Endogenous Social Preferences

JEFFREY P. CARPENTER

Department of Economics, Middlebury College, Middlebury, VT 05753, USA;
e-mail: jpc@middlebury.edu

Received June 4, 2002; accepted June 16, 2003

Abstract

A long-standing discussion in economics asks whether institutions affect people's social predispositions. The current experiment tests whether different aspects of markets affect people's social preferences. The results are that people are less socially minded in more anonymous settings. Additionally, market competition erodes social preferences through two mechanisms. First, market competition encourages opportunistic behavior, and second, the market institution itself decreases the other-regardingness of the participants.

JEL classification: *C72; C92; D64*

Keywords: *endogenous preferences; social preferences; value orientation; ultimatum game; best shot game*

It is almost a general rule that wherever manners are gentle there is commerce; and wherever there is commerce, manners are gentle.

—*The Spirit of Laws*, Montesquieu (1749/1989, vol. 2: 7)

The bourgeoisie . . . has left no other nexus between man and man than naked self-interest, than "callous payment." It has drowned out the most heavenly ecstasies of religious fervour, of chivalrous enthusiasm, of philistine sentimentalism, in the icy water of egotistical calculation.

—*The Communist Manifesto*, Marx and Engels (1847/1955: 12)

1. Introduction

The idea that economic institutions shape the preferences of individual agents has a long history in economic thought even if it has not had much impact on economic theory (Bowles 1998). Preferences that interact with institutions are a problem for theorists

Author's Note: The author thanks Sam Bowles, Corinna Noelke, Christophre Georges, Peter Matthews, Bruce Pietrykowski, and Gil Skillman for comments on an earlier draft and the National Science Foundation (SES-CAREER 0092953) for financial support.

Review of Radical Political Economics, Volume 37, No. 1, Winter 2005, 63-84
DOI: 10.1177/0486613404272330

because of mutual causation, but ignoring the endogeneity of preferences is an even larger problem for the validity of economic discourse. At a minimum, policies written as if changing the rules will have no effect on the attitudes of those people that have to live within the rules will lead to unintended consequences.

Concerning one institution in particular, Hirschman (1982) documents two competing theories of how markets affect the way individuals think of and treat each other in society. One theory, originating in the mid-eighteenth century and related to the first quote above, asserts that markets exert a civilizing influence over individuals, making them, for lack of a better term, nicer. Although the mechanisms by which markets civilize society often look suspiciously functionalist, the basic idea is that, in a society dominated by well-functioning markets and specialization, people are forced to interact because they can no longer individually produce all the things they need to survive. Therefore, a sort of folk theorem results: because individuals are forced to interact repeatedly, an equilibrium arises in which people need to be nice to each other to maintain trade relationships.

On the other hand, a second theory, often associated with Marx (the second quote), states that markets corrode societal values rather than instill them. Here people become nastier because "individual behavior . . . [is] increasingly directed to individual advantage, habits and instincts based on communal attitudes and objectives have lost out" (Hirsch 1976: 117-18). Implicitly, this view contends that markets make interactions more anonymous, not more personal, and this anonymity fosters the competitive behavior necessary for markets to work. The end result is that anonymity and competition drive wedges between individuals, diminishing their preferences to engage in collectively beneficial acts.

In trying to reconcile these points of view, one notices that there are two issues at stake: the effect of markets on the anonymity of interactions, which, in turn, affects people's attitudes toward each other; and the effect of market-induced competition on people's attitudes toward each other. This article measures the effects of economic institutions on people's *social preferences* defined as the way people rank different allocations of material payoffs to themselves and others (Camerer and Fehr 2001).[1] Specifically, with the long-standing debate about the effects of markets in mind, I conducted an experiment to assess the impacts of anonymity and competition on individuals' preferences for other people's well-being.

Before proceeding, it is worth noting two previous experiments that partially provide a foundation for the current experiment in the existing literature. Hoffman et al. (1994) examine the effect on bargaining outcomes of framing interactions as markets. Specifically, the experimenters changed the wording of the instructions from the standard context-free reference to "persons A and B" to one that labeled players as buyers or sellers. The results show that simply framing an interaction as a market has a significant effect on the distribution of the surplus. Sellers offer much less of the surplus to buyers than when the bargainers are called persons A and B. In this case, it appears that markets trigger more egoistic behav-

1. The tradition in economic theory has been to assume agents who only care about their own well-being. However, there is now an overwhelming amount of evidence that real decision makers act altruistically, cooperatively, reciprocally, and even spitefully towards each other. These deviations from *homo economicus* have been termed social preferences. See Fehr and Schmidt (1999), Bolton and Ockenfels (1999), Falk and Fischbacher (1998), or Charness and Rabin (2002) for an introduction.

ior in people, which might also suggest that social preferences deteriorate in markets. Considering the effect of competition, Schotter, Weiss, and Zapater (1996) show that introducing competition also reduces offers in the same bargaining game. In this experiment, first-movers had to compete to survive to a second round of play. Comparing the first round offers of the survival treatment with a control and first round offers with second round offers of those who survived, the authors provide marginally significant evidence that competition also appears to make participants behave more egoistically.

One is tempted to conclude that these two studies illustrate how markets erode participants' other-regardingness. However, we need to be careful here. Actually, these studies *only* show that offers fall when interactions are framed as markets or when competition is allowed; we do not know, however, whether participants' preferences have changed. In fact, offers in the ultimatum game may be particularly bad measures of social preferences because it is well known that egoistic first-movers tend to balance payoffs against the subjective probability of rejection (Forsythe et al. 1994; Hoffman, McCabe, and Smith 1996; Carpenter 2002b). That is, the same egoistic first-mover in the standard game who offers half because she thinks low offers will be rejected may offer a lower amount when the interaction is framed as a market or when competition works in her favor because her estimate of the likelihood of being rejected is lower.[2] The point is that ultimatum offers may change without social preferences changing.[3]

To assess whether social preferences are affected by factors associated with markets, I discuss an experiment that measures participants' social preferences both before and after interacting with other participants using a method that is incentive consistent and provides a measure of preference strength. Economic experiments are *incentive consistent* because people are paid based on the actions they (and the other participants) choose, and choosing actions that result in more money is always in one's self-interest.

Briefly, the results suggest that reducing anonymity does make people *more social* in the way theorized by the folk theorem: reducing anonymity (i.e., repeating interactions) reduces people's ability to engage in opportunistic acts such as taking advantage of being in a more powerful bargaining position because those in less powerful positions will react spitefully. More important, market competition erodes social preferences through two mechanisms. First, market competition encourages opportunistic behavior, creating a less friendly atmosphere; and second, controlling for the first effect, the market institution, perhaps because of its framing effects, itself decreases the other-regardingness of the participants.

---

2. A pertinent example of this comes from Barr (2003), who, in a field experiment, shows that resettled villagers in Zimbabwe with much more access to markets make higher offers in the ultimatum game than villagers who have not been resettled. One might take this as evidence of the socializing effect of markets, but Barr shows that another reason for this behavior is that offers increase in settlements when players are more uncertain about what will be an acceptable offer to their counterparts.

3. Schotter, Weiss, and Zapater (1996) also ran dictator games (the same as the ultimatum game except that second movers cannot reject an offer; they just get whatever they are given) using the same setup and found that transfers to a second party drop significantly when competition is introduced. However, without pretesting subjects' social preferences, these results only provide between-subject comparisons. One important strength of the current experiment, discussed below, is that it provides much stronger within-subject comparisons.

2. The Experiment[4]

The experiment was designed to test, first, the hypothesis that preferences are endogenous. That is, do economic institutions affect people's social preferences? Second, if the endogenous preferences hypothesis holds, do specific aspects traditionally associated with large markets—the absence of repeated interaction (as a proxy for anonymity), material incentives, and competition—erode or instill goodwill among individuals? The exercises used to elicit people's social preferences are discussed, and then five treatments are described which were used to assess the impact of anonymity, the incentive to treat one's partner nicely, and the impact of competition on people's social orientations.

The same two preference revelation mechanisms were used for each treatment, and they were always presented in the same order. A *preference revelation mechanism* is an allocation rule for which truthfully revealing one's preferences is incentive consistent. In the specific context, this reduces to decision tasks that monetarily reward self-interest. Furthermore, in these tasks social preferences (e.g., altruism, or cooperativeness) reduce one's payoff, implying that only people with social preferences will incur the costs. The first mechanism pretested players' social orientations and the second tested for changes in player's social preferences during each of the five treatments. Both measures are based on a series of dictatorial choices over the division of a monetary pie. In the standard dictator game (Forsythe et al. 1994), one player, the "dictator," is given the choice of how to divide a sum of money between herself and another anonymous participant. This choice is made with impunity because players are anonymous and the recipient has no "veto" power over the dictator's choice. Given this structure, the amount the dictator transfers to the recipient is a measure of her social preferences or other-regardingness.

Because it is important to be careful when measuring preferences, the simple dictator game was not used as it provides only one observation per participant. Instead, both measures categorize participants by how other-regarding their responses to a series of dictator choices were. Using this method not only gets more than one observation per participant but can also construct a measure of how consistent participants' social preferences are. The social orientation exercise developed in Griesinger and Livingston (1973) was used as a pretest, and the posttest was the dictator GARP (generalized axioms of revealed preference) exercise developed in Andreoni and Miller (2002) (see below).[5]

The social value orientation (VO) exercise was first used in social psychological research (Shure and Meeker 1967; Liebrand 1984; McClintock and Liebrand 1988; Kramer and Goldman 1995), but it has now been adopted by economists (Carpenter 2002b; Buckley et al. 2001). In the VO exercise, participants make binary dictator choices over combinations of *own* and *other* monetary payoffs. Own amounts are kept by the dictator and other amounts are given to another anonymous participant. Because measures of social preferences are needed for everyone, the participants were matched in groups of three for this exercise. The reason for this triadic design was to eliminate any strategic thinking among

4. The experimental instructions are available from the author upon request (jpc@middlebury.edu).

5. It was also important to not use the same mechanism to both pretest and posttest preferences because it would be too obvious to participants what was happening. Furthermore, both exercises are based on a series of dictator choices, which provides the basis for a natural interpretation of any differences between the pretest and posttests.

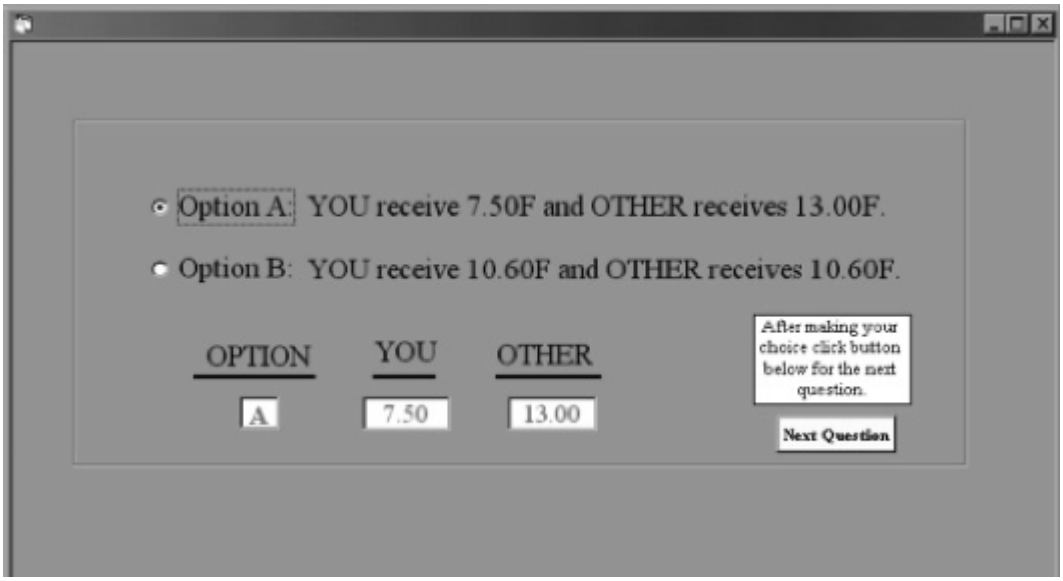

Figure 1.
Screen Shot of Value Orientation Exercise.

the participants. That is, one dictator's transfers were sent to another dictator, who sent her transfers to a third dictator, and the third dictator completed the circuit by sending to the first. This way, there was no reason for individuals to think about, or anticipate, reciprocity between themselves and another participant to whom they sent money and from whom they received money. In this way, altruism was captured without any confounding effects of reciprocity.

The participants received no feedback about how much they were sent until the end of the entire experiment. In total, participants made twenty-four decisions and their payoffs were the sum of the twenty-four amounts kept plus the sum of the twenty-four amounts sent by another dictator. This is obviously an incentive consistent way of eliciting social preferences because it is always costly for the dictator to transfer money to the recipient and, given the anonymity of interactions, there is no possible material benefit from doing so.

An example of the choice problem is illustrative. Figure 1 presents one of the twenty-four choices each participant made. They chose either option A or option B, and the order in which the payoffs were presented was randomized each time (i.e., it was either *own, other* or *other, own*) so that players could not just focus on their own payoffs; at a minimum they needed to look at the consequences of their choices for the recipient. Payoffs were listed in terms of experimental francs and then translated into dollars (the exchange rate was $1 = 5F) at the end of the experiment.

The sum of the payoffs is not constant across options. This is an asset of the VO because it implies that the cost of giving is not constant. Specifically, the twenty-four outcome pairs lie evenly spaced on a circle with radius of 15 experimental francs, and each choice was between two contiguous options on this circle. The center of the circle is the origin of the two-dimensional space where the horizontal axis measures own francs and the vertical axis measures other francs.

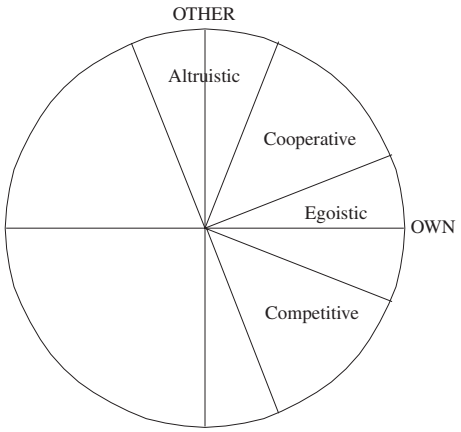

Figure 2.
Social Value Orientation Diagram.

The *motivational vector* for the VO exercise is calculated by adding all the participant's responses. This vector is then mapped back onto the original circle and is used to place the subject into one of four categories based on how much she kept and how much she transferred to the recipient. Figure 2 shows the VO circle divided into four classifications, from most social to least: altruistic, cooperative, egoistic, and competitive.[6]

Another benefit of the VO approach is that the length of the motivation vector measures the consistency of each subject's choice pattern. If, for example, an individual's motivation vector were calculated to be (15,0) and the individual chose (7.5,13) over (3.9,14.5), then the individual would have made a consistent choice because she picked the choice closest to her final vector. Subjects who choose randomly will have very short motivation vectors, and subjects whose behavior is completely consistent will have vectors twice as long as the circle radius. The measure of consistency will be each subject's vector length as a fraction of the maximal length.

The second preference mechanism, completed after participating in one of the treatments described below, I call the GARP mechanism because it was developed by Andreoni and Miller (2002) to test the extent to which social preferences adhere to revealed preference axioms. In the GARP exercise, participants make eight dictator choices about how to divide a variable number of tokens that have differing values to the dictator and the recipient. Therefore, as with the VO exercise, the GARP mechanism alters the relative price of giving. Figure 3 shows the screen used to collect participant responses in the GARP phase of the experiment. As one can see, there were four possible "pie" sizes: 40, 60, 75, and 100 tokens for each decision, and the relative price of giving was respectively 1/3, 1/2, 1, 2, 3. Eight budget constraints were formed by different combinations of pie sizes and relative prices from which participants made their choices.[7]

---

6. More specifically, according to the convention in the psychological literature, motivational vectors that fall in the range of 112.5 to 67.5 degrees are classified as altruistic, between 67.5 and 22.5 as cooperative, between 22.5 and –22.5 as egoistic, and between –22.5 and –67.5 as competitive.

7. Notice that in the lower left of the screen is a simple calculator that was provided for the participants to use. This matches the protocol of Andreoni and Miller (2002).

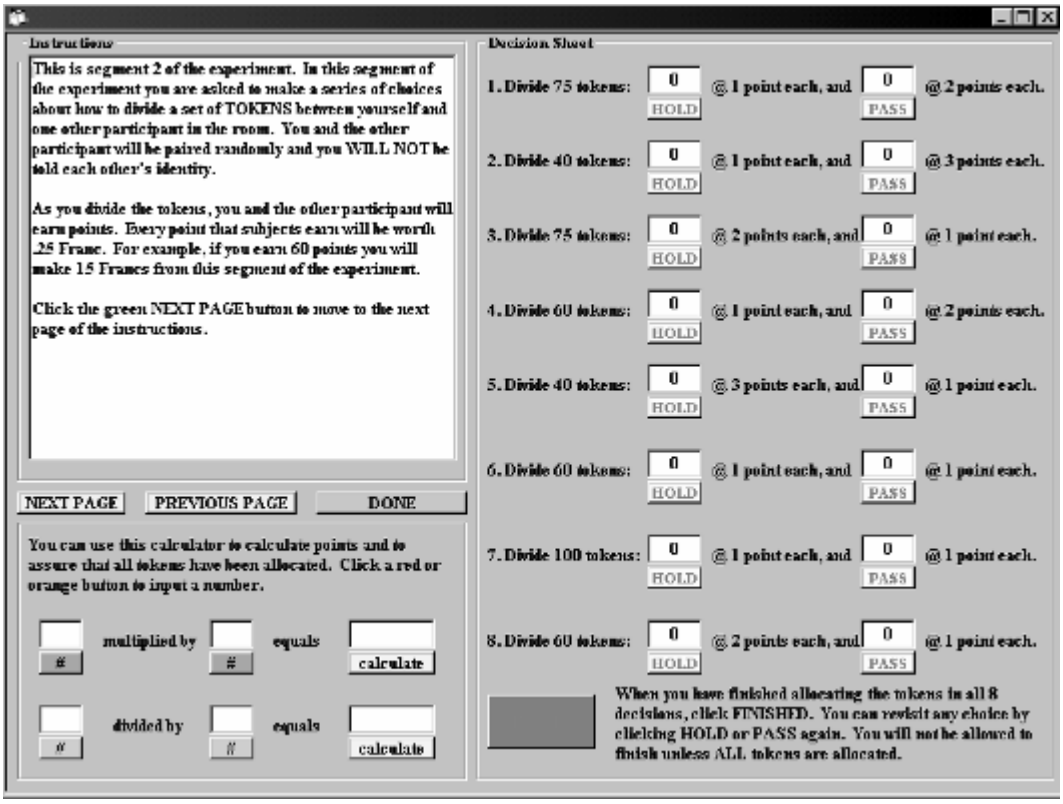

Figure 3.
The GARP (Generalized Axioms of Revealed Preference) Decision Screen.

Like the VO, players in the GARP exercise are randomly organized in triads to eliminate any opportunity for strategic behavior. Furthermore, the exercise was not completed until participants had filled in each of the "Hold" and "Pass" input boxes where they typed amounts that they would keep for themselves and amounts they wanted to pass to their partner. However, they could always change any decision before finally submitting the entire series.

By simply minimizing the distance between player choices and three models of play in this exercise, players can be categorized as altruistic, cooperative, or egoistic as in the VO exercise. Model altruists are those whose preferences for their own payoff and the payoff of the other player are substitutes because, for a given price ratio, they assign all the tokens to whomever benefits the most. Model cooperative players are those who exhibit Leontief preferences in that they equalize payoffs regardless of the pie size or relative price (i.e., they value fair outcomes). Finally, model egoists keep all of the pie for themselves and do not react to the relative cost of giving or the size of the pie. Notice that the mapping from behavior to preference categories is identical to the VO: those with high other scores will also exhibit substitutes as preferences, those who are cooperative in the VO will have Leontief preferences in the GARP exercise, and so on.[8]

8. However, the mapping between the value orientation and GARP (generalized axioms of revealed preference) is not perfect because there is no equivalent of a competitor in the GARP framework. The simple reason is

The discussion of the preference revelation mechanisms can be summarized by reiterating the similarities between the two exercises. First, both mechanisms are based on the dictator game, which economists have now come to understand as a way to elicit social preferences (Carpenter 2002c; Camerer and Fehr 2001). Second, both games involve a series of dictator choices rather than a single choice which means the consistency of preferences can be assessed. Third, both measures test the robustness of social preferences to changes in the size of the pie and the relative price of giving. Finally, there is a clear relationship between VO categories and GARP categories, which makes their comparison meaningful.

Now I redirect attention to the five treatments that participants took part in. The first treatment was a control to calibrate the relationship between VO categories and GARP categories. In this treatment, thirty-six subjects participated in only the VO and GARP exercises, with the VO immediately preceding the GARP exercise. The four other treatments were composed of inserting a specific game that varied the anonymity of interactions, the off-equilibrium incentives (i.e., which actions should one follow to increase her payoff when not at an equilibrium), or the level of competition between the VO and the GARP exercises.[9] Specifically, while the control consisted of participants playing VO then GARP, each of the four treatments consisted of playing VO, then playing one of the games described below, then playing GARP. Preference changes were tested for by comparing GARP behavior to VO behavior in each treatment against any control differences in the two exercises.

Two of the treatments consisted of ultimatum bargaining. In the *ultimatum game* (Gueth, Schmittberger, and Schwarze 1982), a first-mover is provisionally allocated some pie, say $10, and is asked to propose a distribution of the pie between herself and a second-mover. The second-mover is then asked to accept or reject the proposed distribution. Accepted distributions are implemented, but if the second-mover rejects, neither player gets anything. The standard solution concept for this sort of sequential game is subgame perfection. Essentially, a *subgame perfect* equilibrium is a strategy profile that does not rely on empty threats. For example, in the ultimatum game the second-mover could threaten to reject low offers, but such a threat is incredible because it is never in the second-mover's material interest to carry through on the threat because she would lose money. Because the first-mover realizes that it is in the interest of the second-mover to accept any offer, she will never offer the second-mover any more than the smallest possible amount. Hence, the subgame perfect equilibrium occurs where the first-mover demands the entire pie minus some small amount and the second-mover accepts because the small amount is better than nothing.

In the *same bargaining* treatment, twenty-four participants were assigned to an unchanging role and played the ultimatum game for ten periods with the same partner. In the *random bargaining* treatment, thirty-two participants played the ultimatum game for ten periods, but they were randomly repaired at the beginning of each round. In the *best shot* treatment, twenty-four participants played the best shot game (defined below) for ten rounds,

that while the value orientation asks players to make decisions in both gain space and loss space, GARP only asks about allocating gains. However, this problem is small because few players end up being categorized as competitors.

9. Three of the four games were the same games used by Prasnikar and Roth (1992) to understand how off-equilibrium incentives affect the expression of fairness in games.

Table 1
Best Shot Game Payoff Table

	Redemption Values		Expenditure Values	
Project Level (Units)	Redemption Value of Specific Units	Total Redemption Values of All Units	Number of Units You Provide	Cost of the Units You Provide
0	0.00	0.00	0	0.00
1	1.00	1.00	1	0.82
2	0.95	1.95	2	1.64
3	0.90	2.85	3	2.46
4	0.85	3.70	4	3.28
5	0.80	4.50	5	4.10
6	0.75	5.25	6	4.92
7	0.70	5.95	7	5.74
8	0.65	6.60	8	6.56
9	0.60	7.20	9	7.38
10	0.55	7.75	10	8.20

and at the beginning of each round they were randomly repaired. Finally, in the *market* treatment, twenty-four subjects participated in a ten-period market.

The best shot game (Harrison and Hirshleifer 1989) and the specific market mechanism (Prasnikar and Roth 1992) must also be described. Players of the best shot game are assigned to be either the first-mover or the second-mover and are then presented with Table 1. The best shot game concerns the provision of a public good where the provision level is determined by the highest individual contribution (i.e., the best shot). If the first-mover provides a level of q_1 and the second-mover a level of q_2, the benefits of the project are determined by $\max\{q_1, q_2\}$. Providing for the public good is costly. Specifically, the cost of providing is symmetric with respect to the players payoffs and linear in the chosen level, $0.82q_i$ for $i = 1,2$, where the costs are measured in experimental francs.

The first-mover has the advantage in this game because identical or lower contributions are wasted, which means she can force the second-mover to contribute by giving nothing. The subgame perfect equilibrium occurs when the first-mover provides a project level $q_1 = 0$ and the second-mover maximizes her payoff by providing $q_2 = 4$. This results in a payoff of 3.70F for the first-mover and 0.42F for the second mover.[10] There is an interesting difference between the best shot game and the ultimatum game that provides the reason for including it as a treatment in the experiment. Notice, as first-movers in the ultimatum game become more social (i.e., offer more), second-movers respond by being more likely to accept offers (Camerer and Thaler 1995; Roth 1995). However, because contributions below the maximum are redundant in the best shot game, as first-movers become more social and increase their level of provision, second-movers have a strong incentive to free ride (Harrison and Hirshleifer 1989; Prasnikar and Roth 1992). Therefore, off the equilibrium

10. I used the "full information" best shot game in which players know they are all using the same payoff table because relative comparisons may be an essential trigger of preference changes and because the other games are full information.

path, first-movers are rewarded for being generous in the ultimatum game, but penalized in the best shot game. Does this structural change affect participants' social preferences?

To induce competition, the ultimatum game was also implemented as a market. In each market, there was one seller and four buyers who bid on an indivisible good.[11] The good costs nothing for the seller to provide and each buyer was allocated a maximum willingness to pay of ten experimental francs. In each period, the four buyers submitted bids simultaneously and then the highest bid was presented to the seller who accepted or rejected the offer (this was similar to Prasnikar and Roth [1992] and Schotter, Weiss, and Zapater [1996]). If there were two equally high bids, the good went to a buyer randomly. Because of the excess demand in the market, the equilibrium is the same as the ultimatum game since one player (in this case the seller) receives most of the pie.

Notice that the four games have very similar perfect equilibrium predictions (i.e., one person earns much more than the other), but the games differ along the dimensions of interest. The same bargaining treatment makes interactions less anonymous compared to the other three games, the best shot game changes the incentive to reward social behavior when compared to the ultimatum game, and the market game introduces competition.

It is useful to conclude this section by making the priors clear about the effect of game structure on preferences. Returning to the "markets make nicer people" hypothesis, repeated interactions would be expected to foster more friendly preferences, and in accordance with the "markets make nastier people" hypothesis, we expect random (seemingly one-shot) interactions to make people more egoistic. Furthermore, the reduced incentive to be social in the best shot game would be expected to affect people's social preferences when compared to the random ultimatum game, that is, people will become more egoistic as a result of the best shot game. Finally, competition would be expected to make people more egoistic when compared to random bilateral negotiations. If the same treatment is S, the random treatment is R, the best shot treatment is B, and the market treatment is M, where $>$ means "elicits more social attitudes," then $S > R$, $R > M$, $R > B$ would be expected, and by transitivity $S > M$ and $S > B$.

3. The Results

The results of the experiment will be discussed in the following order. First, tabulations of the two preference elicitation mechanisms are presented without discussing the relationship between the two. Next, the results of the four games participants played are briefly discussed. Last, the main results are presented respecting the endogeneity of social preferences by analyzing the relationship between the treatments and the participants' sociality.

Table 2 summarizes the social preferences data. There are five tabulation tables (one for each treatment) in which the fraction of participants who fall into a paired (VO, GARP) category are calculated. The various categories are represented as follows: 0 = competitive, 1 = egoistic, 2 = cooperative, 3 = altruistic. For example, in the control treatment, 39 percent of the participants were categorized as egoistic by both the VO and GARP exercises.

11. Each market consists of five participants. Participants keep the same role, but, as in the best shot treatment and the random bargaining treatment, players are reshuffled into new markets at the beginning of each round.

Table 2
Tabulation of Preference Revelation Results

		GARP						GARP			
Control ($n = 31$)	$C = 0.47$, $p < .04$	1	2	3	Total	Best Shot ($n = 20$)	$C = 0.43$, $p < .12$	1	2	3	Total
VO	0	0	0.06	0	0.06	VO	0	0.10	0.05	0	0.15
	1	0.39	0.13	0	0.52		1	0.45	0.05	0	0.50
	2	0.13	0.13	0.13	0.39		2	0.10	0.20	0.05	0.35
	3	0.03	0	0	0.03		3	0	0	0	0
	Total	0.55	0.32	0.13	1		Total	0.65	0.30	0.05	1
		GARP						GARP			
Bargain (Same) ($n = 19$)	$C = 0.27$, $p < .50$	1	2	3	Total	Market ($n = 20$)	$C = 0.27$, $p < .60$	1	2	3	Total
VO	0	0.05	0.11	0	0.16		0	0.05	0	0	0.05
	1	0.47	0.21	0	0.68		1	0.30	0.10	0	0.40
	2	0.11	0.05	0	0.16		2	0.30	0.15	0.10	0.55
	3	0	0	0	0		3	0	0	0	0
	Total	0.63	0.37	0	1		Total	0.65	0.25	0.10	1
		GARP									
Bargain (Random) ($n = 27$)	$C = 0.37$, $p < .15$	1	2	3	Total	Treatment	Tested Direction	Sign Test p Value			
VO	0	0.11	0.04	0.04	0.19	Control	More Social	.27			
	1	0.37	0.07	0	0.44	Bargain (Same)	More Social	.03			
	2	0.11	0.19	0.07	0.37	Bargain (Random)	More Social	.23			
	3	0	0	0	0	Best Shot	More Asocial	.31			
	Total	0.59	0.30	0.11	1	Market	More Asocial	.03			

Note: VO is value orientation preference pretest; GARP (generalized axioms of revealed preference) is revealed preferences posttest. Category 0 = competitive, 1 = egoistic, 2 = cooperative, and 3 = altruistic. C is Cramer's coefficient of association; p is p value on C; and n is number of participants. Tabulated numbers are frequencies of participants based on VO type and GARP type. The Sign Test tests whether preferences change significantly in the hypothesized direction.

The rightmost column in each tabulation reports the distribution of VO types and the bottommost column lists the distribution of types in the GARP exercise. Using pairwise Kolmogorov-Smirnov tests, no statistical difference was found between any of the five VO distributions at the 5 percent level, indicating each treatment began with a similar distribution of types. In general, most participants were classified as egoistic by both preference measures, but there are a significant number of social types (cooperators and altruists) in each treatment. Furthermore, judging by the fraction of participants who are categorized off-diagonally, it appears that the treatments affected participants' preferences; however, the analysis of preference endogeneity will be postponed until after discussing the possible causes of changed behavior in the treatments.

Before moving on, the reader should notice differences in the number of subjects reported per treatment in Table 2 and the corresponding numbers mentioned in section 2. In each case the number of observations is lower in Table 2. As stated above, one of the strengths of using the VO and GARP methods is that the consistency of each player's choices can be assessed. Observations were culled in each treatment when a player demonstrated choice consistency in the VO lower than 60 percent of the maximum.[12] As in many experiments, despite being paid, some subjects do not pay attention to the experiment or are confused. Culling based on choice consistency allows evaluation and elimination of this noise from the data.[13]

A summary of behavior in the four games is presented in Table 3. The second and third columns report the mean offers to the second-mover in the two ultimatum bargaining treatments, the fourth column lists the mean project level chosen by the first-mover in the best shot game (q_1), and the fifth column lists the mean buyer bid in the market treatment. As one can see, the two bargaining treatments elicit similar behavior: on average and across periods the first-mover offers between 42 and 46 percent of the pie to the second-mover. However, pooled mean offers are statistically lower in the random treatment ($z = 2.17$, $p < .03$), and the rejection rate is much higher. These differences account for the difference in average payoffs between treatments and suggest that, as anticipated, increased anonymity in the random treatment creates a less friendly bargaining environment.

The best shot results largely replicate Carpenter (2002a) in that first-movers reduce their contributions, but they never quite reduce their contributions over time to zero. However, the important comparison is between first-mover behavior in the best shot game and the random bargaining game. Because first-movers in the best shot game reduce their contributions over time, while first-movers in random bargaining increase their offers over time, two very different atmospheres develop in these games. Although both games main-

12. Eight percent of the sample was culled. The 60 percent cutoff rule in the value orientation was chosen based on the bimodality of the data, that is, almost all players either recorded levels below 60 percent or well above the 84 percent overall average. Furthermore, the average consistency of random choice in the value orientation (based on a simulation run one hundred times) is 53 percent, which suggests that using the natural break point of 60 percent is reasonable. Admittedly, the results of the uncontrolled statistical tests done in Table 2 are less impressive when done on the unculled sample, but the more important hypothesis tests using regressions in Table 5 are largely unaffected by culling.

13. Readers might worry that culling observations with low consistency may remove people who are just indifferent between many of the value orientation pairs. If people are equally likely to choose either option when indifferent, those who were culled would be expected to be disproportionately egoists; however, slightly less than half of those culled are egoists while egoists account for more than half the observations in the full sample.

Table 3
Mean First-Mover Choices (Standard Deviations)

Period	Same Bargaining (0 ≤ Offer ≤ 10)	Random Bargaining (0 ≤ Contribution ≤ 10)	Best Shot Game Period (0 ≤ Bid ≤ 10)	Market (0 ≤ Bid ≤ 10)
1	4.83 (1.40)	4.02 (1.09)	4.56 (2.88)	7.58 (1.47)
2	4.79 (0.33)	4.21 (1.10)	2.78 (1.39)	7.73 (2.05)
3	4.55 (0.97)	4.29 (0.84)	2.44 (1.88)	8.47 (1.08)
4	4.69 (0.72)	4.28 (0.82)	3.33 (2.65)	8.22 (1.98)
5	4.55 (0.89)	4.18 (0.94)	1.22 (1.39)	8.25 (2.39)
6	4.80 (0.34)	4.36 (0.81)	2.11 (2.26)	9.07 (1.22)
7	4.46 (0.89)	4.18 (1.07)	1.22 (1.48)	9.09 (1.25)
8	4.80 (0.40)	4.22 (0.78)	1.89 (2.71)	9.40 (0.79)
9	4.64 (0.90)	4.38 (1.02)	1.44 (3.24)	9.23 (1.66)
10	4.06 (1.11)	4.54 (0.57)	2.11 (2.93)	9.65 (0.54)
Overall mean	4.62 (0.87)	4.27 (0.89)	2.31 (2.46)	8.67 (1.65)
Rejection rate	.09	.20	.19/.32	.64
Mean payoff	$14.34 ($2.71)	$13.24 ($2.53)	$10.45 ($3.22)	$11.40 ($4.54)

tain the same degree of anonymity, differences in the off-equilibrium incentives of the two games affect reactions to friendly acts. Proposers in the bargaining treatment are rewarded for increasing offers while first-movers in the best shot game are taken advantage of when increasing their contributions.[14] To illustrate this point with the data, the first rejection rate reported for the best shot game (19 percent) records the frequency at which first-movers who contribute positive amounts are matched with second-movers who contribute nothing. The second rejection rate indicates that in 32 percent of the interactions in which the first-mover chose zero, the second-mover also chose zero. This statistic indicates that, compared to the random bargaining treatment, the number of interactions in which both players received zero payoffs is much higher in the best shot game.[15]

The market data illustrate the fairness-dampening effect of competition. The average first period bid is well above the five franc bid that equalizes the surplus between the winning buyer and the seller, and by the last period bids are close to the perfect equilibrium prediction, but never quite reach it according to a one-tailed Wilcoxon nonparametric test, $z = 3.64$, $p < .01$. The rejection rate reported here is the average likelihood that a buyer's bid was not the winning bid. As one can see, buyers must have been frustrated by competition because the average bid leaves them with a small fraction of the surplus and there was a 64 percent chance that they would not even get this amount. Compared to the random bargaining treatment, market outcomes are much closer to the theoretical prediction, and much less fair. Now how anonymity, off-equilibrium incentives, and competition affect social preferences is analyzed.

Returning to Table 2, the degree to which the different treatments affect people's social preferences can be assessed. Begin with the reference point: the relationship between the VO and the GARP exercise in the control treatment. In the upper left of each treatment cell,

14. As hoped, the best shot and random results more or less replicate Prasnikar and Roth (1992).
15. The high rejection rate is also reflected in the low average payoff listed in Table 3.

Table 4
Do the Treatments Make Participants More Social, Less Social, or Do They Reinforce Initial Preferences?

Treatment Ranking	Make More Asocial	Make More Social	Reinforce Preferences	Cramer's Coefficient (*C*)
1	Market[a]	Same,[a] Control	Best Shot	Control[b]
2	Control	Same, Control	Random	Best Shot
3	Same, Random	Random	Same	Random
4	Same, Random	Market	Control	Same, Market
5	Best Shot	Best Shot	Market	Same, Market

Note: This table puts the treatments in order according to the four criteria at the head of each column. Treatments on the same row indicate ties.
a. Implies result is confirmed by the Sign Test at the 5% level or better.
b. Implies Cramer's coefficient of association is significant at the 5% level or better.

Cramer's coefficient of association for categorical data and the significance level of the association (see chap. 9 of Siegel and Castellan 1988) are reported.[16] As hoped, the control treatment elicits the highest degree of association between the two preference measures ($C = 0.47$) and is the only treatment in which the association is significant ($p < .04$).

To be conservative, for the remainder of the analysis based on Table 2, paired VO scores of 0 and GARP scores of 1 are treated as unchanged preferences because there is no competitive category in the GARP exercise. Although the control condition exhibits a high and significant association between the two measures, nearly half the participants change their social orientation. This indicates that either people's preferences are highly volatile or the framing of the two exercises affects preferences. Regardless, what is important for the current discussion is how the other treatment tabulations compare to the control and to the random bargaining treatments.

Table 4 summarizes the endogenous preferences results. As a first pass, the frequency will just be calculated with which players became more asocial (less other-regarding), more social (i.e., more other-regarding), or had their VO preferences reinforced. Recall that the proper reference point for the market, same, and best shot treatments is the random treatment because the market treatment only differs with respect to the existence of competition, the same-bargaining treatment differs only with respect to the absence of anonymity, and the best shot treatment only differs with respect to the incentives for prosocial acts.

The second column of Table 4 shows that the market institution that generates competition and the resulting payoff asymmetries creates more asocial players compared to the random treatment. The third column illustrates that the evolution of a strong sharing rule (recall Table 3) in the less anonymous same treatment extends past the bargaining stage of the experiment and makes players more other-regarding. Another interesting result is that the best shot game, despite reducing the incentive to be social as a first-mover, largely reinforces players' VO preferences (fourth column). That is, much more than any other treatment, players in the best shot game report the same preferences in both preference exercises.

16. Effectively, Cramer's coefficient is a measure of correlation for categorical data.

Table 5
The Determinants of Preference Changes

Dependent Variable = Increase in Other-Regardingness (GARP [Generalized Axioms of Revealed Preference] Score Minus Value Orientation [VO] Score)			
	1	2	3
Value orientation	–1.97*** (0.33)	–2.04*** (0.36)	–2.20*** (0.38)
Role	0.53* (0.29)	0.76 (0.49)	2.82*** (0.83)
Payoff	—	–0.06 (0.05)	–0.10* (0.06)
Non homo economicus	—	–1.30** (0.59)	–3.35** (1.70)
Dissatisfaction	—	–0.83 (0.84)	1.42 (1.03)
Role × *Non homo economicus*	—	—	1.44 (1.77)
Role × Dissatisfaction	—	—	–8.70*** (2.51)
Same bargaining	–2.92*** (0.84)	–1.58 (1.03)	–0.96 (1.02)
Market	–2.36*** (0.86)	–2.72*** (0.91)	–3.21*** (0.85)
Best shot	–2.35*** (0.89)	–2.44*** (0.89)	–1.41 (0.95)
Sex	–0.14 (0.36)	–0.16 (0.40)	–0.55 (0.42)
Foreign born	–0.03 (0.30)	–0.15 (0.29)	–0.42 (0.35)
n	79	79	79
Pseudo R^2	.40	.44	.52
Wald χ^2 (*p* value)	59.44 (<.01)	55.48 (<.01)	65.20 (<.01)

Note: Ordered probits including session fixed effects. Robust standard errors in parentheses.
*Significant at the .10 level. **Significant at the .05 level. ***Significant at the .01 level.

The fifth column of Table 4 lists the treatments in decreasing order of the association between the two preference mechanisms. Notice, the best shot treatment generates the highest association (second only to the control), reinforcing the stability of preferences in this treatment, and the same market treatments exhibit equally low measures of association because less anonymous bargaining generates more social players and competitive markets generate more asocial players. Additionally, using the Sign Test for matched samples (Siegel and Castellan 1988: chap. 5), it is possible to test whether any of the treatments cause significant changes in players' social preferences.[17] The lower right corner of Table 2 reports the hypothesized direction of change and the significance of change in the five treatments. As the reader can see, the two major results are corroborated by Sign Tests: repeated, less anonymous bargaining makes players more social while competitive markets make players more asocial.

To supplement the treatment-level analysis, I dig deeper by decomposing the treatment differences into the individual determinants of preference changes.[18] Table 5 reports ordered probit results where the dependent variable is the difference between player GARP scores and VO scores. Here, positive dependent variables indicate players became more

17. To conduct the Sign Test, competitive and egoistic players are pooled as asocial and cooperative, and altruistic players as social, to test whether players in a treatment are equally likely to become more or less social.

18. Another interesting exercise would be to analyze the degree to which value orientations predict behavior in the treatments. However, such an analysis would obscure the purpose of this article. The interested reader should see Carpenter (2002b) for such an analysis using a similar bargaining environment.

social, and negative values indicate players became more asocial. To bias the results against the findings, the seven players who were categorized as competitive in the VO exercise and egoistic in the GARP exercise were dropped from the analysis because they would show up as becoming more social in this analysis.[19]

Now, the regressors are defined. *Value orientation* is a participant's categorization (0, 1, 2, or 3) from the preference pretest. *Role* divides players between those who have power (i.e., first-movers and sellers = 1) and those who do not (second-movers and buyers = 0). *Payoff* is the final payoff a participant received in the experiment (i.e., the total from all three games). *Non homo economicus* is an index of the frequency of participants acting against their myopic self-interest.[20] Specifically, a non homo economicus act in bargaining occurs when proposers offer an equal split or when responders reject positive offers. For the best shot game players act against their self-interest by contributing positive amounts as a first-mover and by responding to zero contributions with zero contributions as a second-mover. In the market, sellers who reject high offers act against self-interest, as do buyers who refuse to compete and make bids of zero. *Dissatisfaction* is an index of the frequency with which players were not treated kindly. Proposers who have their offers rejected or responders who receive offers less than half qualify as being dissatisfied with the interaction. In either role of the best shot game, when one's partner contributes nothing, that person is dissatisfied, and when sellers receive bids that are less than five francs or buyers have their bids rejected, they are dissatisfied. Also included are three treatment dummies (making random bargaining the reference treatment); the *sex* of each player (1 for female); and as a cultural catch-all, whether each player was born outside the United States, *foreign born*.

All of the regressions discussed were run with robust standard errors and use session-specific fixed effects to control for any unmodeled heterogeneity among players that may have been generated by idiosyncratic occurrences during an experimental session. Overall, the Wald χ^2 statistic on each regression can be seen as highly significant and explains a substantial amount of the variation in preferences changes. The regressors are split into two categories: exogenous effects and endogenous effects. Included in the exogenous effects are the different treatments, players' roles, initial social orientations, and other personal characteristics. The endogenous effects are variables that depend on how an individual and her coparticipants interact. These effects include player payoffs, how frequently she acted contrary to self-interest, and how often she was treated badly by the other participants.

In regression (1), consider only the exogenous effects. As one would expect, players' initial value orientations are strong predictors ($p < .01$) of how preferences will change (i.e., competitive types can only become more social, and altruists can only become more asocial), but neither a player's sex nor ethnicity seems to matter.[21] However, first-movers and

19. Rather than dropping these observations, treating them as having preferences that do not change does not substantially change the results. To save space, the most conservative results only are presented.

20. Another way of defining *non homo economicus* is as not choosing one's component of the perfect equilibrium prediction and not being close enough to have just made a small error.

21. Note, the initial value orientation was added as a regressor to control for the fact that people with high initial value orientations cannot become more other-regarding and players with low value orientations cannot become less other-regarding. All the results in Table 5 are qualitatively identical if one uses the tobit procedure for censored data instead.

sellers (who tend to earn more) become slightly more generous ($p < .10$), and, controlling for one's initial social orientation, strong effects can also be seen of the institutional dummies. Specifically, less anonymity (the same bargaining regressor) causes players to be more likely to become egoistic ($p < .01$) as does being exposed to market competition ($p < .01$) and having less motivation to reciprocate friendly acts ($p < .01$). Notice, these individual-level results corroborate and reinforce the treatment-level analysis (Table 2) in two of the three cases; market competition and reductions in the incentive to reciprocate social acts generate more asocial preferences compared to random bargaining, but in the third case we appear to have a contradiction. The Sign Test on the treatment-level data indicated participants became significantly more social in the same bargaining treatment, but regression (1) shows the opposite result. As can be seen, adding the endogenous effects and a few interactions in regressions (2) and (3) help resolve this puzzle.

In regression (2), the endogenous effects are added. First, interestingly, a player's payoff, by itself, is found to have no influence on her social orientation. But the more a player engages in non-self-interested play and the more dissatisfied she is with the way she has been treated by others, the more asocial she becomes; however, only the first effect is significant ($p < .05$). As in regression (1), in regression (2), controlling for other factors, women and foreign-born players are not more (or less) likely to change preferences and now a player's role has dropped off the list of significant determinants. Second, notice that adding the endogenous factors increases the magnitudes of the market and best shot coefficients but reduces the same bargaining coefficient by almost half and lowers its significance substantially.

What explains the changes in the treatment regressors? It appears that market competition and best shot inefficiencies erode social preferences independently of making interactions less friendly (i.e., the effect of the endogenous variables). The same bargaining coefficient, however, is reduced by the addition of the non homo economicus variable, which is, on average, significantly greater in the same bargaining treatment than in the random treatment.[22] Recall that for first-movers in the bargaining games, a non homo economicus act means making an offer for half the pie, and for second-movers the regressor picks up rejecting unfair offers. Because this variable absorbs much of the variation attributed to the same treatment dummy, it can be concluded that the same bargaining treatment elicits more fair offers and more rejections (controlling for the offer). Hence, the main effect of reducing anonymity in the same bargaining treatment is to elicit higher offers. In other words, the primary difference between the same and random bargaining treatments is that reducing anonymity decreases the amount of opportunistic behavior by first-movers.

To reinforce and expand on this explanation, in regression (3) the differential effect of making fair offers and having offers rejected on first-movers are examined. To do so, I interact *role* with the frequency of non-self-interested acts and with players' dissatisfaction and get the expected effects: making fair offers makes first-movers more social, but the effect is not significant and having one's offer rejected makes first-movers less social ($p < .01$). Notice, adding these two interactions again halves the coefficient on the same bargaining dummy, significantly reduces the best shot coefficient, and further increases the coefficient

22. This claim is based on adding the two variables one at a time, noting the difference in the size and significance of the same bargaining regressor, and testing the difference in the mean levels of non homo economicus in the two treatments.

on the market dummy. Now the story becomes clearer; being exposed to competition drastically erodes social preferences independently of how one is treated by one's peers, but the effects of less anonymous settings reduce to the differential ability of the same bargaining treatment to elicit higher offers and the reactions of proposers who have their offers rejected.

I end this section by summarizing the main results. First, the control study shows that there is a positive and significant association between preferences measured using the VO and the GARP exercises. Compared to the control, each of the treatments generates insignificant and lower levels of association, indicating that economic institutions affect social preferences. As hypothesized, market competition causes players to become less other-regarding and this effect is independent of the effects of payoff disparities and other endogenous determinants of preferences. A puzzle was also discovered concerning the effect of reducing anonymity and changing the incentives to reciprocate kind acts. To reconcile the results, the treatment effects of reducing the anonymity in bargaining revolving around the behavior and treatment of first-movers were demonstrated. Less anonymity matters only because repeated interactions allow second-movers to discipline unfair first-movers. In turn, having offers rejected makes first-movers much more egoistic, but making fair offers and being in the more powerful role make first-movers more altruistic. Combining these effects ends up with a few very unhappy first-movers who become more self-centered in the same bargaining treatment but many other first-movers who make fair offers and become more charitable.

4. Discussion

This article began by reviewing two contradictory theories about how economic institutions affect agents' social preferences. Specifically, the two theories differ on whether markets alienate people because they make interactions more or less anonymous and competitive. The results of the experiment suggest that, if markets are more like large anonymous supermarkets than small intimate farmers' markets, people's social preferences (i.e., their regard for other people's well-being) will diminish over time. Furthermore, an even stronger diminution of social preferences takes place when markets are highly competitive. In this case, those participants who are on the long side of the market (i.e., those who cannot make as many transactions as they would like) end up resenting the market structure and their competitors, which over time leads them to care less about the well-being of others. Perhaps most interestingly, the results also show that the negative effects of competition and anonymity on social preferences are not reducible to individual experiences alone. That is, controlling for how badly individuals are treated, settings that do not reward friendly acts, and competitive markets create atmospheres that themselves erode social preferences, a sort of framing effect.

The current results fit well with other, mostly psychological, studies that shed light on the relationship between markets, institutions, and the endogenous nature of social preferences. For example, Messick and Sentis (1985) show that social preferences in a hypothetical work situation are affected by perceived differences in work achievement. Specifically, this experiment (a within-subject design), though not about markets, illustrates (like the endogenous factors in Table 5) that individual social preferences are determined by the

nature of interactions between people. Furthermore, Breer and Locke (1965), who both pretest and posttest their subjects' social preferences, show that, controlling for initial preferences, those participants who participated in a work treatment that rewarded individual effort became more egoistic while those who were rewarded for collective effort did not. The current results nicely dovetail with this study to the extent that markets reward individual over collective effort.

Additionally, Loewenstein, Thompson, and Bazerman (1989) focus on the framing effects of market interactions. In this experiment participants were presented with three different hypothetical scenarios and were asked to rate a number of monetary outcomes for themselves and another person. The important factor that changed among scenarios (for our purposes) was the relationship between the two people interacting in each scenario. In one scenario, a dispute took place between two neighbors, and in the other the dispute was framed as a market interaction (between a customer and a sales manager). The data revealed that players were generally inequality averse in the neighbor treatment but liked to be better off than the other in the market scenario. While this experiment does not measure differences attributable to actual behavior in markets, it does provide corroboration that just framing an interaction as a market significantly deteriorates social preferences (i.e., this data validates the significant coefficient on the market regressor in Table 5 even when controlling for how well an individual is treated).

While these results are important for moral, theoretical, and institutional design reasons, such discussions are postponed until more, similar results are recorded. However, there are four directly related issues that arise concerning the results. First, in the face of mounting evidence from economic experiments, new theories have been developed that organize the behavioral results from many games based on social preferences for reciprocity, fairness, and inequality aversion (e.g., Falk and Fischbacher 1998; Fehr and Schmidt 1999; Bolton and Ockenfels 1999). A key feature of these models is to posit the kind of heterogeneity of social preferences seen in the preference data presented above and to show how, under certain institutional rules, all players appear to behave egoistically. For example, even cooperative or altruistic individuals behave competitively when they are on the long side of a market (we see this in the data analyzed here). Alternatively, cooperative players may withhold contributions in public goods games when egoists take advantage of their kindness. One contribution of the current results is to show that, with enough exposure, these players not only mimic egoists, they become egoists.

Second, these results are remarkable given the time scale of the interactions. That is, significant changes are seen in players' social orientations after exposure to different economic institutions for only an hour. A critical reader might be suspicious of player motivations that are so malleable. However, preference changes based on such short exposures are far from an anomalous result in the psychological literature. For example, Breer and Locke (1965) note that within four hours they were able to change seemingly robust attitudes toward individuals and society by repeated exposure to a task. Hence, while it is important to know how adaptive social preferences are and how persistent changes might be, it is not unreasonable to expect them to change quickly. In fact, there are situations in which it is reasonable for people to change their attitude toward others instantaneously. For example, blue-collar workers who are promoted to managers often appear to change their perceptions of workplace fairness overnight, and graduate students view comprehensive exams much differently almost immediately after they find out they have passed.

Third, are there other possible explanations of the data? For example, could it be the case that preferences are stable and the behavioral differences simply reflect learning? Or might it be the case that these are not preference changes but a more complicated schedule of stable preferences that depend on the framing of the interaction? Considering the first alternative, if one returns to Table 3, it does not appear that the subjects learned much about subgame perfection. Only in the market does behavior approach the predicted equilibrium, and the obvious lesson learned in this treatment is that it is good to be on the short side of the market. Let us assume however that some learning did occur. What would be expected? If people are learning self-interested strategies, more selfish preferences would be expected in the GARP experiment from all the participants. This is not what was seen. The second alternative states that players have *context dependent* preferences, which means that they are sensitive to how the interaction is framed. From an experimental point of view, this explanation has been controlled for by design because all the interactions used a standard neutral frame. But let us assume that players' preference schedules depend on the rules of the game instead. For example, people's schedules might depend on how much power they perceive they have in the game. The problem with this explanation is that it is hard to justify why framing the VO exercise would be perceived differently than the GARP exercise, and yet there is clearly different behavior in the two exercises. I concede, however, that much has not been said about the mechanism of preference changes. The only clues come from the regression results (Table 5), which support the commonsense hypothesis that a person's social orientation will depend on how well she is treated by others, which, as has been shown, depends on the rules of the game.

Finally, this experiment presents an apparent contradiction with other experimental results. Henrich et al. (2001) in their cross-cultural analysis of ultimatum bargaining behavior among members of nonindustrialized societies find two robust predictors of proposer behavior: the social returns to cooperation and the degree of market integration. Both coefficients are positive, and together they explain 68 percent of the variance in group average proposals. Note that these results suggest that markets correlate with more social individuals rather than less. However, an explanation of this apparent contradiction lies in the type of markets these people participate in. People in nonindustrialized societies participate in the idealized, intimate markets of Montesquieu, who wrote based on his experience in largely preindustrialized Europe. As the experiment here shows, markets are not the sole ingredient of alienation in that institutions must also foster anonymity and competition before a degradation of social preferences can be expected.

References

Andreoni, J., and J. Miller. 2002. Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica* 70: 737–57.

Barr, A. 2003. Trust and expected trustworthiness: Experimental evidence from Zimbabwe. *Economic Journal* 113: 614–30.

Bolton, G., and A. Ockenfels. 1999. ERC: A theory of equity, reciprocity and competition. *American Economic Review* 90: 166–93.

Bowles, S. 1998. Endogenous preferences: The cultural consequences of markets and other economic institutions. *Journal of Economic Literature* 36: 75–111.

Breer, P., and E. Locke. 1965. *Task experience as a source of attitudes*. Chicago: Dorsey.

Buckley, N., K. Chan, J. Chowhan, S. Mestelman, and M. Shehata. 2001. Value orientations, income and displacement effects, and voluntary contributions. *Experimental Economics* 4: 183–95.

Camerer, C., and E. Fehr. 2001. Measuring social norms and preferences using experimental games: A guide for social scientists. In *Foundations of human sociality: Experimental and ethnographic evidence from 15 small-scale societies*, ed. J. Henrich, R. Boyd, S. Bowles, H. Gintis, E. Fehr, and C. Camerer. Oxford: Oxford University Press.

Camerer, C., and R. Thaler. 1995. Anomalies: Ultimatums, dictators and manners. *Journal of Economic Perspectives* 9: 209–19.

Carpenter, J. 2002a. Information, fairness, and reciprocity in the best shot game. *Economics Letters* 75: 243–8.

———. 2002b. Is fairness used instrumentally? Evidence from sequential bargaining. *Journal of Economic Psychology* 24: 467–89.

———. 2002c. Measuring social capital: Adding field experimental methods to the analytical toolbox. In *Social capital and economic development: Well-being in developing countries*, ed. J. Isham, T. Kelly, and S. Ramaswamy, 119–37. Northampton, MA: Edward Elgar.

Charness, G., and M. Rabin. 2002. Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117: 817–70.

Falk, A., and U. Fischbacher. 1998. A theory of reciprocity. University of Zurich, Institute for Empirical Research in Economics Working Paper 006, Zurich, Switzerland.

Fehr, E., and K. Schmidt. 1999. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114: 769–816.

Forsythe, R., J. Horowitz, N. E. Savin, and M. Sefton. 1994. Fairness in simple bargaining experiments. *Games and Economic Behavior* 6: 347–69.

Griesinger, D., and J. Livingston. 1973. Toward a model of interpersonal motivation in experimental games. *Behavioral Science* 18: 173–88.

Gueth, W., R. Schmittberger, and B. Schwarze. 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3: 367–88.

Harrison, G., and J. Hirshleifer. 1989. An experimental evaluation of weakest link/best shot models of public goods. *Journal of Political Economy* 97: 201–25.

Henrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath. 2001. In search of homo economicus: Behavioral experiments in 15 small-scale societies. *American Economic Review* 91: 73–8.

Hirsch, F. 1976. *Social limits to growth*. Cambridge, MA: Harvard University Press.

Hirschman, A. 1982. Rival interpretations of market society: Civilizing, destructive, or feeble? *Journal of Economic Literature* 20: 1463–84.

Hoffman, E., K. McCabe, J. Shachat, and V. Smith. 1994. Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior* 7: 346–80.

Hoffman, E., K. McCabe, and V. Smith. 1996. On expectations and the monetary stakes in ultimatum games. *International Journal of Game Theory* 25: 289–301.

Kramer, R., and L. Goldman. 1995. Helping the group or helping yourself? Social motives and group identity in resource dilemmas. In *Social dilemmas: Perspectives on individuals and groups*, ed. D. Schroeder, 49–67. Westport, CT: Praeger.

Liebrand, W. 1984. The effect of social motives, communication and group size on behavior in an n-person multi-stage mixed-motive game. *European Journal of Social Psychology* 14: 239–64.

Loewenstein, G., L. Thompson, and M. Bazerman. 1989. Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology* 57: 426–41.

Marx, K., and F. Engels. 1847/1955. *The communist manifesto. With selections from the eighteenth Brumaire of Louis Bonaparte and Capital, by Karl Marx,* ed. S. H. Beer. New York: Appleton- Century-Crofts.

McClintock, C., and W. Liebrand. 1988. Role of interdependence structure, individual value orientation, and another's strategy in social decision making: A transformational analysis. *Journal of Personality and Social Psychology* 55: 396–409.

Messick, D., and K. Sentis. 1985. Estimating social and nonsocial utility functions from ordinal data. *European Journal of Social Psychology* 15: 389–99.

Montesquieu, C. de S., Baron de. 1749/1989. *The spirit of the laws*. Trans. A. M. Cohler, B. C. Miller, and H. S. Stone. Cambridge: Cambridge University Press.

Prasnikar, V., and A. Roth. 1992. Considerations of fairness and strategy: Experimental data from sequential games. *Quarterly Journal of Economics* 107: 865–88.

Roth, A. 1995. Bargaining experiments. In *The handbook of experimental economics*, ed. J. Kagel and A. Roth, 253–348. Princeton, NJ: Princeton University Press.

Schotter, A., A. Weiss, and I. Zapater. 1996. Fairness and survival in ultimatum and dictatorship games. *Journal of Economic Behavior and Organization* 31: 37–56.

Shure, G., and R. Meeker. 1967. A personality attitude schedule for use in experimental bargaining studies. *Journal of Psychology* 65: 233–52.

Siegel, S., and N. J. Castellan. 1988. *Nonparametric statistics for the behavioral sciences*. Boston: McGraw-Hill.

Jeffrey P. Carpenter completed his doctoral degree at the University of Massachusetts in 1999 and has been teaching at Middlebury College ever since. His research fields include experimental and behavioral economics, game and bargaining theory, and theoretical institutional economics. *His research is funded by the MacArthur Preferences Network and the National Science Foundation*.