Analyzing Choice with Revealed Preference:
Is Altruism Rational?

by
James Andreoni
Department of Economics
University of Wisconsin
Madison, Wisconsin 53706

and

John H. Miller
Department of Social and Decision Sciences
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

1.   Introduction

Consumer theory is based on three axioms of choice: a binary preference ordering is complete, reflexive and transitive. If preferences adhere to these axioms, then they can be characterized by a utility function. If preferences are well behaved—that is, they are convex and monotonic—then the utility function will generate smooth downward sloping demand curves.

How do we know if our fundamental assumptions are valid? Samuelson (1948) gave us an elegant answer to this question, the theory of Revealed Preference:[1]

> **Definition: Directly Revealed Preferred.** *An allocation X is directly revealed preferred to a different allocation Y if Y was in the budget set when X was chosen.*

Then if a well-behaved utility function could have generated the data, the data will satisfy *WARP*:

> **Weak Axiom of Revealed Preference (WARP):** *If allocation X is directly revealed preferred to Y, then Y cannot be directly revealed preferred to X.*

WARP is a necessary condition on choices to be consistent with utility theory. However, it is not a sufficient condition. For this we need a stronger axiom.

> **Definition: Revealed Preferred.** *If an allocation A is directly revealed preferred to B, B is directly revealed preferred to C, C is directly revealed preferred ... to Z, and A and Z are not the same bundle, then A is revealed preferred to Z. That is, the Revealed Preferred relation is the transitive closure of the Directly Revealed Preferred relation.*

> **Strong Axiom of Revealed Preference (SARP):** *If allocation X is revealed preferred to Y, then Y will never be revealed preferred to X.*

If preferences are strictly convex, then choices will conform to SARP. Moreover, if choices conform to SARP then there exists a well-behaved preference ordering that could have generated the data. That is, utility theory is valid for the data observed.

SARP is a strong tool for economists to use to verify that an individual's behavior is "rational;" that is, it is consistent with neoclassical choice theory. However, SARP is still a bit restrictive in that it requires preferences to be strictly convex. Afriat (1967) and Varian (1982) showed that a fully general axiom that is both necessary and sufficient for the existence of a utility function—even one that has flat spots on indifference curves—is GARP:

> **Generalized Axiom of Revealed Preference (GARP):** *If an allocation X is revealed preferred to Y, then Y is never strictly directly revealed preferred to X, that is, X is never strictly within the budget set when Y is chosen.*

---

[1] See Varian (1992) for a more detailed discussion of revealed preference.

In this paper we discuss an application of revealed preference to a common occurrence in experiments—kindness among subjects. In many experiments, including Prisoner's dilemma, public goods, and bargaining experiments, subjects are often found to act benevolently toward each other. The immediate reaction when these findings began appearing was that subjects were "irrational" because they did not choose to maximize their own monetary payoffs. Some suggested that neoclassical theory had failed, and others suggested that economics needed to appeal to other behavioral sciences to understand this "non-economic" behavior. But the axioms of choice indicate that what is "rational" is what is consistent, that is, it can be characterized by convex preferences. Hence, whether this benevolent behavior is rational is an empirical question that the experimental economist is perfectly suited to answer.

The hypothesis to explore here is that subjects have consistent preferences for altruism. To address this, we designed an experiment that would measure a subject's simple preferences over allocations between themselves and another subject. Most social dilemma experiments, like Prisoner's dilemma, public goods, or alternating offer bargaining, can be decomposed into unilateral allocation problems. Since we see people allocating some of the payoff to themselves and some to the other subject, we pose this question: Can choices over this allocation process be "rational"?

Let π_s be the payoff a subject allocates to "self" and π_o be the payoff the subject allocates to the "other." Then the research question can be restated as, "Can behavior in experiments be characterized by a quasi-concave utility function of the form $U_i = u_i(\pi_s, \pi_o)$?"

2.   The Choice Task

The data we describe here is a subset of that used our paper, Andreoni and Miller (2002). Subjects were presented eight different allocation tasks, in random order. Each choice endowed the subject with a budget of tokens which were worth different numbers of points to the two subjects. Points were all worth $0.10 to all subjects. The budget of tokens was either 100, 75, 60 or 40 tokens, and tokens were worth either 1, 2 or 3 points each. Hence, by varying the endowments and points, we were able to create various budgets of payoffs with different relative prices. The budgets were chosen to intersect often so as to give the strongest test of revealed preference. Table 1 presents the eight choices presented to subjects.

[TABLE 1 HERE.]

After subjects made allocation decisions in all eight budgets, the experimenter randomly chose one to carry out with another subject chosen at random from the room. We ran four sessions of the experiment, each with 35 or 36 subjects, for a total of 142 subjects, and took great pains to protect the anonymity of all subjects. For details of the experiment and procedures, see Andreoni and Miller (2002).

3. Checking GARP

Table 2 shows the violations of GARP. Over all 142 subjects, fewer than 10 percent had a single violation. For most of these subjects, however, the violations were minor. Altering a single choice by a single unit would eliminate all violations for all but 2 subjects, and moving two choices by 1 unit each would eliminate all but 1.[2] The only severe violation of revealed preference was subject 40. His choices are shown in Figure 1. For example, bundles *A* and *B* violate WARP, as do bundles *A* and *C*. However, bundles *B* and *C* also violate SARP and WARP.

[TABLE 2 HERE.]

[FIGURE 1 HERE.]

Given that virtually everyone's preferences can be rationalized, what do indifference curves look like? It turns out that 22 percent of subjects were perfectly selfish, keeping all of the endowment. Hence preferences of the form $U = x_s$ could characterize these people. Another 16 percent always split the payoffs exactly; hence Leontief utility, $U = \min\{x_s, x_o\}$, could generate this data. Finally, 6 percent always allocated tokens to maximize the total payoff of subjects; that is, they were social maximizers. $U = x_s + x_o$ could represent these people.

This covers over a third of all subjects. The rest of the subjects were similar to these extreme cases but didn't fit them exactly. Figure 2 illustrates an example of what we call a weakly selfish person—someone with a bit more price sensitivity than a strong free rider. Figure 3 is a weakly Leontief person. She splits payoffs evenly or nearly evenly on all budgets. Finally, Figure 5 is a weak social maximizer. This person has very flat (but not perfectly flat) indifference curves.

[FIGURE 2 HERE.]

[FIGURE 3 HERE.]

[FIGURE 4 HERE.]

---

[2] We also conducted more sophisticated analysis of violations, including applying Afriat's Critical Cost Efficiency Index, which is a measure of how costly a violation of revealed preference is for the subject. This analysis, which can be found in Andreoni and Miller (2002) yielded a similar interpretation, that is, only subject 40 had severe violations.

Overall we can characterize all individuals as one of these six types.[3] Table 3 shows the distribution of these preferences.

[TABLE 3 HERE.]

4. Conclusion

This work illustrates that not all "non-economic" behavior is beyond economic analysis. Our maintained assumption as economists is that individual behavior is consistent with self-interest. At its weakest, self-interest only means that choices conform to some underlying preference ordering that is complete, reflexive and transitive, and, hence, some utility function can be used to describe behavior. However, the assumption of self-interest does not tell us what variables are in that utility function. What does? Our methodology dictates that people themselves, through their actions, will do so. What we have shown here is that unselfish behavior in experiments can indeed be captured by a model of self-interested agents, but that self-interested agents are not always money-maximizing. When we define the choice set appropriately, unselfish acts can be described and predicted with the standard neoclassical model of choice.

---

[3] This characterization is based on minimizing the Euclidian distance between a subject's choices and those of one of the three exact utility functions, selfish, Leontief, and perfect substitutes. We also categorized people using a Baysian criterion and by using an adaptive search algorithm. Each produced similar results.

TABLE 1. This table lists the eight allocation decisions presented to the subjects. The decisions were presented in random order to each subject. Each decision can be thought of as choosing payoffs along a budget constraint. For each budget, a subject allocates the token endowment between himself and another subject. Tokens are redeemable for points at different rates for the two subjects. Hence, each allocation problem is a choice of final payoffs for the two subjects, where the price of the other's payoff in terms of self payoff varies across budgets. We can then check to see that there are no violations of revealed preference, e.g. GARP, in which case the data on the subject is consistent with rational choice.

| Budget	Token Endowment	Points received for each token allocated:	
		Self	Other
1	40	3	1
2	40	1	3
3	60	2	1
4	60	1	2
5	75	2	1
6	75	1	2
7	60	1	1
8	100	1	1

TABLE 2. Here we list the subjects who showed at least one violation of the revealed preference axioms. Of the 142 subjects, fewer that 10% had any violations. Of those violations, most would disappear if one choice on one budget were moved by one token. Hence, only subject 40 showed serious violations of revealed preference. This implies that virtually all subjects display a rational demand for altruism.

	Number of Violations		
Subject	WARP	SARP	GARP
3	1	3	3
38	2	4	4
40	3	10	10
41	1	1	1
47	1	1	3
61	1	4	4
72	1	1	1
87	1	1	1
90	1	1	1
104	1	2	2
126	1	3	1
137	1	1	1
139	1	1	1

TABLE 3. Given that preferences are rational, what utility functions could have generated the observed behavior? By employing several different search algorithms, we found that six categories of utility functions could best characterize the data. There were three utility functions that fit a large fraction of the data precisely. The three other categories of preferences had preferences similar to the exact utility function, but differed some, as is illustrated in Figures 2, 3, and 4.

| Utility Function	Exact Fit	Weak Fit
Selfish: $U_i = \pi_s$	31	31
Leontief: $U_i = \min \{\pi_s, \pi_o\}$	23	26
Social Maximizer: $U_i = \pi_s + \pi_o$	8	22
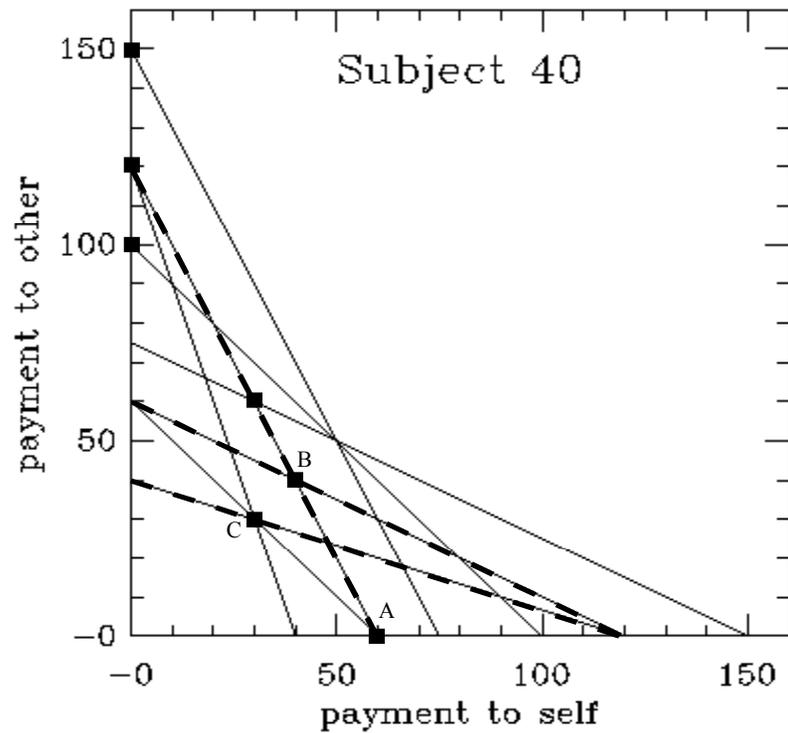
FIGURE 1: Of the subjects who had any violations of revealed preference axioms, only Subject 40, shown above, had severe violations that could not be eliminated with small adjustments to choices. For instance, A is directly revealed preferred to B, but B is directly revealed preferred to A. Likewise, A and C are directly revealed preferred to each other. Hence, A and B, and A and C violate WARP. Notice C is revealed preferred to B, but since B is directly revealed preferred to C, then B and C violate GARP as well. For this subject, no quasiconcave utility function could rationalize the data.
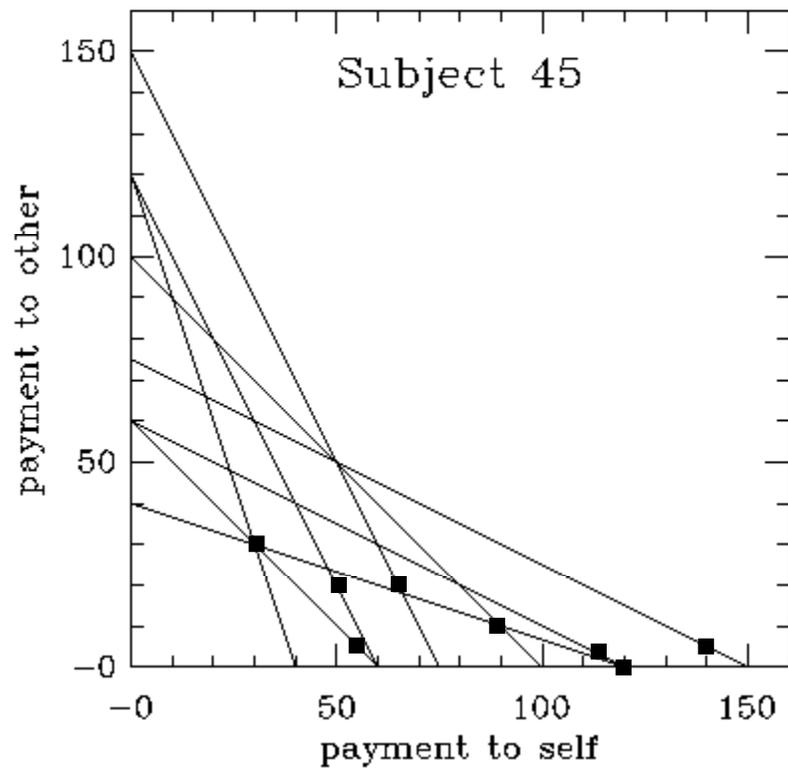
FIGURE 2: Over all 142 subjects, 22% were perfectly selfish, so $U = x_s$ could rationalize these choices. Another 22% were close to being perfectly selfish. Subject 45 shown here is typical of these. Most of the payoff is kept for himself, but there is still considerable price sensitivity.
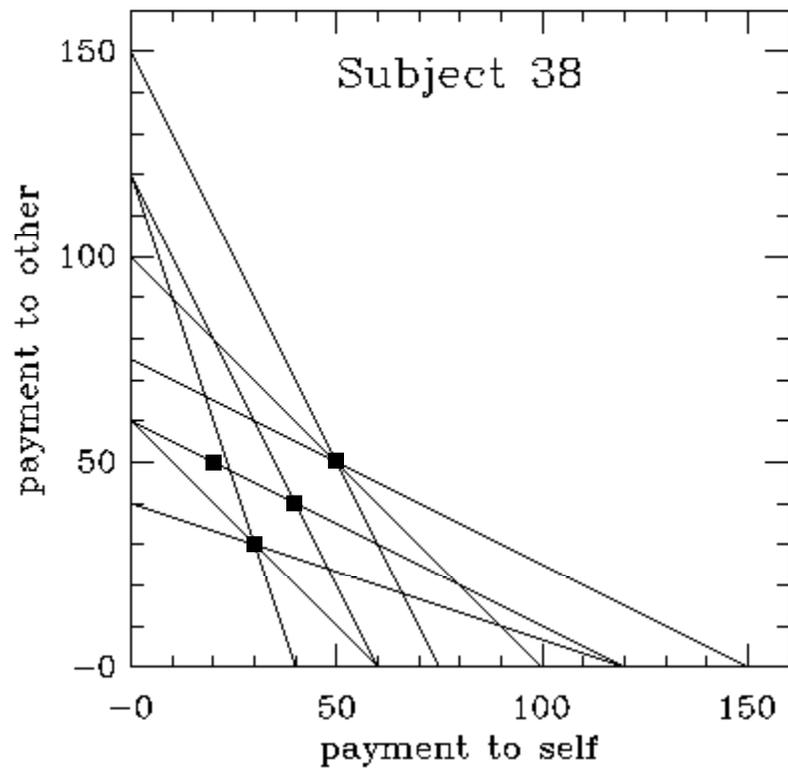
FIGURE 3: We found that 16% of subjects always divided the payoffs equally, hence $U = \min\{x_s, x_o\}$ could rationalize these choices. Another 18% had preferences that were very near to Leontief preferences but, like Subject 38 in this figure, deviated slightly from perfect Leontief Preferences. Note that Subject 38 also violated WARP. However, the deviation (which appears likely to be an honest error) is not severe—moving two of his choices by one token each would remove all violations.
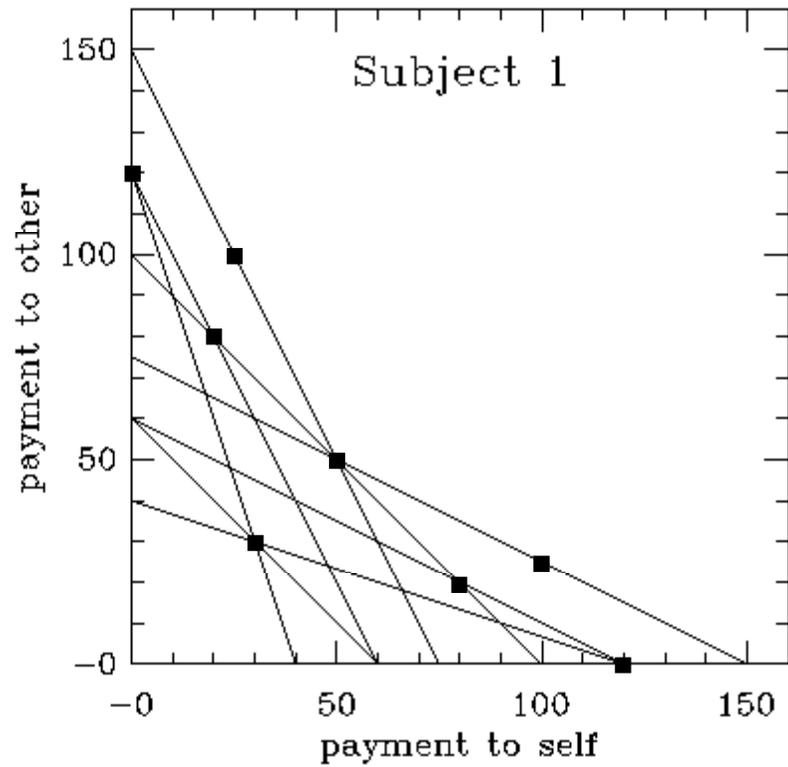
FIGURE 4: Six percent of subjects always gave all of the endowment to the subject with the highest redemption value, meaning the utility $U = x_s + x_o$ would rationalize this data. Another 15% of subjects had preferences like Subject 1 above. They appear to have very flat indifference curves that generate choices near the socially maximal payoffs.

References

Afriat, S. "The Construction of a Utility Function from Expenditure Data."
Econometrica, 6, 1967, 67-77.

Andreoni, James and Miller, John H., "Giving According to GARP: An Experimental
Test of the Consistency of Preferences for Altruism." *Econometrica*, 70, 2002, 737-
53.

Samuelson, Paul A, "Consumption Theory in Terms of Revealed Preference."
Econometrica, 15, 1948, 243-253.

Varian, Hal R., "The Nonparametric Approach to Demand Analysis." *Econometrica*, 50,
1982, 945-72.

Varian, Hal R., *Microeconomic Analysis, Third Edition*. New York: Norton, 1992.