjects (ID 85, 98, 102, and 109) appear to be governed by difference aversion, as $\pi = \pi^e$ is chosen frequently. For these subjects, deviations from equality are dominated by allocations π with $\pi_s > \pi_o$, and the extent of inequality is increasing in π_s^s and decreasing in π_o^s. By contrast, inequality is uncorrelated with either π_s^o or π_o^o. Thus, these subjects made choices that may reflect a combination of selfishness and difference aversion. This is illustrated in the first column of Table 2, which reports the results of a regression predicting the extent of inequality $|\pi, \pi^e|$ for the *self*-damaging allocations chosen by these four subjects.

Additionally, five subjects (ID 100, 103, 111, 114, and 132) choose many *self*-damaging allocations $\pi \in \Pi^2$, all of which decrease inequality, though the distribution of allocations is more dominated by allocations with $\pi \neq \pi^e$. As before, the extent of inequality is increasing in π_s^s and decreasing in π_o^s. Thus, the choices made by these subjects also correspond to selfishness and difference aversion, though with a greater weight put on *self*. A linear regression analysis confirms these results. This is summarized in the second column of Table 2. Finally, the choices of the remaining subject (ID 127) are distributed among π^s, π^o, Π^2, and Π^3 and thus correspond to a combination of selfish, difference aversion, and social welfare preferences. As with our social welfare subjects, the relative surplus $(\pi_s^s - \pi_s^o)/(\pi_o^o - \pi_o^s)$ is highly correlated with this subject's choice of π^s versus π^o; and as with our subjects whose choices fit with a combination of selfish and difference averse preferences, inequality in allocations $\pi \in \Pi^2$ are increasing in π_s^s and decreasing in π_o^s.

[Table 2 here]

Selfish and competitive Lastly, we consider the subjects that almost exclusively chose allocations with $\pi_s = \pi_s^s$. We first consider the seven subjects whose choices fit with selfish preferences. Of these, six subjects chose $\pi_s = \pi_s^s$ and $\pi_o \leq \pi_o^s$ in all (non-interior) decision-rounds and one additional subject (ID 95) chose $\pi_s = \pi_s^s$ in 45 rounds. We say that these subjects made choices that reflect selfish preferences if the choice of π_o is random due to apparent indifference to *other*. We examined the possibility that there may be a competitive element to behaviors of these subjects by noting that if this were the case then the Charness-Rabin model predicts that π_o should increase with π_s^s, for any given π_o^s. However, a simple regression analysis indicates no relation between potential inequality and π_o for the pooled sample and the behavior of no individual subject exhibits a significant relation between π_s^s and π_o.

Finally, three subjects (ID 78, 125, and 126) choose the competitive allocation $\pi^c = (\pi_s^s, 0)$. We argue that these three subjects are best classified as competitive, since some effort is required in navigating the mouse to π^c rather than choosing randomly on Π^1. However, we note that the step-shaped sets we employ are not ideally suited to identifying competitive preferences. To distinguish more effectively between selfish and competitive preferences, a useful modification would be to present subjects with choice sets where the subset Π^1 of the constraint is upward sloping. This comes at a cost, however, since this would confound our identification of selfish versus lexself preferences. Since this distinction involves at most a small fraction of subjects, we leave this extension for future work.

5 Conclusion

Our results emphasize both the prominence and the heterogeneity of other-regarding behaviors. In the budget set experiments, the existence of a well-behaved rationalizing preference ordering is confirmed by the data for a significant majority of our subjects. By contrast, less well-behaved cases, which also display features that cannot be accommodated by prominent models of distributional preferences, are confirmed by the data from the step-shaped experiment. We emphasize that these results are not in conflict:

consistency with a well-behaved utility function is a direct consequence of the linear budget constraint, as Afriat's (1967) theorem makes clear.

We also find that subjects overall lean toward a social-welfare conception of preferences and that there is a strong correlation between the equality-efficiency tradeoffs subjects make in their preferences for giving and social preferences. We thus conclude that subjects' special concern for themselves seems not to distort impartiality with respect to efficiency-equity tradeoffs (ρ in the CES model) nearly as much as it does with respect to the indexical weights that they place on payoffs to $self$ versus *others* (α in the CES model). And insofar as this is so, it suggests that at least with respect to preferences concerning efficiency versus equity, subjects actually act on unified distributive principles.

References

[1] **Afriat, Sidney N.** 1967. "The Construction of a Utility Function from Expenditure Data." *International Economic Review* 8(1): 67-77.

[2] **Afriat, Sidney N.** 1972. "Efficiency Estimation of Production Functions." *International Economic Review* 13(3): 568-598.

[3] **Andreoni, James and John H. Miller.** 2002. "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism." *Econometrica*, 70(2): 737-753.

[4] **Bolton, Gary E.** 1991. "A Comparative Model of Bargaining: Theory and Evidence." *American Economic review*, 81: 1096-1136.

[5] **Bolton, Gary E. and Axel Ockenfels.** 1998. "Strategy and Equity: An ERC-Analysis of the Gueth-van Damme Game." *Journal of Mathematical Psychology*, 42: 215-226.

[6] **Bolton, Gary E. and Axel Ockenfels.** 2000. "ERC: A Theory of Equity, Reciprocity, and Competition." *American Economic Review*, 90: 166-193.

[7] **Bronars, Stephen G.** 1987. "The power of nonparametric tests of preference maximization." *Econometrica* 55(3): 693-698.

[8] **Camerer, Colin F.** 2003. "Behavioral Game Theory: Experiments in Strategic Interaction." Princeton University Press.

[9] **Charness, Gary and Matthew Rabin.** 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics*, 117: 817-869.

[10] **Charness, Gary and Matthew Rabin.** 2005. "Expressed preferences and behavior in experimental games." *Games and Economic Behavior*, 53: 151-169.

[11] **Choi, Syngjoo, Ray Fisman, Douglas M. Gale, and Shachar Kariv.** 2007. "Consistency, Heterogeneity, and Granularity of Individual Behavior under Uncertainty." Mimo.

[12] **Fehr, Ernst and Klaus M. Schmidt.** 1999. "A Theory of Fairness, Competition and Co-operation." *Quarterly Journal of Economics*, 114: 817-868.

[13] **Houtman, Martijn, and J. A. H. Maks.** 1985. "Determining all Maximal Data Subsets Consistent with Revealed Preference." *Kwantitatieve Methoden* 19: 89-104.

[14] **Karni, Edi and Zvi Safra.** 2002. "Individual Sense of Justice: A Utility Representation." *Econometrica*, 70: 263–284.

[15] **Levine, David K.** 1998. "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics*, 1: 593-622.

[16] **Loewenstein, George E., Leigh Thompson and Max H. Bazerman.** 1989. "Social utility and decision making in interpersonal contexts." *Journal of Personality and Social Psychology*, 57(3): 426-441.

[17] **Rabin, Matthew.** 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, 83: 1281-1302.

[18] **Varian, Hal R.** 1982 "The Nonparametric Approach to Demand Analysis." *Econometrica* 50(4): 945-972.

[19] **Varian, Hal R.** 1991. "Goodness-of-Fit for Revealed Preference Tests." mimeo.

Table 1: The relative surplus of subjects whose choices correspond to social welfare preferences

ID	π^s obs.	π^s mean	π^s sd	π^o obs.	π^o mean	π^o sd
80	45	1.514	1.296	5	0.359	0.366
101	44	1.506	1.044	6	0.212	0.087
105	45	1.275	1.057	5	0.262	0.130
106	48	1.229	0.966	2	0.333	0.018
121	45	1.692	1.607	5	0.325	0.243
132	36	1.804	0.982	10	0.851	0.670
134	43	1.779	1.583	3	0.349	0.217

Table 2: Estimation results for subjects that exhibit *self* -damaging behavior

	(1)	(2)
π_s^s	0.221*	0.473*
	(0.063)	(0.042)
π_o^s	-0.254	-0.513*
	(0.154)	(0.084)
π_o^o	-0.023	-0.010
	(0.051)	(0.032)
π_s^o	0.129	-0.028
	(0.124)	(0.069)
obs.	87	122
R^2	0.44	0.70

Subjct ID: (1) 85, 98, 102, 109. (2) 100, 103, 111, 114, 132. Standard errors in parentheses. * 1 percent significance level.

Figure 1: The step-shaped set

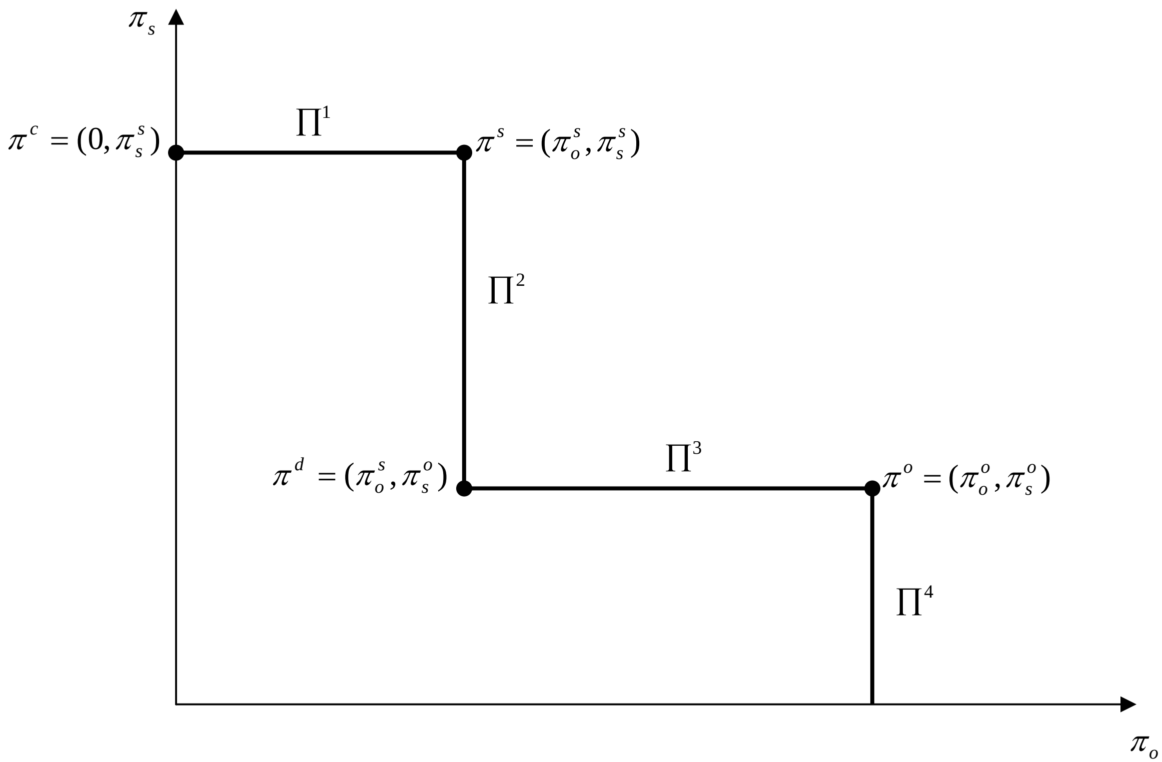

Figure 2: The distribution of expenditure on tokens given to *other* as a fraction of total expenditure

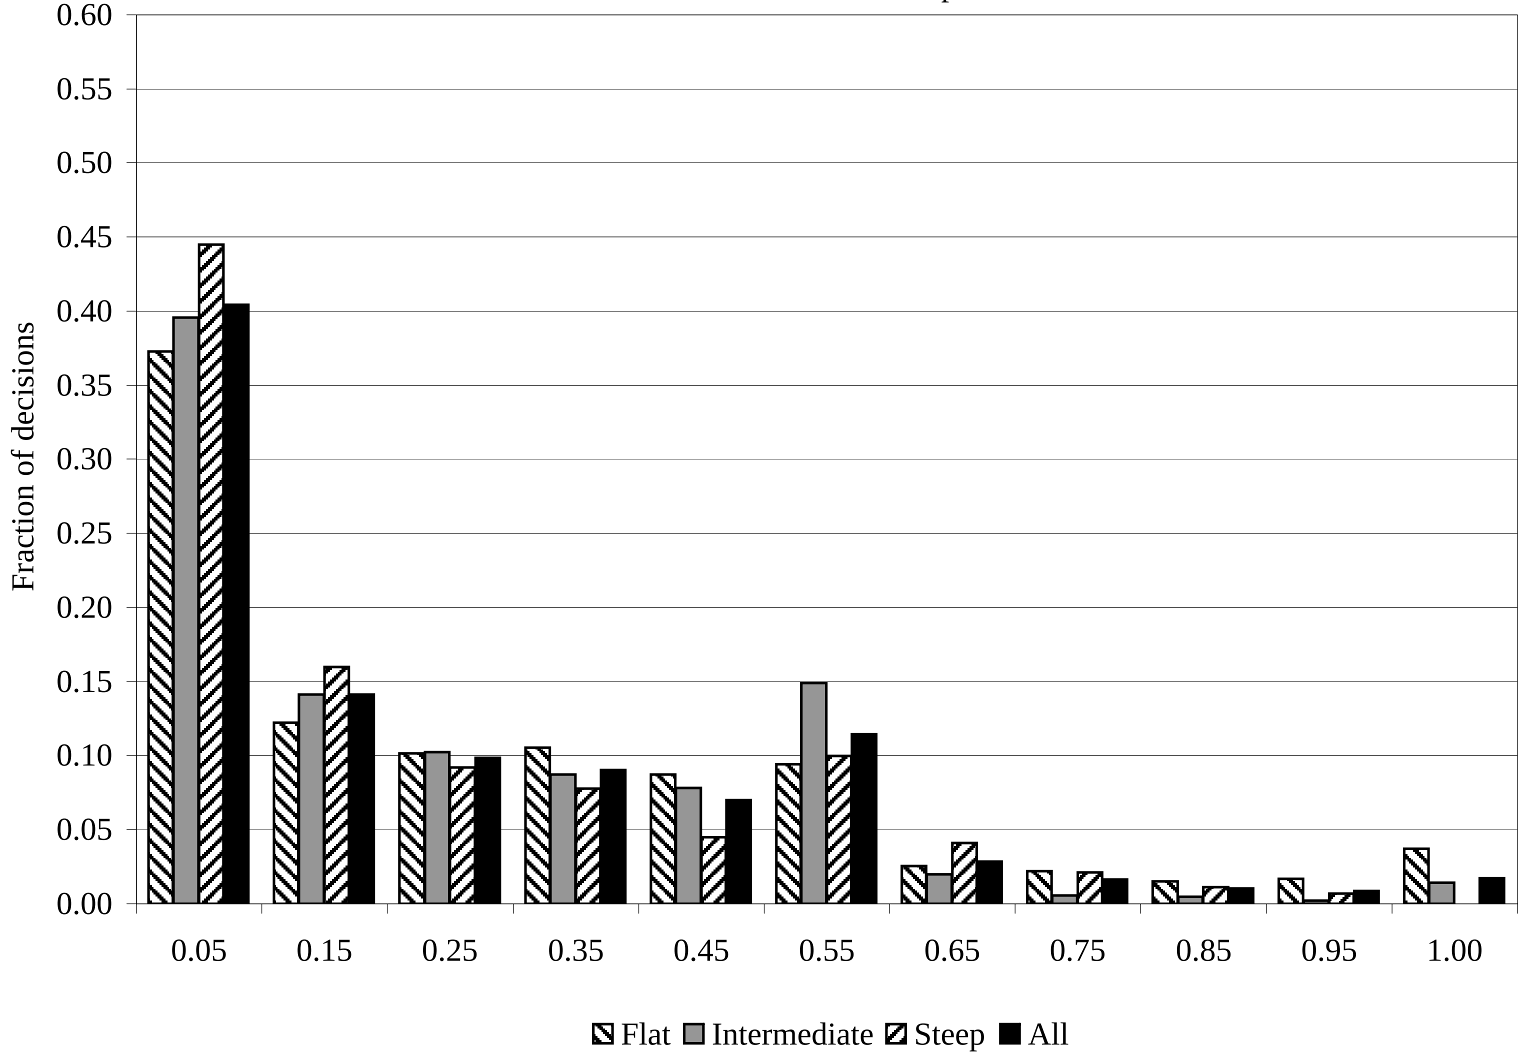

Figure 3: The distributions of Afriat's (1972) critical cost efficiency index (CCEI)

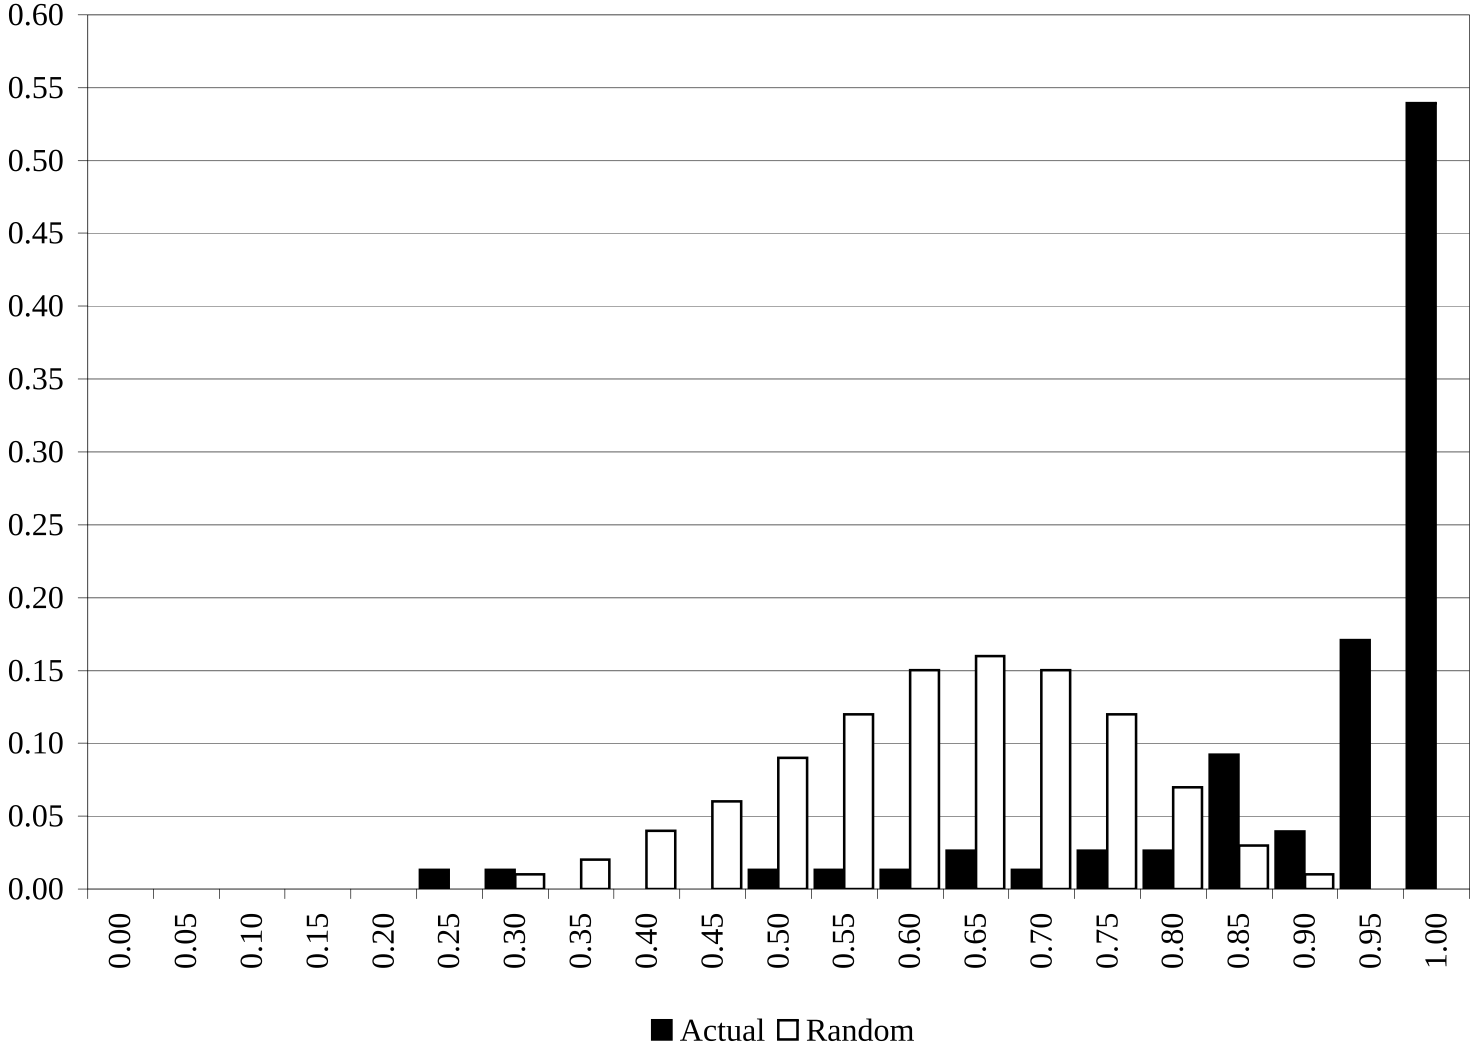

Figure 4: The distribution of expenditure on tokens given to *others*
as a fraction of total expenditure in the two- and three-person budget set experiments

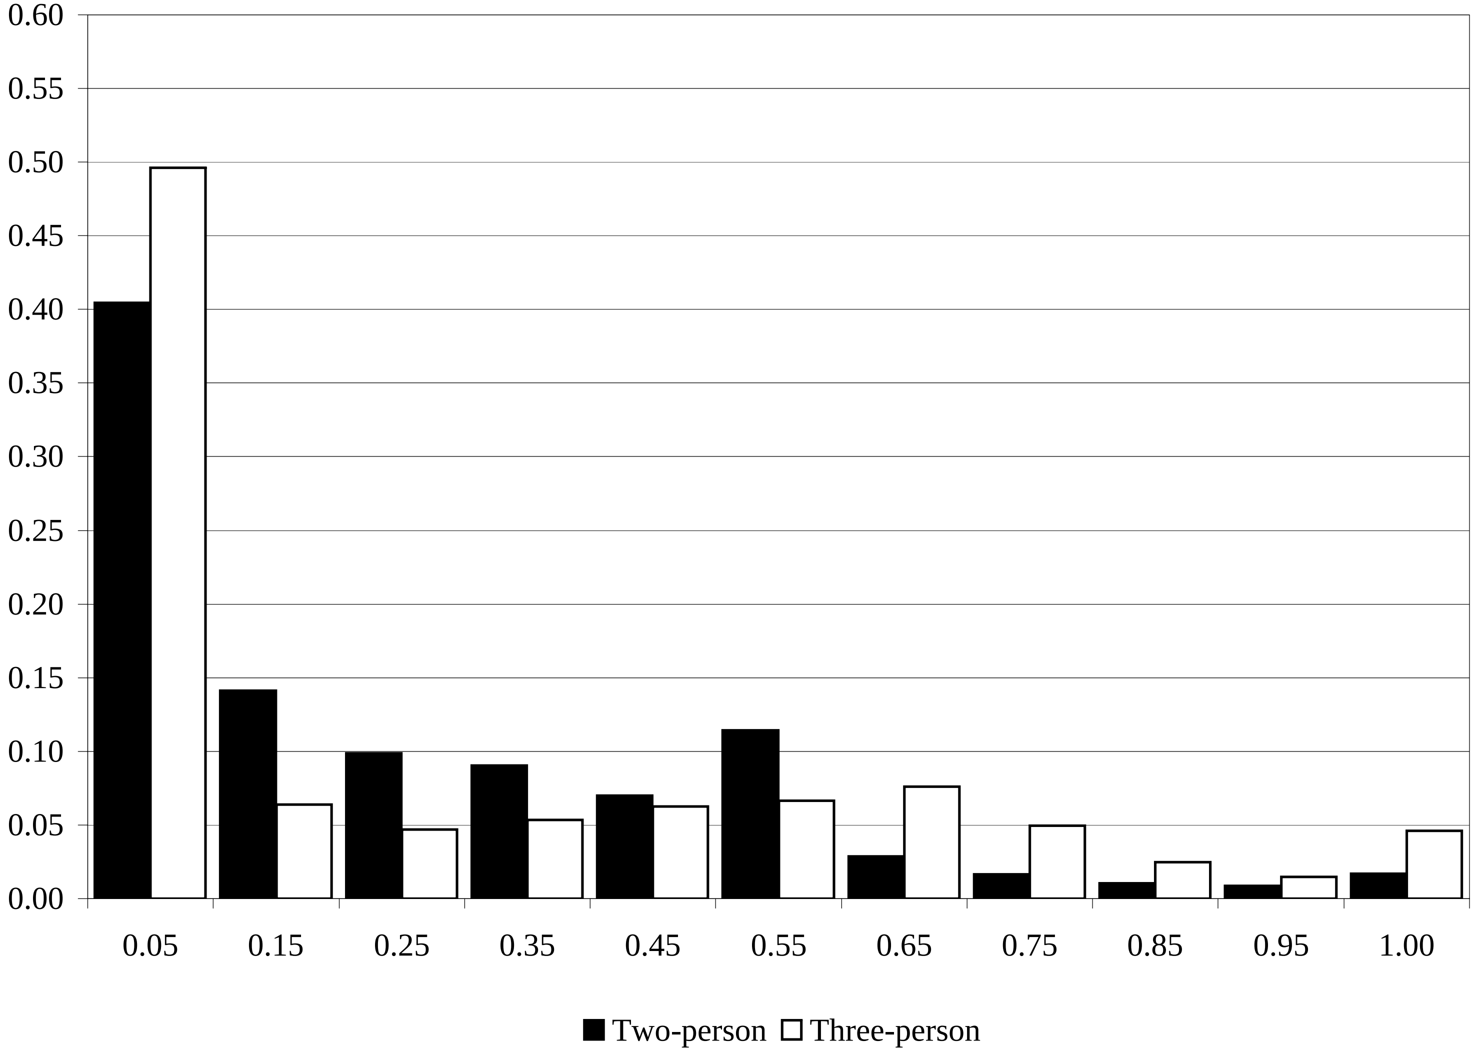

Figure 5: The distribution of expenditure on tokens given to person *A*
as a fraction of total expenditure on tokens given to *others*

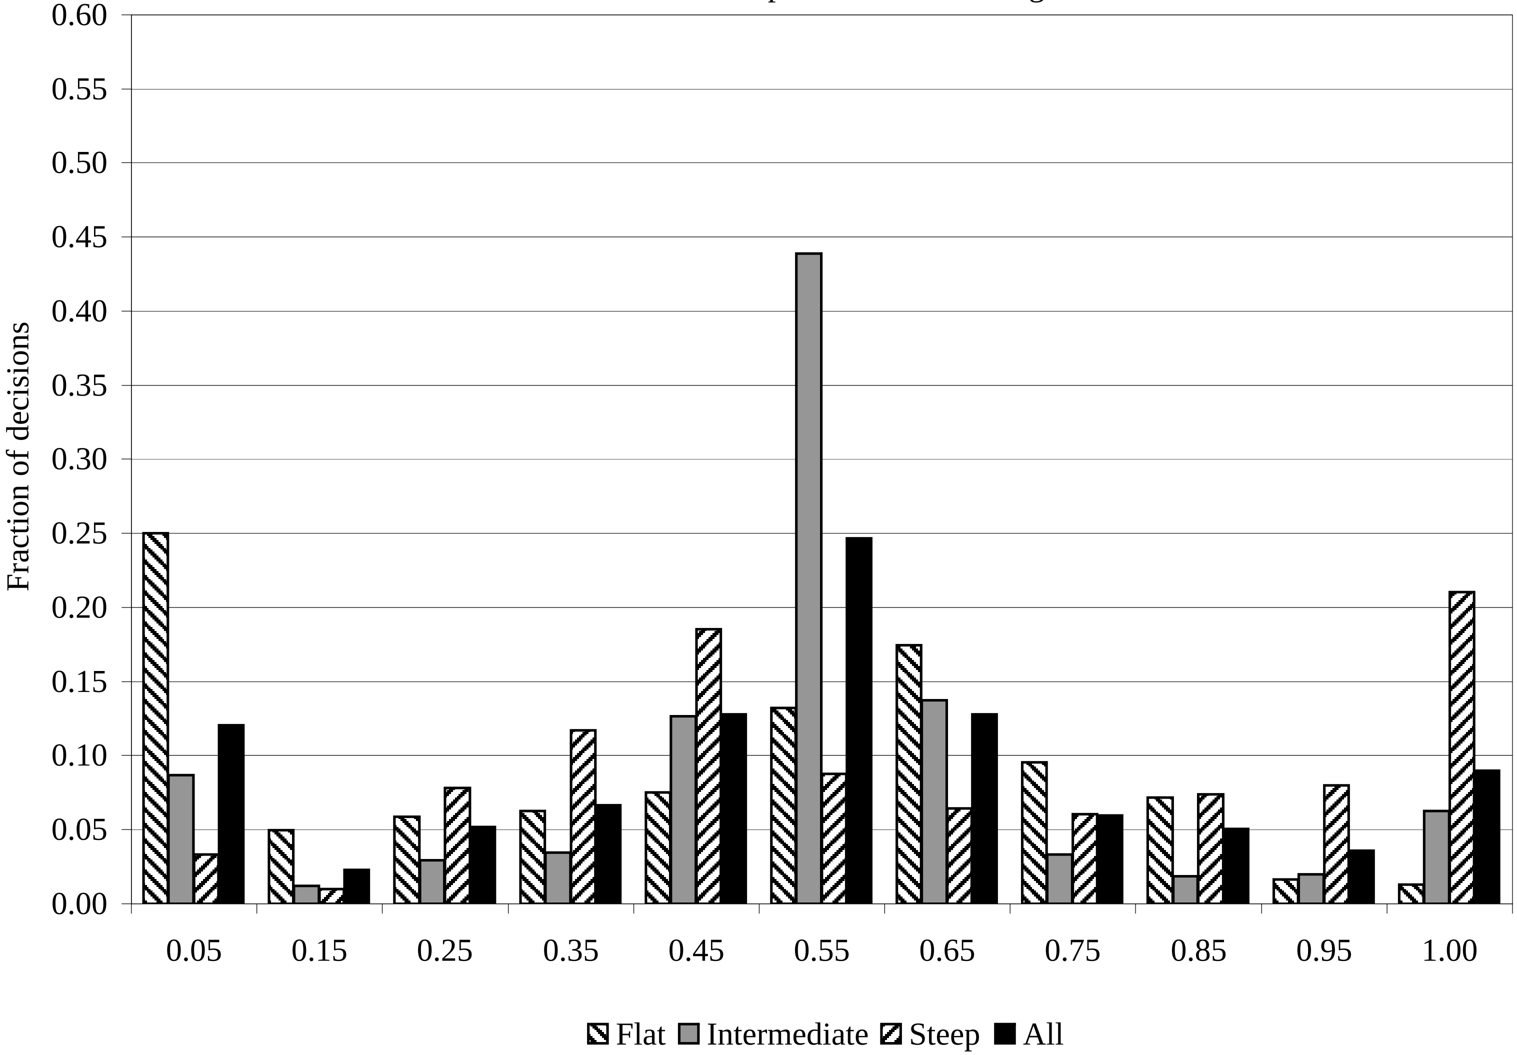

Figure 6: Scatterplot of the CES estimates α and ρ in the two- and three-person budget set experiments

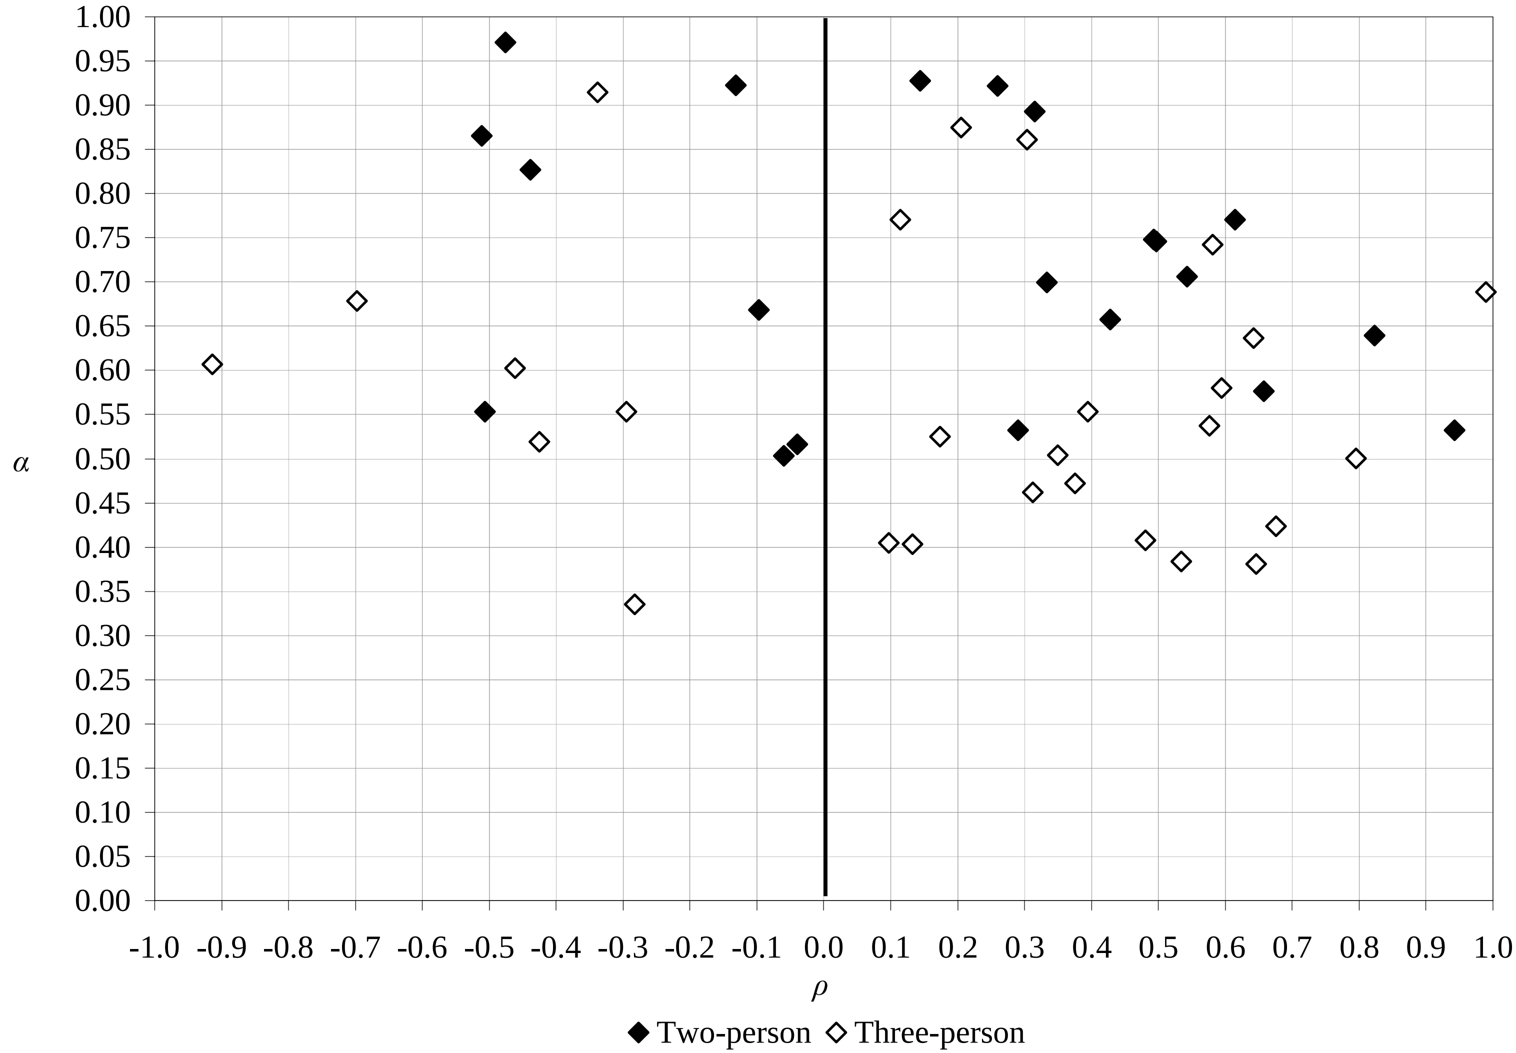

Figure 7: The distribution of the CES parameter ρ in the two- and three-person budget set experiments

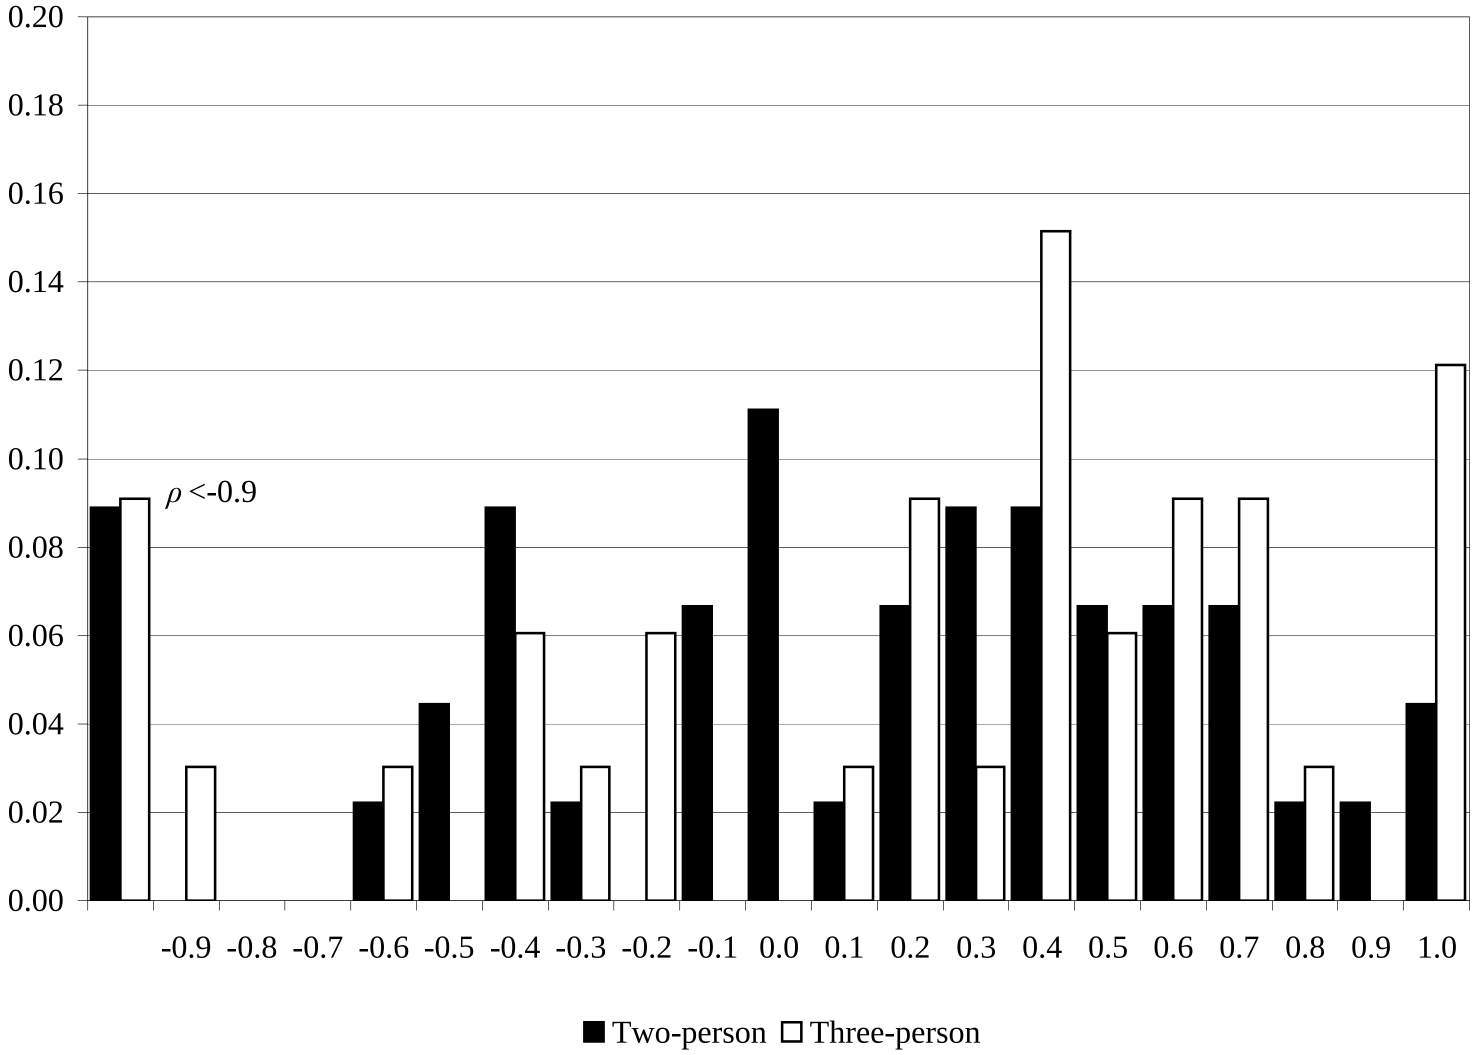

Figure 8: The distribution of the CES sub utility parameter ρ' in the three-person budget set experiment

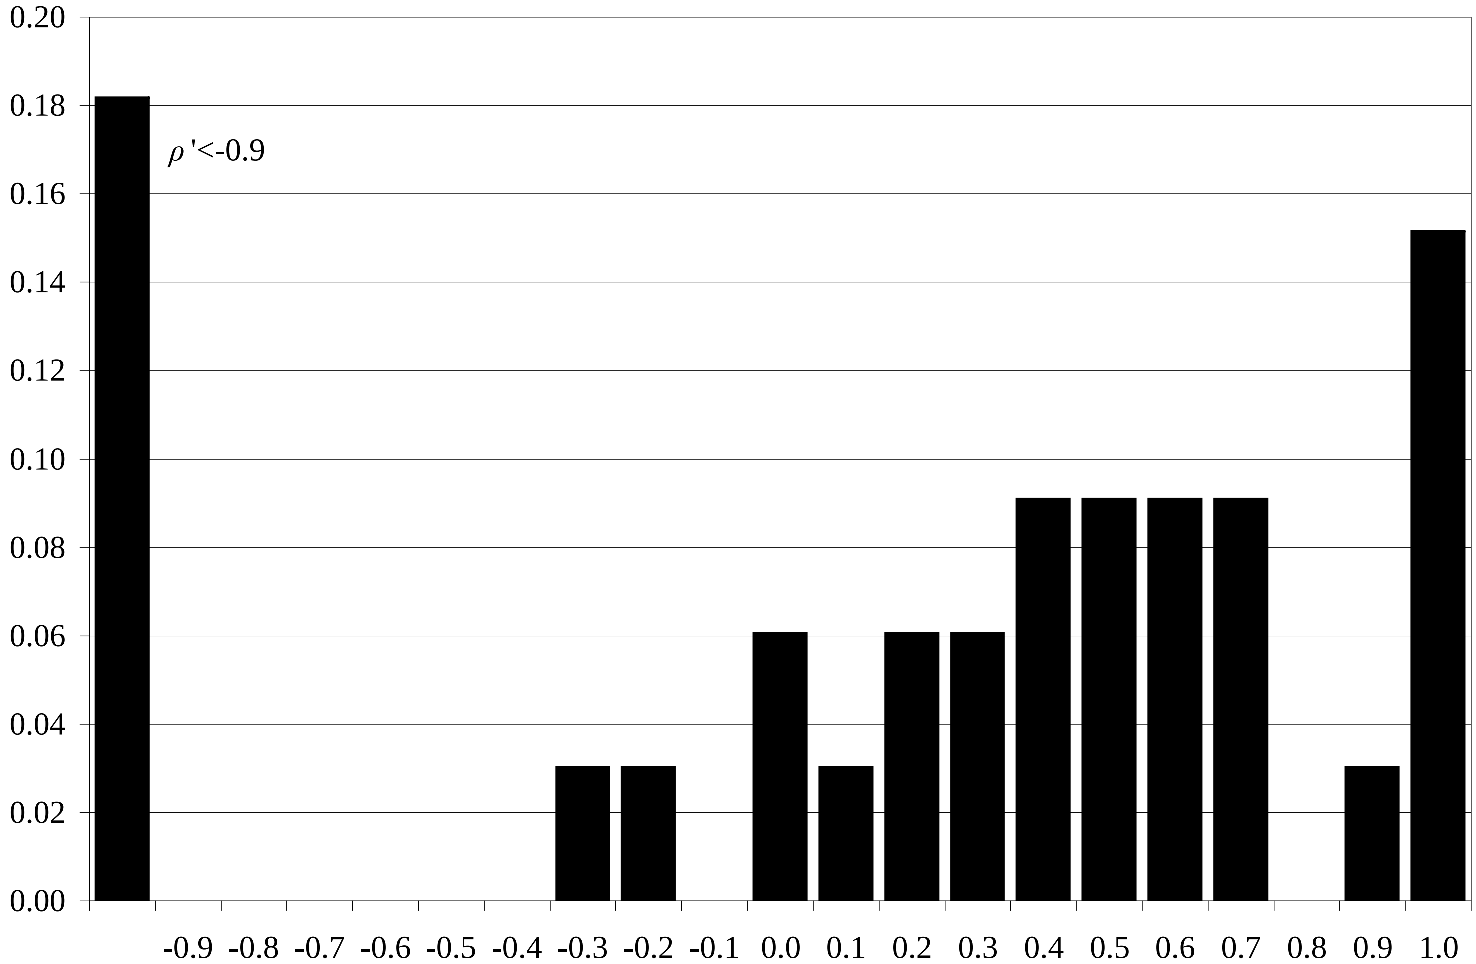

Figure 9: Scatterplot of the CES estimates ρ and ρ' in the three-person experiment