The Demand for Punishment

by

Jeffrey P. Carpenter

September 2002

MIDDLEBURY COLLEGE ECONOMICS DISCUSSION PAPER NO. 02-43

DEPARTMENT OF ECONOMICS
MIDDLEBURY COLLEGE
MIDDLEBURY, VERMONT 05753

http://www.middlebury.edu/~econ

The Demand for Punishment[*]

Jeffrey P. Carpenter
Department of Economics
Middlebury College
Middlebury, VT 05753
(jpc@middlebury.edu)

September 30, 2002

Abstract: While many experiments demonstrate that actual behavior is different than predicted behavior, they have not shown that economic reasoning is necessarily incorrect. Instead, these experiments illustrate that the problem with *homo economicus* is that his preferences have been mis-specified. Modeled with social preferences, agents who forgo material gains can often be called rational. The current experiment illustrates this point with an example. Assuming self-interested agents, punishment is not credible in social dilemmas, yet people are often willing to incur costs to punish free riders. Despite this seeming irrationality, we show that these same people react to changes in the price of punishing and income as if punishment was an ordinary and normal good.

Keywords: public good, social dilemma, experiment, punishment

JEL Codes: C72, C92, H41

---

[*] I thank Julia Assael and Marla Weinstein for research assistance. I also thank Steve Burks, Simon Gaechter, Herb Gintis, and Peter Matthews whose comments substantially improved this paper. This research is supported by Middlebury College and the National Science Foundation (SES-CAREER 0092953).

Introduction

At this point in the evolution of experimental and behavioral economics laboratory experiments have provided more new questions about economic behavior than answers. Instead of confirming the standard tenets of neoclassical economics, experiments have identified decision-making anomalies (Camerer [1995]), preference reversals (Tversky et al. [1990]), and non-standard or "social" preferences (Camerer and Fehr [2001], Carpenter [2001], Charness and Rabin [2001]).[1] Expanding on the area of social preferences, experiments have shown that, instead of everyone being selfish and myopic, the average participant is much better described as trusting and trustworthy (Berg et al. [1995]), fair (Gueth et al. [1982], Fehr et al. [1993]), and cooperative (Isaac et al. [1984]), but also vindictive (Camerer and Thaler [1995], Fehr and Gaechter [2000]).

Notice, however, that the fact that many economic models predict behavior that is at odds with what we observe in experiments may be because we have misspecified peoples' preferences not because the methodology of economics is fundamentally flawed. Although people behave as if they have preferences for cooperation and retaliation, they may still react to incentives in ways predicted by standard economic logic. For example, if we hypothesize that peoples' observed preferences for cooperation operate like preferences for more standard consumption goods, then we might expect people to cooperate less when the implied price of cooperation increases just as they tend to buy fewer ordinary goods when the price increases.

In addition to being predisposed to cooperate, recent experiments have demonstrated that people retaliate against perceived injustices, even

[1] In fact, competitive markets is one of the few areas where experiments have come close to confirmed existing theories (Davis and Holt [1993]).

when doing so is costly and the benefits of doing so are small. This evidence (reviewed below) leads one to believe that many participants have a preference for punishing asocial behavior. In the experiment reported on herein, we use standard tools to test, in a controlled setting, whether such a nonstandard preference behaves according standard economic reasoning.

This research is unique because it is the first to explicitly examine whether standard economic tools can explain behavior motivated by the nonstandard preference to punish free riders. However, this research is linked to other recent work in behavioral economics. One area of research examines the sacrifices that people are willing to endure to assure fair outcomes and, in this sense, examines the price responsiveness of fairness preferences. Examples of this literature include Eckel and Grossman [1996], Suleiman [1996], and Zwick and Chen [1999]. In a second related project, James Andreoni and his coauthors (Andreoni and Vesterlund [2001], Andreoni and Miller [2002]) empirically recover utility functions that are based on social preferences from observed behavior. Variants of these utility functions could, in principle, generate the sort of demand for punishment functions that we estimate from the current data.

We proceed by briefly reviewing the literature on social dilemma experiments in which players were given the opportunity to punish each other. Hopefully, this review will convince the reader that cooperation and retaliation are robust behaviors. We then discuss the current experiment which was designed to examine whether peoples' preferences for punishment behave according to standard economic logic. Specifically, the experiment provides us with data which we use to estimate the demand for punishment. Our analysis indicates that punishment is both ordinary and normal, but is also relatively inelastic with respect to both price and income.

Fairness, Cooperation, and Punishment

The first evidence of a preference for punishing asocial behavior came from one-shot ultimatum games in which a first-mover makes an offer to share a sum of money with a second-mover who accepts or rejects this offer (Gueth et al. [1982], Camerer and Thaler [1995]). Although any division of the pie can be supported as an equilibrium of this game, subgame perfection leads one to expect that the first-mover will receive all (or almost all) of the money because selfish second-movers will always accept small offers rather than reject them and get nothing. Despite this unambiguous prediction, nearly all small offers are rejected and the most common explanation given by second-movers is that they are retaliating against greedy first-movers (Pillutla and Murnighan [1996]).

Punitive behavior has also been witnessed in social dilemma games in which individual and group incentives are at odds, and therefore, free-riding is expected from selfish players. One of the first of these experiments was conducted by Elinor Ostrom and her colleagues (Ostrom et al. [1992]). In this common pool resource experiment players cooperate with each other by not extracting too much from an open-access and subtractable resource. Resource use is problematic because by extracting, one player imposes a negative externality on all the other players. Under these incentives, the authors showed that when costly punishment was allowed cooperative players used it to regulate the behavior of over-extractors (i.e. free riders) and the gross efficiency of extraction increased, especially when communication was allowed too.

Considering positive rather than negative externalities, Fehr and Gaechter [2000] tested whether costly punishment could curtail free riding in a public goods experiment. In the voluntary contribution mechanism players emit a positive externality every time they contribute to a group project, the benefits of which are shared by the entire group.

Given this structure, selfish players should contribute nothing and free ride on the contributions of others. Fehr and Gaechter's results mirror those of Ostrom et al. [1992] in that they find that many contributors are willing to pay to punish those who contribute less than the average. Further, the (theoretically incredible) threat to punish reduces free riding dramatically. These results suggest that when subjects punish free riders they are expressing a social preference for retaliation because they punish despite having to pay to do so and despite the negligible material benefits that are expected to follow punishment.

Experimental Design

While the following experiment is based on the voluntary contribution mechanism (Isaac et al. [1984]), to test whether we can explain punishment in terms of standard economic logic we made a few changes. Our changes were designed to provide us with the data to estimate the demand for punishment. First, we allowed players to monitor and punish each other. Second, punishment was costly to impose and the price of punishment changed during the course of the experiment. This feature allows us to estimate the price elasticity of the demand for punishment. Third, the level of provision of the public good during each round determines an income for each player from which players paid to punish each other. This feature allows us to estimate the income elasticity of demand. Also note that because players' earnings and the price of punishment varied over the course of the experiment we are able to analyze the demand for punishment using a (more powerful) within subject design. The specifics of our experiment are as follows.

Define the *price of punishment, r,* as the amount a punisher must pay, in experimental monetary units (EMUs), to remove one EMU from the target. Our experiment was 15 periods long and each session was split into five blocks, each block lasting three periods. The price of punishment varied from block to block such that $r \in \{0.25, 0.5, 1, 2, 4\}$. We ran two treatments to balance the effect of changing prices. In the *decreasing price* treatment *r* equaled 4 for the first three periods, meaning the punisher spent 4 EMUs to remove 1 EMU from the target, *r* equaled 2 in periods four through six, and so on until in periods thirteen through fifteen the price was 0.25. In the increasing price treatment *r* started at 0.25 and cycled upward to 4. Our players were randomly assigned to a treatment and we ran a total of six sessions (three for each treatment). This design resulted in a total of 18 four-person groups.

We used the familiar *strangers* protocol (Andreoni [1988]) under which players are randomly reshuffled from group to group at the beginning of each period because we wanted to control for any strategic reasons to punish. For example, players may perceive that they would benefit later if they also anticipate that punishment will cause free riders to contribute more in the future. However, if the target of one's punishment is likely to be in a different group next period, the expected benefit of punishing is negligible. Controlling for strategic punishment is important because doing so allows us to focus on punishment as the expression of a social preference.

The payoff function for the voluntary contribution mechanism was augmented to account for punishment.[2] Imagine groups of *n* players, each of whom can contribute any fraction of their *w* EMU endowment to a public good and keep the rest. Say player *i* free rides at rate $0 < \sigma_i < 1$ and contributes $w(1-\sigma_i)$ to the public good, the benefits of which are shared equally among the members of the group.

Each player's contribution was revealed to one other player in

[2] The instructions (see the appendix) referred to "reductions" with no interpretation supplied.

the group who could punish this person at a price of r EMUs per sanction. Let rs_{ij} be the expenditure on sanctions assigned by player i to player j and let s_{ki} be the sanctions player i receives from player k (the instructions explicitly mentioned that $j \neq k$), then the payoff to player i is

$$\pi_i = w[\sigma_i + nm(1-\bar{\sigma})] - rs_{ij} - s_{ki}$$

where $\bar{\sigma} \equiv (\sum \sigma_i)/n$ is the average free riding rate in the group. The variable m is the marginal per capita return on contributions to the public good (see Ledyard [1995]). In all sessions n equaled 4, m was set to 0.5, and w was 25 EMUs.

Because $1/n < m < 1$ the game without punishment is a social dilemma: group incentives are at odds with individual incentives. Each contributed EMU returns only 0.5 to the contributor which means free riding is a dominant strategy. But if $\bar{\sigma} = 1$ then everyone free rides fully and each player's payoff is lower than it would be if everyone contributed fully. The game is finitely repeated which implies that subgame perfection predicts complete free riding on every round.

Notice that adding the possibility of punishment does not change the subgame perfect prediction. Because sanctions are costly to impose and any potential benefits from getting a free rider to contribute cannot be fully internalized by the punisher, punishment is incredible and therefore can not be a component of any subgame perfect equilibrium. Without credible punishment, free riding is still subgame perfect.

As noted above, each player monitored and was able to sanction only one other member of the group. This design feature was added to control for any possible strategic or coordination reasons that might affect players' punishing propensities. For example, if each player monitors and can punish all the other members of the group, there are at least two problematic scenarios that may arise. First, from a strategic perspective, a player may be less likely to punish a free rider because she thinks she can free ride on the punishment of others. Second, a player may be less likely to punish because she can not explicitly coordinate her punishment efforts with the rest of her group. For example, she may feel that the free rider should be punished, but also that there is an appropriate level of punishment that fits the infraction. If she does not know, or can not estimate, how much others will punish she may withhold sanctions to be sure that the punishment does not exceed the offense. If each player sees only one other player and knows that the person they are monitoring is not monitoring them, we control for any strategizing and coordination problems. People should only pay to punish if they wish to express their preferences.

Overview of the Data

We recruited 72 participants (36% were female) in our six experimental sessions. Participants earned an average of $26.26, including a $5 show-up fee. The typical session lasted a little less than an hour. We begin our analysis by giving the reader a broad sense of the data and then we focus on our estimates of the demand for punishment.

By reviewing previous punishment experiments (e.g. Fehr and Gaechter [2000]) we see that the typical time path of contributions, averaging across treatments, starts near half the endowment and then increases at a decreasing rate. However, as seen in Carpenter [2002], punishment has less of an effect on contributions when players only monitor a subset of their group-mates. Figure 1 illustrates the time paths from the current experiment. In one sense our contributions data look similar to the other limited monitoring data because contributions do not increase monotonically. However, in another sense the current data is markedly different because there seems to be a treatment effect.

Figures 1 and 2 about here

The treatment effect in our contributions data makes sense from an economic point of view. When the price of punishment starts at a relatively low level and then increases over the course of the experiment, contributions fall steadily and more dramatically than when the price is constant (as in Carpenter [2002]). One explanation, which we will confirm below, is that our players based their punishment decisions on price, as well as, on how egregiously the target free rode. On the one hand, when the price increased over time players bought less punishment causing the threat of punishment to abate. This lead to more free riding. On the other hand, when the price fell over time players responded by buying more punishment per offense. In this case, the effectiveness of punishment increased over the course of the experiment and, although we see an initial drop in contributions, they recover as the price of punishment continues to fall.

Figure 2 presents the time paths of the average expenditure on punishment. Even though this graph does not control for other factors that might have affected our players' punishment decisions (e.g. income or average level of free riding), it provides evidence consistent with the hypothesis that players reacted to the price of punishment and this affected the credibility of punishment and the level of contributions. As the price increased, our players spent less on punishment. In fact, by the last three rounds of the increasing price treatment when it cost 4 EMUs to remove 1 EMU from the target, the players stopped punishing completely. In the other treatment, as the price fell, players spent more on, and bought more, punishment.

Because the instructions explicitly mentioned the order in which the price of punishment would change (see the appendix), one might worry that players anticipated and reacted in advance to the direction of the price change. For example, in the increasing price treatment, it might be reasonable to think that players spent more on early punishment than they would have had they not known that the price was going to increase. Or, players in the decreasing price treatment might have delayed punishment to later rounds when they knew it would be cheaper. If this is true then the slopes of the graphs in figure two are steeper than they would have otherwise been.

Because we do not want to detract from our main purpose of estimating the demand for punishment by explicitly modeling player expectations, we will present evidence indicating that expectations did not affect players' choices. In both treatments, the per sanction price of punishment was 1 EMU during periods 7, 8, and 9. If expectations played a significant role then, by the above logic, we would not expect expenditures in the two treatments to be the same, controlling for the amount of free riding. If the expenditures per offense are the same then we have evidence that expectations did not matter. When we calculate the ratio of the expenditure on punishment to the amount kept by the target during these three rounds and then compare these ratios between treatments, we find no significant difference using either the Wilcoxon test or the Kolmogorov-Smirnov test ($Z=0.77$, $p=0.44$; $ks=0.05$, $p=0.99$).

The Demand for Punishment

We now proceed by econometrically estimating a demand for punishment equation. One valuable benefit of using an experiment to elicit the data for our estimation is that we control for most of the problems that typically plague demand estimates. Specifically, simultaneity and identification are not problems for us because price is, by design, completely exogenous. However, we do face other issues. Because our experiment is 15 periods long, we generate a panel of data. To control for individual heterogeneity, all our regressions include random effects. Because there are a lot of observations where our players showed no preference for punishment, our dependent variable, the quantity of punishment purchased, is truncated from below at zero. For this reason, we use the Tobit procedure.

Before we present the fully controlled estimate of the demand for punishment, we present the reader with a graphical presentation of the main result. Figure 3 illustrates an uncontrolled demand for punishment function. In figure 3 observations are represented by numbers which indicate the average amount contributed by the targets who received the designated amount of punishment at the corresponding price. Further, the size of each number is determined by how many observations there are at each location. For example, there are a lot of observation where the price of punishment was 0.25 and the quantity was 0. Considering all these observations, the average contribution by the targets was 13 EMUs. Likewise, there was one case in which the punisher purchased 22 units of punishment at a price of 1 to direct at a person in his or her group who contributed only 3 EMUs.

Figure 3 about here

Figure 3 illustrates three facts about the demand for punishment. First, by the size of the numbers we see along the quantity equals zero axis, most players were not punished. In fact 78% of our observations were for zero punishment. However this makes sense given the value of these numbers. The average contribution across sessions and periods is 7.57 EMUs and, based on Fehr and Gaechter [2000], we know that punishment is directed, primarily, at those who contribute less than the average. Second, notice that, for any given price, as one reads up to higher quantities of purchased punishment, one reads smaller valued numbers. This fact demonstrates that punishment is proportional to how much one free rides. Third, and most importantly for our purposes, the regression line indicates that punishment exhibits the standard substitution effect: as the price increases, people punishment less.

We can demonstrate the robustness of figure 3 by considering the regression results presented in table 1. We build our econometric model in three stages. In stage one we estimate the uncontrolled price and income elasticities. In stage two we add controls for how much the target contributed (the null hypothesis being one is punished less the more one contributes), and for how much the punisher contributed (the null, in this case, being that people who contribute more, punish more). Finally, in stage three we control for the time trend apparent in our data (i.e. figures 1 and 2) by modeling a simple dynamic which says that the amount of punishment purchased in period t is a function of the change in group cooperativeness between periods $t\text{-}1$ and t.[3]

In the upper half of table 1 we report the coefficients and standard errors of our regressions and in the bottom half we report elasticities calculated at the regressor means.[4] Equation (1) introduces

[3] Notice that we could have included period fixed effects as an alternative to modeling a dynamic, but this approach would be incorrect because doing so assumes there are time idiosyncrasies while what we need to control for appears to be a dynamic process.

[4] Calculating elasticities from tobit coefficients is not straightforward because, when one calculates the marginal effect, one has to account for the probability that a change in the regressor will push one past the "kink," and the impact of a change in the regressor on the dependent variable, given it is uncensored. However, we can use the McDonald and Moffitt [1980] decomposition to calculate elasticities. With latent variable, p_{it} we have the following marginal effect.

$$\frac{\partial E(p_i \mid x_i)}{\partial x_{it}} = E(p_i^* \mid x_i, p_i^* > 0) \bullet \frac{\partial \Pr(p_i^* > 0)}{\partial x_{it}} + \Pr(p_i^* > 0) \bullet \frac{\partial E(p_i^* \mid x_i, p_i^* > 0)}{\partial x_{it}}$$

our main result. Both the price and income elasticities are negative which indicates that, given the average participant prefers to punish free riders, people react to economic incentives in what economists would consider reasonable fashion. However demand appears to be inelastic with respect to price and elastic with respect to income. Specifically, a one percent increase in price reduces the quantity of punishment demanded by 0.90 percent and a one percent increase in income decreases the amount of punishment demanded by 1.24 percent. In sum, punishment is ordinary and inferior according to our simplest model.

Table 1 about here

Equation (2) indicates that only one of our initial elasticity estimates is robust to the inclusion of other punishment determinants. Part of the variation in punishment previously attributed to changes in price is actually caused by changes in how egregiously the target free rides and how much the punisher contributed, but the coefficient remains negative and highly significant. At the same time, there is a dramatic change in our estimate of the income elasticity. There is a simple explanation for the change in the sign of the marginal effect of income. Income and Target's Contribution are correlated (rho=0.63). Without controlling for the target's contribution, the income regressor picks up the variation due to both how much the target free rides and how much income is generated by the group. Because free riders are punished less the more they contribute and because punishers sanction more when their incomes increase, the combined effect is negative and results in the inferior characterization of punishment in equation (1).

We now find that punishment is normal, and inelastic with respect to income – a one percent increase in income increases the amount of punishment purchased by 0.89 percent. We also find that a one percent increase in the contribution made by the target reduces the amount of punishment he or she can expect by 0.62 percent and a one percent increase in the amount that the punisher contributes to the group account increases the amount he or she will punish by 0.58 percent.

In equation (3) we test whether our elasticity estimates are robust to an accounting of the trend in group cooperativeness. As the reader can see, controlling for the dynamic effect of contributions on punishment has little effect on any of the elasticities. Finally, in an unreported regression, we allowed for the fact that punishers tend to sanction those who give less than them more severely than those who give more. Although we find evidence to this effect (as do Falk et al. [2001]), namely the marginal effect on those who contribute less is 0.16 ($p<0.01$) and the marginal effect on those who contribute more is 0.08 ($p<0.01$), the price and income marginal effects change by less than 5%, and therefore the elasticities are also relatively unchanged. We conclude that the demand for punishment is ordinary and normal, but inelastic.

A Slutsky Decomposition

Averaging across periods, treatments, and individuals those participants who punished spent 14% of their per period income on sanctions. Given this is a significant fraction of their earnings, a change in the price of punishment has a dramatic effect on their *real* budget constraints in the experiment. For this reason, we would like to further decompose our data to ascertain how much of the change in behavior attributed to a change in price is due to an income effect that is not picked up in our income elasticity, and how much is due to a pure substitution effect.

Figure 4 about here

Consider the choice between punishment, *s*, and a composite good, *x*. As above, the price of punishment is *r* and the per period income is π. Lastly, assume the composite good is normalized to be the numeraire. In figure 4, we derive the substitution effect by isolating the change in *s* resulting from an adjustment in the participant's income after a reduction in *r*. Income is adjusted so that the pre-change chosen bundle (s^*, x^*), purchased for π^*, is just affordable. This change will be due entirely to the change in the price of punishment. We calculate the income effect by multiplying the needed income change by the marginal effect of income on the choice of *s*. In other words, differentiating the identity,

$$s^s(r, s^*, x^*) \equiv s(r, rs^* + x^*)$$

and rearranging terms, where s^s is the Slutsky demand function for punishment, leads to the following punishment price effect decomposition.

$$\frac{\partial s(r,\pi^*)}{\partial r} = \frac{\partial s^s(r,s^*,x^*)}{\partial r} + \frac{\partial s(r,\pi^*)}{\partial \pi} s^*$$

The first term is the total effect of a change in the price of punishment (i.e. the observed change is *s*), the first term after the equal sign is the substitution effect and the last term is the income effect.

Table 2 about here

Table 2 summarizes our empirical approximation of the substitution and income effects. We calculate estimates of the income and substitution effects for each price transition and for both price treatments. Further, we offer estimates for two data sets: all the data and just the uncensored data. We begin by calculating the average quantity of punishment demanded before a price change, $\bar{s}_0$, and the average after the price change, $\bar{s}_1$. The difference $|\bar{s}_0 - \bar{s}_1|$ is the average observed effect of a price change. For each price change and $\bar{s}_0$, we calculate the income compensation needed to make the original bundle just affordable by finding, $\Delta\pi = |r_0\bar{s}_0 - r_1\bar{s}_0|$. Multiplying $\Delta\pi$ by our estimate of the marginal effect of income on punishment from table 3, equation (3) yields our estimate of the income effect of a given price change. We then back out the substitution effect by subtracting the income effect from the total observed change. In the final column of table 2 we calculate how much of the total observed change is due to the substitution effect.

With the exception of the transition from $r=4$ to $r=2$, the substitution effect explains, on average, more than 90% of the observed change in punishment following a price change. Although the magnitudes of the substitution effects are larger in the uncensored data, so are the observed changes. Therefore, the relative size of the substitution effect remains about 90% of the observed change. In other words the income effect is relatively small and the substitution effect is relatively large.

Concluding Remarks

At the beginning of this paper we pointed out that, while laboratory experiments in economics have provided more puzzles than answers, we should not be too quick to conclude that the standard methodology of economics is inherently flawed. The results of the current experiment give us hope that after documenting and understanding anomalies like social preferences, economic tools are informative.

Specifically, our analysis has demonstrated three things: One, we have replicated and extended the experiments suggesting that the average economic decision-maker will, at some personal cost, punish free riders who reduce the social efficiency of group interactions. Adding the current evidence to that of a number of other experiments illustrates that positing a preference for punishing free riders appears to be a reasonable addition to standard, selfish, preferences. Two, given we accept that people prefer to punish free riders, we have shown that the most basic economic analysis, the estimation of demand, illustrates that people react to price and income changes when they consider punishing free riders just as they react to changes in these variables when they consume more standard commodities. Specifically, the demand for punishment slopes downward and is relatively inelastic with respect to price and income. Third, we have shown that punishers are sensitive to both the price of punishment and the fact that more income allows one to punish more severely. However, decomposing the effect of a price change, we see that most of the change in punishment is due to substitution; changes in real income play only a small role.

Appendix – Participant Instructions

You have been asked to participate in an experiment. For participating today and being on time you have been paid $5. You may earn an additional amount of money depending on your decisions in the experiment. This money will be paid to you, in cash, at the end of the experiment. When you click the BEGIN button you will be asked for some personal information. After everyone enters this information we will start the instructions for the experiment.

During the experiment we will speak in terms of Experimental Monetary Units (EMUs) instead of Dollars. Your payoffs will be calculated in terms of EMUs and then translated at the end of the experiment into dollars at the following rate: 25 EMUs = 1 Dollar.

In addition to the $5.00 show-up fee, each participant receives a lump sum payment of 10 EMUs at the beginning of the experiment.

The experiment is divided into 15 different periods. In each period participants are divided into groups of 4. The composition of the groups will change randomly at the beginning of each period. This means that in each period your group will consist of different participants.

Each period of the experiment has two stages.

Stage One

At the beginning of every period participants receive a 25 EMU endowment. In stage one participants decide how much of their 25 EMUs to contribute to a group project and how much to keep for themselves. Participants' payoffs are determined by the total contribution of their specific group and how much they individually keep.

To record their decisions, participants will type EMU amounts in two text-input boxes, one for the group project labeled GROUP ALLOCATION and one for themselves labeled PRIVATE ALLOCATION. These boxes will be yellow. Once a participant makes a decision, he or she will record his or her decision by clicking on the green SUBMIT button.

After all the participants have made their decisions, you will each be informed of your gross earnings for the period.

Participant Gross Earnings will consist of two parts:

(1) Earnings from the Private Allocation. Individuals are the only beneficiary of EMUs they keep. Specifically, each EMU a participant keeps increases that person's earnings by one.
(2) Earnings from the Group Project. Each member of a group gets the same payoff from the group project regardless of how much he or she contributed. The payoff from the group project is calculated by multiplying 0.5 times the total EMUs contributed by the members of the group.

Participant Gross Earnings can be summarized as follows:

$1 \times$ (EMUs you keep) + $0.5 \times$ (Total EMUs contributed by your group)

Let's discuss three examples.

Example 1: Say each member of a group contributes 15 of the 25 EMUs. In this case, the group total contribution to the project is $4\times15 = 60$ EMUs. Each group member earns $0.5\times60 = 30$ EMUs from the project. The gross earnings of each member will then be the number of EMUs kept, 25-15 = 10, plus the earnings from the group project, 30 EMUs, for each member. In total, each member would earn 10+30 = 40 EMUs.

Example 2: Now say everyone in the group contributes 5 EMUs. Here the group total contribution will be 20 and each member will earn 0.5×20 = 10 EMUs from the group project. This means that the total earnings of each member of the group will be 20 (the number of EMUs kept) plus 10 (earnings from the group project) which equals 30 EMUs.

Example 3: Finally, say three group members contribute all their EMUs and one contributes none. In this case, the group total contribution to the project is 3×25 = 75 EMUs. Each group member earns 0.5×75 = 37.5 EMUs from the project. The three members who contributed everything will earn 0+37.5 = 37.5 EMUs and the one member who contributed nothing will earn 25+37.5 = 62.5 EMUs.

Stage Two

In stage two participants will be shown the allocation decision made by one other randomly selected member of their group. Everyone's choice will be seen by exactly one other group member and the person you see is different from the person seeing you. In addition to seeing another group member's choice, at this stage participants can reduce the earnings of the group member they see, if they want to.

Participants will be shown how much one member of their group kept and how much this person allocated to the group project. Participants will also see their own allocation decision and this decision will be labeled 'YOU'.

At this point participants will decide how much (if at all) they wish to reduce the earnings of the other group member they are seeing.

Participants reduce someone's earnings by typing the number of EMUs they wish to spend to reduce that person's earnings into the input-text box that appears below the other group member's allocation decision.

Participants can spend as much of their accumulated earnings as they want to reduce the earnings of the other group member. For each EMU spent by a participant the earnings of the other group member will be reduced by R EMUs. The value of R will change during the experiment.

[Price Decrease] The experiment is divided into 5 blocks of 3 periods and the value of R will change every 3 periods according to the following sequence {0.25, 0.5, 1, 2, 4}. For example, during the first 3 periods of the experiment R will be 0.25 so spending 1 EMU will reduce the other group member's earnings by 0.25 EMUs. During the third block of periods R will equal 1 and spending 1 EMU will reduce the other group member's earnings by 1 EMU. During the final block R will equal 4 and spending 1 EMU will reduce the other group member's earnings by 4 EMUs.

[Price Increase] The experiment is divided into 5 blocks of 3 periods and the value of R will change every 3 periods according to the following sequence {4, 2, 1, 0.5, 0.25}. For example, during the first 3 periods of the experiment R will be 4 so spending 1 EMU will reduce the other group member's earnings by 4 EMUs. During the third block of periods R will equal 1 and spending 1 EMU will reduce the other group member's earnings by 1 EMU. During the final block R will equal 0.25 and spending 1 EMU will reduce the other group member's earnings by 0.25 EMUs.

Consider this example: suppose someone spends 2 EMUs to reduce the earnings of the other group member when R is 0.5. This expenditure reduces the other group member's earnings by 1 EMU ($2\times0.5=1$). When participants have finished stage two they will click the blue DONE button.

Participant Net Earnings in each period will be calculated as follows:

> (Gross Earnings from Stage One) – (R times the number of EMUs spent on reductions directed towards the participant) – (the participant's expenditure on reductions directed at someone else).

If you have any questions please raise your hand. Otherwise, click the red FINISHED button when you are done reading.

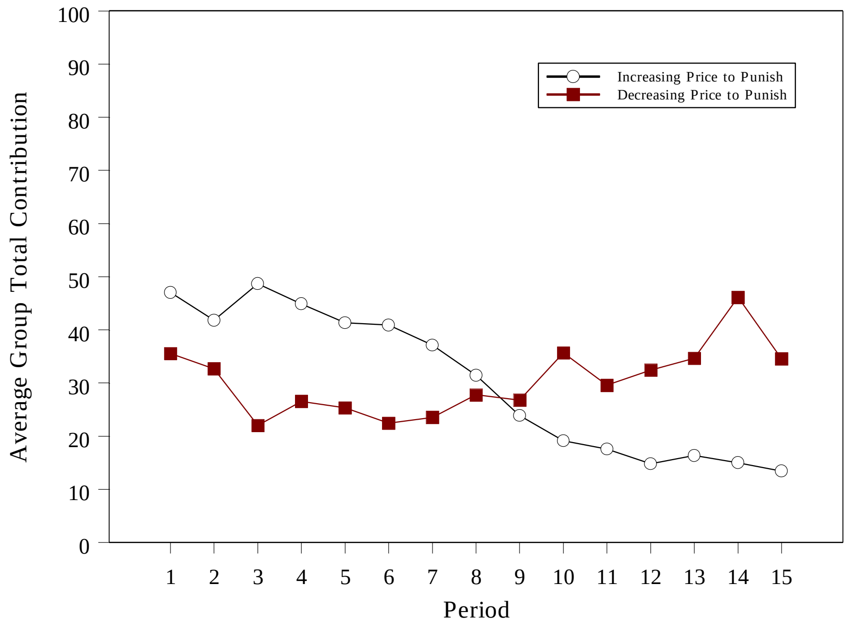

Figure 1 – The evolution of average group contributions over time (note: increasing price indicates that the price per sanction increased from 0.25 to 4 while decreasing price means the opposite).

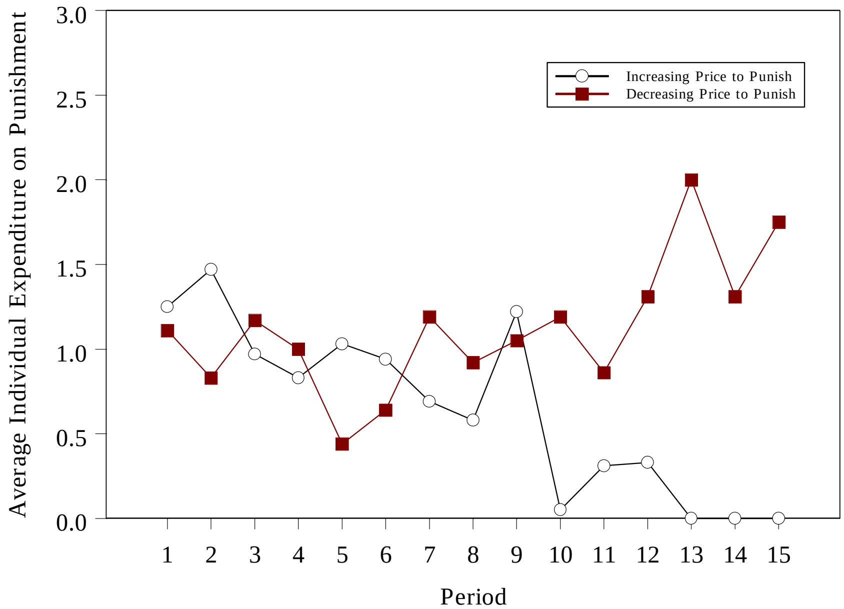

Figure 2 – The time path of punishment expenditures (note: this figure includes all the cases where players choose to not punish at all and does not control for how much free riding occurred).

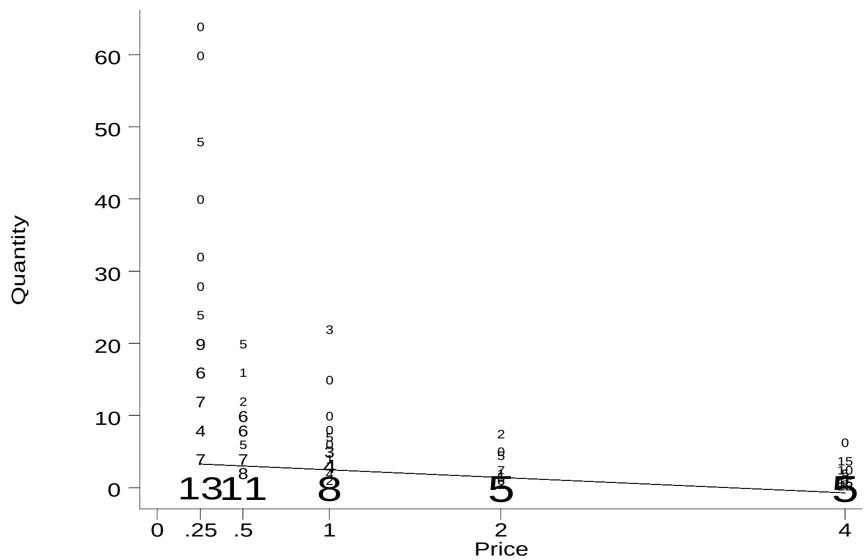

Figure 3 – The demand for punishment (note: each observation is denoted by a number. The size of the number is weighed by the number of identical observations. The value of the number denotes the average contribution of the punishment target for all the identical observations. The line is the result of a linear regression of quantity on price).

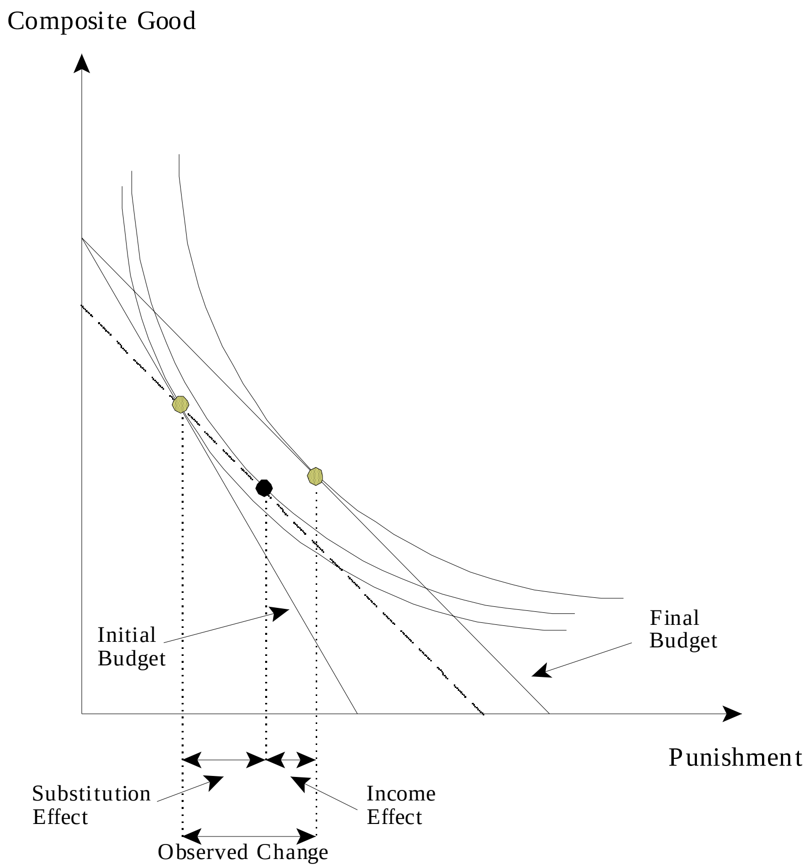

Figure 4 – Punishment income and substitution effects (assuming punishment is ordinary and normal).

Dependent Variable = Quantity of Punishment Inflicted on Target $_{i,t}$	(1)	(2)	(3)
Price $_{i,t}$	-1.07***	-0.73***	-0.75***
	(0.14)	(0.12)	(0.13)
Income $_{i,t}$	-0.07***	0.05***	0.05***
	(0.02)	(0.02)	(0.02)
Target's Contribution $_{i,t}$		-0.15***	-0.13***
		(0.02)	(0.02)
Punisher's Contribution $_{i,t}$		0.14***	0.12***
		(0.02)	(0.02)
Group Total $_{i,t}$ - Group Total $_{i,t-1}$			-0.01**
			(0.005)
Price Elasticity	-0.90	-0.61	-0.64
Income Elasticity	-1.24	0.89	0.91
Target Elasticity		-0.62	-0.54
Contribution Elasticity		0.58	0.50
Change in Total Elasticity			0.01
N	1080	1080	1008
Wald χ^2, p-value	107, <0.01	208, <0.01	205, <0.01

Notes:
(1) Each regression is a Tobit and includes individual random effects.
(2) The coefficients are the marginal effect of a change in the regressor on the expected value of the observed quantity purchased.
(3) The elasticities are calculated at the regressor means.
(4) The N in equation (3) is lower because we lose period 1 to differencing.
(5) Significance: * indicates 0.10, ** indicates 0.05, *** indicates 0.01.

Table 1 – Calculating the elasticities of demand.

Punishment Slutsky Calculations								
All Data								
Direction of Price Change	Price Values	$\overline{s}_0$	$\overline{s}_1$	Compensating Income ($\Delta\pi$)	Income Effect $\left(\frac{\partial s}{\partial \pi} \bullet \Delta\pi\right)$	Observed Change $\lvert \overline{s}_0 - \overline{s}_1 \rvert$	Substitution Effect $\lvert \overline{s}_0 - \overline{s}_1 \rvert - \left(\frac{\partial s}{\partial \pi} \bullet \Delta\pi\right)$	Substitution Effect as percentage of total change
Increase	0.25 -> 0.50	4.93	1.87	1.23	0.06	3.06	3.00	98%
	0.50 -> 1.00	1.87	0.83	0.94	0.05	1.04	0.99	95%
	1.00 -> 2.00	0.83	0.11	0.83	0.04	0.72	0.68	94%
	2.00 -> 4.00	0.11	0	0.22	0.01	0.11	0.10	91%
Decrease	4.00 -> 2.00	0.26	0.35	0.52	0.03	0.09	0.06	67%
	2.00 -> 1.00	0.35	1.05	0.35	0.02	0.70	0.68	97%
	1.00 -> 0.50	1.05	2.24	0.53	0.03	1.19	1.16	97%
	0.50 -> 0.25	2.24	6.74	0.56	0.03	4.50	4.47	99%
Uncensored Data Only								
Increase	0.25 -> 0.50	13.30	6.12	3.33	0.17	7.18	7.01	98%
	0.50 -> 1.00	6.12	4.74	3.06	0.15	1.38	1.23	89%
	1.00 -> 2.00	4.74	1.79	4.74	0.24	2.95	2.71	92%
	2.00 -> 4.00	1.79	0	3.58	0.18	1.79	1.61	90%
Decrease	4.00 -> 2.00	1.56	2.68	3.12	0.16	1.12	0.96	86%
	2.00 -> 1.00	2.68	4.56	2.68	0.13	1.88	1.75	93%
	1.00 -> 0.50	4.56	6.54	2.28	0.11	1.98	1.87	94%
	0.50 -> 0.25	6.54	17.33	1.64	0.08	10.79	10.71	99%

Table 2 - Decomposing price related punishment changes into substitution and income effects (note: $\overline{s}_0$ is the average level of punishment before a price change, $\overline{s}_1$ is the average level after the change, and π is per period income).

Bibliography

Andreoni, J. (1988). Why free ride? Strategies and learning in public good experiments. *Journal of Public Economics* 37: 291-304.

Andreoni, J. and J. Miller (2002). Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica* 70(2): 737-757.

Andreoni, J. and L. Vesterlund (2001). Which is the fair sex. *Quarterly Journal of Economics* 116(February): 293-312.

Berg, J., J. Dickaut and K. McCabe (1995). Trust, reciprocity and social history. *Games and Economic Behavior* 10: 122-142.

Camerer, C. (1995). Individual decision making. *The handbook of experimental economics*. J. Kagel and A. Roth Eds. Princeton, Princeton University Press: 588-683.

Camerer, C. and E. Fehr (2001). Measuring social norms and preferences using experimental games: A guide for social scientists. *Foundations of human sociality: Experimental and ethnographic evidence from 15 small-scale societies*, Oxford University Press. forthcoming.

Camerer, C. and R. Thaler (1995). Anomalies: Ultimatums, dictators and manners. *Journal of Economic Perspectives* 9(2, Spring): 209-219.

Carpenter, J. (2001). Measuring social capital: Adding field experimental methods to the analytical toolbox. *Social capital and economic development: Well-being in developing countries*. S. Ramaswamy Ed. Northampton, Edward Elgar. forthcoming.

Carpenter, J. (2002). Punishing free-riders: How group size affects mutual monitoring and collective action. mimeo.

Charness, G. and M. Rabin (2001). Understanding social preferences with simple tests. *Quarterly Journal of Economics* forthcoming.

Davis, D. and C. Holt (1993). *Experimental economics*. Princeton, Princeton University Press.

Eckel, C. C. and P. J. Grossman (1996). The relative price of fairness: Gender differences in a punishment game. *Journal of Economic Behavior & Organization* 30: 143-158.

Falk, A., E. Fehr and U. Fischbacher (2001). Driving forces of informal sanctions. University of Zurich, Institute for Empirical Research in Economics(working paper 59).

Fehr, E. and S. Gaechter (2000). Cooperation and punishment in public goods experiments. *American Economic Review* 90(4): 980-994.

Fehr, E., G. Kirchsteiger and A. Riedl (1993). Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics*: 437-459.

Gueth, W., R. Schmittberger and B. Schwarze (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3: 367-388.

Isaac, R. M., J. Walker and S. Thomas (1984). Divergent evidence on free-riding: An experimental examination of possible explanations. *Public Choice* 43(1): 113-149.

Ledyard, J. (1995). Public goods: A survey of experimental research. *The handbook of experimental economics*. J. Kagel and A. Roth Eds. Princeton, Princeton University Press: 111-194.

McDonald, J. and R. Moffitt (1980). The uses of tobit analysis. *Review of Economics and Statistics* 62: 318-321.

Ostrom, E., J. Walker and R. Gardner (1992). Covenants with and without a sword: Self-governance is possible. *American Political Science Review* 86: 404-417.

Pillutla, M. and K. Murnighan (1996). Unfairness, anger, and spite: Emotional rejections of ultimatum offers. *Organizational Behavior and Human Decision Processes* 68(3): 208-224.

Suleiman, R. (1996). Expectations and fairness in a modified ultimatum game. *Journal of Economic Psychology* 17: 531-554.

Tversky, A., P. Slovic and D. Kahneman (1990). The causes of preference reversals. *American Economic Review* 80: 204-217.
Zwick, R. and X.-P. Chen (1999). What price for fairness? A bargaining study. *Management Science* 45(6): 804-823.