Promoting Cooperation in

Nonlinear Social Dilemmas through Peer Punishment*

Timothy N. Cason
Department of Economics
Purdue University

and

Lata Gangadharan
Department of Economics
Monash University

January 2014

Abstract

Many social dilemmas exhibit nonlinearities and equilibrium outcomes in the interior of the choice space. This paper reports a laboratory experiment studying whether peer punishment promotes socially efficient behavior in such environments, which have been ignored in most experimental studies of peer punishment. It compares the effectiveness of peer punishment in a linear public good game to the effectiveness of this decentralized enforcement mechanism in two nonlinear social dilemma games: a piecewise linear public good game and a common pool resource game. While peer punishment improves cooperation in these new environments, the impact of punishment is weaker and takes longer to be effective. This appears to be due to the greater complexity of the nonlinear settings, which makes socially optimal choices more difficult to identify.

Keywords: Experiment; Voluntary Contributions Mechanism; Common Pool Resources; Public Goods; Nonlinear Returns.
JEL Classification: C90, D70, H41

***** We are grateful for valuable comments provided by two anonymous referees, Marco Casari, Nikos Nikiforakis, Elinor Ostrom, Michael Price, Steven Tucker, Jimmy Walker, and audiences at Indiana University (Bloomington), the Universities of Tennessee, Warwick and Zurich, the European Workshop on Experimental and Behavioral Economics and the Australia-New Zealand Experimental Economics Workshop. We retain responsibility for any errors. This research has been supported by a grant from the U.S. Environmental Protection Agency's National Center for Environmental Research (NCER) Science to Achieve Results (STAR) program. Although the research described in the article has been funded in part through EPA grant number R833672, it has not been subjected to any EPA review and therefore does not necessarily reflect the views of the Agency, and no official endorsement should be inferred. Funding from the Cooperative Research Centre for Water Sensitive Cities (CRC grant number 20110044) is also acknowledged. Part of this research was conducted while Cason was a visiting scholar with the Department of Economics, Monash University. Daniel Nedelescu provided very capable research assistance.

1. Introduction

Costly peer punishment to promote cooperation in social dilemmas such as public goods provision is an important, influential and well-studied phenomenon in experimental economics. Allowing individuals to punish their peers for socially bad behavior can improve cooperation in groups and sometimes allow them to achieve more efficient outcomes in social dilemmas by instilling positive social norms (Ostrom, 1990; Fehr and Gaechter, 2000; Chaudhuri (2011) provides a recent survey). Norm enforcement by peers rather than a centralized authority could therefore reduce social inefficiencies, and this suggests that self-governance can replace costly enforcement of regulations in some circumstances.

Many studies have followed the Fehr and Gaechter (2000) approach and constructed a body of knowledge about the impact and robustness of alternative peer punishment institutions, but few have deviated from the simplified linear environment reported in that seminal paper. Our paper makes a novel contribution by examining whether peer punishment promotes socially efficient behavior in environments that exhibit nonlinearities and have equilibrium outcomes in the interior of the choice space. We find that though peer punishment improves cooperation in these new environments, the impact of punishment is weaker and takes longer to become effective. This appears to be due to the greater complexity of the nonlinear settings, which makes socially optimal choices more difficult to identify.

Ostrom et al. (1992) report one of the first experiments on peer-punishment.[1] Although they found that peer-punishment had some impact on behavior, it was quite ineffective when introduced in isolation. Ostrom et al. highlight this by summarizing, "…covenants, even without

[1] Yamagishi (1986) is the earliest influential experimental paper on peer punishment. The contribution structure in this study is different from most social dilemma games examined subsequently. Individuals did not receive any return from their own investment in the public good, but their contributions provide positive (and linear) returns to other group members.

a sword, have some force, while *swords without a covenant may be worse than the state of nature*" (page 414, emphasis added). While this paper is widely cited, the main result that peer punishment only reduced free-riding when it was combined with communication is overshadowed by the more prominent interpretation that peer punishment helps to resolve social dilemmas. Ostrom et al. do not find that peer punishment alone increases cooperation, as the authors emphasize in their book providing more detail on the experimental results (Ostrom et al. 1994).[2] This early study, however, incorporates peer punishment differently than most of the subsequent literature. For example, subjects in Ostrom et al. could only punish one peer per period, they participated in relatively large groups (n=8), and they could not vary their intensity of punishment.

The Ostrom et al. (1992) and Fehr and Gaechter (2000) studies initiated the broad literature on decentralized mechanisms for groups to address social dilemmas, but they employ specific environmental features that exist in distinct sub-literatures. Both are public goods provision problems, but Ostrom et al. is based on the Common Pool Resource (CPR) environment with nonlinear payoff functions, while Fehr and Gaechter's study is based on the linear Voluntary Contribution Mechanism (VCM). The effectiveness of peer-punishment for improving cooperation, even without communication, is seen to be a robust empirical finding in the VCM sub-literature, and researchers have explored variations in the strength, relative cost, type, and information available for punishment. Surprisingly, however, almost no research has examined the robustness of these results to alternative nonlinear environments such as the CPR

[2] Ostrom et al. (1994) note that "in only one of the eight experiments do we see the sanctioning mechanism having a major impact on the level of net yield accrued" (page 178), and although they remark upon "the lack of a significant improvement in net yield accrual with the introduction of a sanctioning mechanism" (page 179) they do not report a statistical test. We extracted from the detail provided in their book the raw data on group contributions and net yield from the public good to conduct such tests. Consistent with the statements quoted above, these tests reveal no significant increase in yield due to sanctioning using either non-parametric Mann-Whitney tests comparing across-sessions (p-value=0.22) or within-session increases due to the introduction of sanctioning (p-value=0.41), or when using a parametric regression model (p-value=0.29). Details are available upon request.

setting.[3] For some reason, the compelling results of Fehr and Gaechter (2000) for the simple linear VCM environment steered research on decentralized peer punishment away from other more realistic nonlinear settings.

This study aims to address this gap in the literature by examining the effectiveness of costly peer-punishment in a richer set of environments that include nonlinearities. The linear structure employed in nearly all previous experiments leads to the starkest kind of social dilemma because it leads to boundary predictions, where the socially-optimal and self-interested (Nash equilibrium) outcomes are at the extreme opposite locations in the choice set. This makes noncooperative behavior easy to identify for targeting punishment. Nonlinear environments may be considered as more representative of many practical situations, since they typically lead to equilibrium and socially optimal choices that are not on the boundaries of the choice space. This is consistent with the many social dilemmas in the field that feature non-zero self-interested contributions but socially-optimal choices that do not devote all resources to the public good. As we discuss in the next section, nonlinear environments are common for social dilemmas in the field, such as when firms in an industry decide on production levels or when researchers decide how much time and effort to spend on their own research versus others' research via refereeing or mentoring.

Our experiment studies behavior in linear and nonlinear social dilemma games with and without costly punishment opportunities, integrating these two sub-literatures in an otherwise identical institution (e.g., identical choice space, matching, feedback, etc.). The linear VCM

[3] Three exceptions are reported in Casari and Plott (2003), van Soest and Vyrastekova (2006), and Janssen et al. (2010). Casari and Plott show that punishment improves cooperation, but for a special (historical) institution in which successful punishers receive the "fines" paid by the punished. van Soest and Vyrastekova employ a quadratic CPR and compare reward and punishment mechanisms. Peer punishment is observed to be more effective than rewards in improving cooperation. Janssen et al. do not find punishment to be effective in increasing cooperation in a CPR game with complex spatial and temporal dynamics.

design is used as a baseline treatment. A nonlinear VCM treatment features a piecewise linear private return such that the marginal per capita return decreases for greater contributions to the public good. A CPR treatment (in which higher contributions imply harm to the group rather than benefit) introduces nonlinearities in the group return similar to Ostrom et al. (1992), but with punishment introduced in a more "standard" way used in recent research.

Our results provide the first clear evidence that peer punishment can effectively reduce free-riding in nonlinear environments.[4] The experiment shows that the main conclusions of the peer punishment literature are observed in new and more complex environments, including nonlinear private or nonlinear public returns, and when reversing whether "allocations to the public account" create positive or negative externalities on others. Importantly, however, we also find that the impact of punishment is weaker and takes more time to emerge in the nonlinear environments. This is particularly so in the VCM treatment where the private returns are nonlinear. The magnitude of punishment's impact on cooperation is also smaller in the CPR treatment than in the linear VCM, with outcomes often closer to the Nash equilibrium than the social optimum even when punishment is available.

The impact or effectiveness of peer punishment is usually measured by its effect on contributions. Contributions are influenced by the intensity with which individuals punish fellow group members for free riding and by how the group members react to being punished. Our results indicate that the weaker impact of punishment in nonlinear environments is mainly due to the reduced intensity of punishment. Our comparison of the results across treatments suggests

[4] While van Soest and Vyrastekova (2006) use a quadratic CPR and punishment technology similar to the present study, and find that peer punishment improves cooperation, their results are unable to establish a clear causal link between peer punishment and improved cooperation. Punishment opportunities are always introduced after rounds without punishment in their design, so the marginal impact of punishment cannot be identified independently from learning and the time trend towards greater cooperation. In addition, as mentioned in the previous footnote the focus of their work in on the relative effectiveness of rewards and punishment. In contrast to their paper, our study can isolate the causal effect of peer punishment independent of learning and we also compare different nonlinear environments.

that this may be due to the greater difficulty for subjects to identify free riding actions that deserve punishment.

2. Public Good Dilemmas in Nonlinear Environments

The social dilemma studied in economics as the public good problem has been modeled in many different ways. It is essentially an externality problem in which agents allocate scarce resources, and some allocations benefit the decision-maker while other allocations benefit others. Individuals who focus on their own benefits and costs ignore the externality that their allocation imposes on others' benefits and costs, and so the self-interested allocation is socially inefficient.

This general description does not specify the relative sizes of the benefits and costs, and in fact many different economic environments feature this classic social dilemma. While standard theoretical models of this problem do not assume specific functional forms, for convenience most of the experimental literature has used constant benefits and costs, which leads to a linear payoff structure that is simple to describe to experimental subjects. As noted in the introduction, this linear special case has been used almost exclusively in the literature on public goods provision with peer punishment. This is typically called the linear Voluntary Contributions Mechanism (VCM), characterized by the following payoff function:

$$\pi_i = C_i + \alpha(E_i - g_i) + \beta \sum_{j=1}^{n} g_j$$

The parameters C_i, E_i, α, β denote (respectively) a fixed payment, the endowment that agent i is to allocate, returns to the endowment not allocated to the public good, and the returns to allocations towards the public good. The agent's choice variable is g_i, the amount of the endowment to contribute to the public good. This payoff function represents a social dilemma if

$\beta < \alpha < n\beta.$ The first inequality implies that individual returns from the public good are lower than private returns, so that zero public good contribution is the dominant strategy Nash equilibrium. The second inequality implies that total returns received by the group from any public good allocation exceed the private returns, so the social optimum is for every agent to contribute all resources to the public good. The ratio β/α is usually referred to as the marginal per-capita return (MPCR) (Isaac and Walker, 1988).

Clearly this representation of the social dilemma with constant and linear returns is a special case, and since the unique Nash equilibrium ($g_i^* = 0$) and the social optimum ($g_i^{s.o.} = E_i$) are at extreme ends of the choice space it is likely to make free-riding particularly easy to identify. This could significantly facilitate agents' ability to use decentralized peer punishment to enforce social norms of high contributions, and indeed Fehr and Gaechter (2000) document that a large fraction of individual contributions reach the social optimum of $g_i^{s.o.} = E_i$ in their fixed groups ("partners") treatment.

Many social dilemma and public goods provision problems encountered in the field, however, do not have this "all-or-nothing" structure where self-interest and socially-optimal outcomes are at the extremes of the available choice space. Economic problems often have equilibria and socially-efficient outcomes that are in the interior of the choice space, such as Cournot quantity choice games to take a prominent example from the field of Industrial Organization. It is also not hard to think of simple examples from everyday life, such as the amount of journal refereeing or uninspiring committee service effort that academics contribute to the public good. For many people a zero contribution of effort is not optimal, even for those who are strictly self-interested, since this service provides private as well as public benefits. Moreover, devoting all of one's efforts towards these activities is clearly not socially optimal as the costs of contribution are convex. Many social dilemmas share characteristics of this example and therefore feature equilibria and efficient allocations that are on the interior of the choice space.

The linear VCM could be modified in a variety of ways to move the self-interested equilibrium and socially-efficient outcomes to the interior. One simple way is to make the private and/or public returns α and β non-constant. This is an approach used by Sefton and Steinberg (1996), for example, who employ decreasing marginal returns to the public good in one treatment to induce non-dominant strategy Nash equilibria in the interior, and decreasing marginal private returns in another treatment to create a unique interior dominant strategy Nash equilibrium. Isaac and Walker (1998) use simple quadratic public good returns to generate interior equilibria. In the most simple nonlinear environment we employ, described in the next section in more detail, we adapt the piecewise linear private returns used in Bracha et al. (2011) that lead to a dominant strategy unique interior Nash equilibrium and an interior social optimum.

A richer but more complex way to introduce nonlinearities is the Common Pool Resource (CPR) environment studied by Ostrom et al. (1992). As in the linear VCM, allocations to a "private account" have a constant rate of return. Allocations to a "public account" represent effort exerted in extracting from a CPR, and therefore the group return obtained by an agent depends on the total allocations to the group account by all agents with access to the CPR. For example, in a congested fishery the total catch depends on the total effort exerted in the group, $G = \sum_{j=1}^{n} g_j$, and agent i's share of the catch depends on her share of the effort g_i/G. This leads to the following payoff function for this CPR setting:

$$\pi_i = C_i + \alpha(E_i - g_i) + \frac{g_i}{G}F(G)$$

Drawing on the literature from natural resource management Ostrom et al. use a nonlinear, concave group return function $F(G)$, and our design follows their lead in employing a quadratic functional form. (This quadratic structure has also been adopted throughout the experimental CPR literature.) They argue that "while no formal game or laboratory experiment ever captures all the nuances of field settings, this *n*-person CPR is a far more realistic environment… than many of the dilemma games previously explored" (page 405).

Besides adding some realism, there are several other reasons why studying decentralized peer punishment in this version of the social dilemma is important. As in the environments discussed above, it leads to Nash equilibria that are in the interior of the choice space, and are not in dominant strategies. Thus, also more realistically, a self-interested individual's allocations to the public good depend on the allocations made by others. This also makes the equilibria, and the (interior) social-optimum, more challenging to identify, so the levels of free riding and cooperation are less transparent. Consequently, enforcing social norms of cooperation through peer punishment may become more demanding (Reuben and Riedl, 2013).[5]

This CPR environment also creates a more challenging setting for cooperation since the payoff gains to deviating from a cooperative allocation are greater, and the ability of a subgroup of agents to gain from mutual cooperation are weaker, compared to the VCM environment. This is because the *individual distribution factor* (Maier-Rigaud and Apesteguia, 2006) in the CPR makes the group payoff *rival*, as indicated by the g_i/G in the previous equation. Gains to cooperation and benefits of defection are known to influence the frequency of cooperative and non-cooperative equilibrium play, as Isaac and Walker (1988) discovered for the linear public goods environment when they identified the influence of the MPCR on cooperation and free-riding. In the CPR, the free rider can extract nearly all of the gains from the others' cooperation

[5] Along similar lines, Nikiforakis et al. (2012) find that asymmetries in public goods returns make peer punishment less effective in a linear VCM environment, which they attribute to conflicting social norms. Xiao and Kunreuther (2013) find that punishment is less effective in a stochastic than a deterministic prisoner's dilemma, and they provide evidence that this is due to normative conflict because in the stochastic environment it is less clear whether actions or outcomes should be punished.

because public good returns are determined by agents' share of total allocations. For the CPR parameters used in the present experiment a set of cooperators can increase their payoffs by only 1 or 2 percent if just a single self-interested agent free rides on their cooperation.[6] In other versions of the CPR and other social dilemmas (e.g., the centipede game), cooperation must be unanimous for cooperators to receive positive returns (e.g., Murphy et al., 2006).

The share-based returns to the public good also create much higher gains to an agent who defects from a cooperative agreement in the CPR, relative to the VCM. In the two versions of the VCM used in our experiment, deviating from the social optimum with a best-response public good allocation (which is a dominant strategy in these cases) raises the deviator's payoff by 14 percent. These gains can be easily taken away from the deviator by a single peer punisher who incurs punishment costs for only a single period. By contrast, in the CPR a deviator can extract most of the public good returns, raising her payoff by 43 percent. Thus, the gains from defection are tripled in the CPR relative to the VCM. While these specific figures are based on the particular parameters chosen for this experiment, these structural differences in payoffs between the environments are general (Maier-Rigaud and Apesteguia, 2006). Thus it is plausible that peer punishment may not increase cooperation in nonlinear public goods settings such as the CPR.

3. Experimental Design

The experiment includes six different treatments: three environments summarized in Table 1, each conducted with and without peer punishment opportunities. The design is unorthodox, but necessary whenever the environments considered are not nearly identical. Unlike standard

[6] In the VCM, by contrast, cooperators can "tolerate" free riding by one individual and still reap substantial gains from cooperation over these periods. For the parameters implemented in the experiment, a subset of 3 cooperators can increase their payoffs by 19 percent in the linear and nonlinear VCM environments, although of course the free rider benefits even more from the others' cooperative public good contributions.

treatment comparisons that change only one characteristic at a time, the comparisons of this study necessitate changes in multiple factors to reflect the underlying change in the environment. As noted at the bottom of the table, however, all institutional features were kept fixed across all treatments, including the use of n=4 individuals per group, 10 periods of fixed-group repetition, and a strategy space of $g_i \in \{0, 1, \ldots, 11, 12\}$ tokens that could be allocated to the group or the private account. The framing used in the instructions was similar to that in the literature and it was maintained exactly consistent across all treatments. We also chose parameters so that the payoff gain to cooperation was (approximately) equal across the three environments, since we are primarily interested in how peer punishment can promote cooperation.

We also made the maximum gains to defection from the social optimum approximately equal across the two VCM environments. This helps in the interpretation of the results. The previous section sketched two mechanisms through which nonlinearities may affect the effectiveness of peer punishment in enhancing cooperation. First, nonlinearities move the social optimum and noncooperative Nash equilibrium contribution levels to the interior of the choice space. This increases *complexity* because it is less clear to subjects what constitutes free riding and the actions that should be punished, and could also result in the emergence of different contribution norms across groups. Second, nonlinear environments can increase *payoffs from deviation* from the social optimum, which are of course important for the likelihood of sustained cooperation in repeated games. The deviation returns are (approximately) equal in the linear and nonlinear VCM environments, so any differences in peer punishment and its effectiveness across the two VCM treatments should be attributed to complexity differences. The deviation returns are much greater in the CPR setting, so any differences in outcomes between the CPR and nonlinear VCM are likely due to differences in payoffs from deviation.

The linear VCM environment we employ is standard and serves as the baseline. The MPCR is 0.5 and the Nash equilibrium of zero contributions to the public good for own-payoff maximizing agents is in dominant strategies. As noted in the previous section, the nonlinear VCM employs a piecewise linear private return so that the Nash equilibrium (a per-subject contribution of 3 tokens to the public good) is also in dominant strategies. The MPCR is 1.25 for the first 3 tokens, 0.5 for the next 7 tokens, and 0.2 for the final 2 tokens contributed to the group account.[7] The low MPCR for the final 2 tokens moves the social optimum to the interior of the choice space, and corresponds to each subject contributing 10 of her 12 tokens to the public good.

In the CPR environment subjects also had to allocate 12 tokens to the group and their private account. In this environment, however, an allocation to the group account (although not framed to the subjects this way) creates negative externalities as it is effort spent extracting from a common pool resource. In the results section we will refer to lower extraction from the CPR as greater contributions to the public good, because greater extraction imposes a negative externality on others. The instructions explained that "If a total of X tokens are placed in the group account by group members, then the total group payoff is $18X-0.4X^2$ Experimental Dollars" and that "Your share of this total group payoff equals your number of tokens placed in the group account as a fraction of the total tokens you and the others in your group place in the group account." This explanation was followed by numerical examples, and as shown in the instructions Appendix, the total group returns were displayed numerically on the subjects' input screens. Group returns, as well as each subject's share of the group returns, were also displayed on the input screens in the two VCM treatments.

[7] In the instructions subjects were told that "Your first 2 tokens provide a payoff of 5 Experimental Dollars each. Tokens 3 to 9 provide a payoff of 2 Experimental Dollars each. Tokens 10, 11 and 12 provide a payoff of 0.8 Experimental Dollars each."

The experiment employed quadratic group returns, consistent with the CPR literature following Ostrom et al. (1992). It is straightforward to show that the Nash equilibrium total group investment G is $[n/(n+1)](a\text{-}w)/b$, where n is the number of group members (4), w is the return to tokens retained in the private account (2), and a and b are coefficients of the group return function (18 and 0.4). It is also easy to show that the socially optimal group investment is $(a\text{-}w)/2b$. Therefore, the ratio of the Nash equilibrium allocation to the social optimum allocation is completely determined by the number of group members n in this quadratic case: $2n/(n+1)$. As the number of group members increases the externality caused by greater extraction increases, causing a larger divergence between the group and individually optimal extraction levels. Because we kept the number of group members fixed across environments at n=4, the ratio of Nash equilibrium allocation to the social optimum allocation was fixed at 1.6; specifically, allocations of 8 and 5 for our parameter choices.

Although the range between the equilibrium and socially optimal allocations is smaller in the CPR than in the two VCM environments, as noted above due to the nature of the different payoff functions subjects' gains from deviating from the social optimum are substantially larger for the CPR. Figure 1 shows that deviations from the social optimum allow the deviator to increase his payoff by at most 6 or 7 experimental dollars in the two VCM environments. By contrast, in the CPR a deviator can place all of his tokens in the group account (i.e., devoting all available effort to resource extraction) and increase his payoff by 22.4 experimental dollars. Thus, the temptation to defect from the social optimum is over three times higher in the CPR, which could make coordination on this group optimum considerably more difficult. Since the incentives to deviate from the cooperative allocation are very similar in the two VCM environments, the comparison of outcomes across environments will help reveal whether any differences are due to nonlinearities per se or to the stronger incentives to deviate in the CPR relative to the nonlinear VCM environment.

In three treatments, subjects had only one decision each period: how many tokens to allocate to the group account. The other three treatments added a second decision stage each period, in which subjects observed the allocations by all individuals in their group and had an opportunity to assign "deduction points" to others. The framing used to explain this feature to subjects is standard in the peer punishment public goods literature, and was taken from Gaechter et al. (2008). Each deduction point cost the assigner 1 experimental dollar, but it reduced the target's payoff by 3 experimental dollars. A subject could assign up to 5 deduction points to each of the other 3 group members, for a maximum of 15 points. Potentially a subject could receive 15 deduction points, resulting in a payoff loss of 45 experimental dollars. Payoffs could be negative in a given period, although this was very uncommon (it occurred in only 8 out of 1360 possible subject-periods). Subjects did not know who assigned deduction points to them, and the order that individual contributions were displayed was randomly shuffled each period so that subjects could not acquire individual reputations. These design choices are also standard in the peer punishment public goods literature.

All sessions were conducted at the Vernon Smith Experimental Economics Laboratory at Purdue University, using z-Tree (Fischbacher, 2007). Subjects were drawn from the undergraduate student population, broadly recruited across the university by email using ORSEE (Greiner, 2004). Although some had participated in other economics experiments, all were inexperienced in the sense that they had never participated in a similar experiment featuring public goods or common pool resource characteristics and incentives. While subjects interacted anonymously in 4-person fixed groups, multiple groups under the same treatment conditions were conducted simultaneously in the laboratory, employing 16 to 24 subjects in each session. As the final row of Table 1 shows we have 24 groups (independent observations) for the nonlinear VCM and the CPR treatments, 12 each with and without punishment opportunities. In the baseline linear VCM treatments we collected data from 20 groups, with an equal number of punishment and non-punishment groups.

Subjects were randomly assigned to different groups upon arrival, and they remained in those positions throughout the experiment. Subjects made decisions in one or two stages as described above, depending on the treatment. They participated in 10 periods and this number of periods was common knowledge and announced in the instructions.[8] At the beginning of each experimental session an experimenter read the instructions aloud while subjects followed along on their own copy. All accumulated experimental dollars were converted to U.S. dollars at a pre-announced 40-to-1 conversion rate and paid in cash at the conclusion of the experimental session. The 272 subjects earned US$19.75 each on average, with an interquartile range of $18.00 to $21.75. Including the instruction and payment distribution time, sessions usually lasted about one hour.

4. Results

We report the results in five subsections, corresponding to five key results. The first subsection compares contributions in the three environments, with and without punishment opportunities. The second subsection contrasts end-period contributions for the two nonlinear environments with their interior equilibrium predictions and the third subsection focuses on overall efficiency.

[8] An additional 10 periods were conducted following this as a "surprise restart," either adding or removing punishment depending on the condition used initially. Results were very similar in these restart periods so we do not report them here. The data and statistical tests reported in the paper are all based on between-subject comparisons that employ only the initial 10 periods, so as to maintain independence assumptions required for the statistical tests.

The fourth subsection documents similarities in punishment behavior across all environments, and the final subsection explores whether contributors could profitably deviate from cooperative contribution levels given the observed punishment behavior.

4.1 Contributions

Result 1: (**Contributions**) Contributions are significantly higher with punishment opportunities than without in all three environments, although the impact of punishment is weaker and requires more time to emerge in the nonlinear environments.

Support: Figure 2 summarizes the average contributions to the group account in the three environments, reported separately for the treatments with and without punishment opportunities. Panel A shows that the experiment replicates the standard result for the linear VCM environment that contributions begin at approximately half of the endowment, but they fall gradually without peer punishment opportunities and rise over time when subjects can punish. (As we document below in Subsection 4.4, the free-riders are the individuals who are typically punished.) Panel B indicates relatively steady average contributions in the nonlinear VCM for the first 5 or 6 periods, before contributions diverge in the punishment and non-punishment treatments. By the tenth period, average contributions are 57 percent higher with punishment opportunities than without. Panel C illustrates that contributions also begin near the midpoint of the choice interval (6) in both CPR treatments. "Contributions" after the first few periods tend to be lower when punishment is available, and in this environment these lower contributions provide a positive externality to others in the group.

Table 2 indicates that the differences in average contributions across the punishment and non-punishment treatments are highly significant for the linear environment, considering either all 10 periods or only the last 5 periods. (The nonparametric Wilcoxon tests reported in the lowest row employ one observation—the average contribution—from each statistically independent group of 4 subjects.) Consistent with the visual impression of Figure 2, treatment effects due to punishment do not emerge in the nonlinear VCM until the last 5 periods. The main impact of punishment in this treatment is to prevent the onset of decay in contributions, rather than to increase contributions over time as it does in the linear VCM. Since the gains to cooperation and the gains to defection are almost identical in the linear and nonlinear VCM environments (see Table 1), this finding that cooperation takes longer to emerge in the nonlinear VCM is likely due to the nonlinearities per se and the resulting interior location of the Nash equilibrium and social optimum.

Contributions are also significantly different between the punishment and no punishment conditions in the CPR environment, although near the conventional margin of 5-percent both for all periods and in the late periods. The magnitude of the punishment influence is also smaller in this environment. Similar to the nonlinear VCM, punishment mostly stems the decline in contributions. Recall, however, that the range of the expected contributions between the Nash equilibrium and the Social Optimum is compressed in the CPR (8 - 5) relative to the nonlinear VCM (3 - 10) and especially to the standard linear VCM (0 - 12). To account for the different expected contribution ranges across treatments, we also calculated a 0-to-1 index to assess the change in contributions when peer-punishment is allowed – normalizing by the predicted difference in the contributions in the Nash and the Social Optimum. The results based on this index provide identical conclusions (so we do not report them here) and confirm that punishment is least effective across all periods in the nonlinear VCM.

Separating the contributions in the three treatments into early (periods 1-5) and late (periods 6-10) contributions, we find that when peer punishment is allowed, the contributions in the late periods are significantly higher than in the early periods in the linear VCM (paired observations Wilcoxon test p-value = 0.012), but not in the two nonlinear treatments (Wilcoxon test p-values of 0.195 and 0.239) Hence while in the linear VCM punishment seems to both prevent decay and promote higher contributions over time, in the nonlinear environments it only prevents the breakdown of cooperation over time. Since both nonlinear environments lead to weaker and slower impacts of punishment, the results overall are consistent with the interpretation that the complexity introduced by nonlinearities makes it more difficult for some subjects to identify the social optimum and recognize the free riding actions that should be punished.

4.2 Equilibrium and Social Optimum Benchmarks

In the large literature on the linear VCM game, Nash equilibrium and Social Optimum contribution levels are at the boundaries of the choice space and therefore contributions are always between these theoretical benchmarks—as illustrated in our data in Panel A of Figure 2. Since deviations from these benchmarks are necessarily one-sided, statistical tests can typically reject those theoretical predictions easily. The interior benchmarks in our nonlinear environments, however, allow contributions to be above and below the equilibrium and social optimum targets so it is reasonable to compare behavior to these theoretical predictions statistically.

Result 2: (**Nash Equilibrium and Social Optimum**) In the final period of the two nonlinear environments, average contributions are significantly different from the social optimum both with and without punishment opportunities, and are significantly different from the stage game Nash equilibrium except in the CPR environment without punishment opportunities.

Support: For the nonlinear VCM, as shown in Panel B of Figure 2 average contributions are always well below the Social Optimum of 10, and are above the Nash equilibrium of 3. A

Wilcoxon signed-rank test (based on one observation per 4-player group) rejects the null hypothesis that average contributions in the final period are 10 (*p*-values<0.01) in both the punishment and non-punishment treatments; and this test rejects the null that average contributions are 3 in the punishment (*p*-value<0.01) and non-punishment (*p*-value=0.025) treatments. Panel C of Figure 2 illustrates that contributions in both treatments of the CPR environment remain well above the Social Optimum of 5, but for the non-punishment condition they approach the Nash equilibrium of 8. A Wilcoxon signed-rank test rejects the null hypothesis that average contributions in the final period are 5 in both the punishment and non-punishment treatments (both *p*-values<0.01); and while this test rejects the null that average contributions are 8 in the punishment condition (*p*-value<0.01), in this final period the contributions are not significantly different from 8 in the non-punishment condition (*p*-value=0.365).

4.3 Efficiency

As highlighted originally in Ostrom et al. (1992), the deadweight losses from costly peer punishment can significantly erode gains from improved cooperation of public good contributions. In order to compare efficiency with and without punishment across the three environments, we employ an index based on realized payoffs, denoted π_{actual}, compared to environment-specific Nash equilibrium and socially optimal payoffs, denoted π_{Nash} and $\pi_{Soc.Op.}$. These realized payoffs include losses incurred when assigning and receiving punishment points. The efficiency index is

$$Efficiency = \frac{\pi_{actual} - \pi_{Nash}}{\pi_{Soc.Op.} - \pi_{Nash}}$$

This index varies between 0 and 1 as payoffs vary between levels earned at the Nash equilibrium and the social optimum, but it can also move outside this range when subjects earn more extreme payoffs.

Result 3: (**Efficiency**) Efficiency is not significantly different with and without punishment opportunities in any of the three environments.

Support: Figure 3 summarizes the average efficiency in the three environments, with separate averages reported for the treatments with and without punishment opportunities. In all three panels efficiency is lower initially in the punishment condition, but eventually rises to reach or exceed efficiency in the non-punishment condition. The efficiency time pattern is quite erratic in the CPR environment with punishment. Statistical tests indicate that punishment opportunities do not have a significant influence on efficiency in any of the three environments, regardless of whether all 10 periods or only the final 5 periods are considered. (The lowest two-tailed p-value for the six Wilcoxon rank-sum tests is 0.225.)

4.4 Punishment Behavior

Figure 4 displays the time series of average punishment points that subjects assigned for the three treatments. The averages are similar, and are not statistically different across any of the treatment comparisons. Although punishment tends to decline in all three treatments, this decline is only statistically significant in the nonlinear VCM treatment (Wilcoxon p-value=0.003 comparing periods 1-5 to periods 6-10). This punishment decline is the main cause of the efficiency increase in this treatment (shown in Panel B of Figure 3). The erratic punishment in the CPR treatment is a major reason for the variation in efficiency shown in Panel C of Figure 3.

As discussed earlier, adding nonlinearities to the payoff function makes the environment more complex, so it is less transparent to subjects what contributions correspond to self-interested and socially-efficient behavior. One implication of this reduced transparency is that punishment may not be applied to those who impose negative externalities (or fail to provide positive externalities) on others. Subject reactions to being punished may also be more confused in the more complex, nonlinear settings.

Result 4: (**Punishment**) As in the standard linear environment, subjects apply punishment to strong free riders, and subjects who are punished in the new nonlinear VCM and CPR environments respond in the subsequent period by free riding less. Free riders are punished significantly less, however, in the nonlinear environments.

Support: Figure 5 shows that subjects who deviate from the group norm (i.e., the group allocation chosen by the other group members) receive punishment points, and they change their contributions in response to receiving punishment. In order to compare the environments on a common metric, we employ a contribution index that is the difference between the chosen public allocation and the Nash equilibrium allocation, normalized by the difference between the Nash equilibrium and the social optimum for that treatment. This index is 1 at the social optimum and 0 at the Nash equilibrium in all treatments. It has a range outside [0, 1] for the nonlinear environments, especially for the CPR treatment because of the close proximity of the Nash equilibrium and social optimum in that environment.

The left side of the figure shows that when subjects' contribution index is below the group average index, they receive a greater number of punishment points; moreover, greater deviations from the group average receive more punishment points, especially in the linear VCM. That is, those who fail to provide a positive externality to others are punished. The right side of this figure shows that subjects on average raised their contribution index to increase positive externalities in the subsequent period in response to receiving punishment, indicating that most were not confused about why they were punished.

Table 3 reports results from a tobit regression model of the punishment points received, similar to those in the literature (e.g., Fehr and Gaechter, 2000, Table 5). It pools the data from the three treatments and provides statistical support for the conclusion that free-riders receive punishment in all three environments. Punishment levels are lower when the average contribution index for the others in the group is higher as there is less reason to punish. The three interaction terms confirm that negative (absolute) deviations from average group contribution index result in significantly greater punishment in all three treatments, as negative deviations from the group's average all result in positive and statistically significant punishment.[9] The magnitudes of the coefficients however indicate that punishment points received are higher for deviations in the linear VCM as compared to the other two treatments. An F-test confirms that the intensity of punishment significantly differs in the linear and nonlinear treatments.

Table 4 reports regressions to explore how subjects change contribution levels in the immediate subsequent period in reaction to being punished (as in Masclet et al. (2003) Table 4). The models are estimated separately for subjects who contributed above and below the mean contribution index for a given period. Both models show a strong negative relationship between a subject's deviation from the average contribution of her group, indicating a strong tendency to choose contributions closer to others' contributions. (This is frequently referred to as conditional cooperation.) Controlling for this movement towards the mean, punishment significantly raises contributions among low contributors in all three treatments.[10] Punishment does not statistically impact the change in contributions among high contributors.

[9] Our subjects do not exercise systematic anti-social punishment, i.e., sanction individuals who behave pro-socially, hence we do not include positive deviations from the group average as an explanatory variable.

[10] Although the response is larger for the CPR than the two VCM environments, as illustrated on the right side of Figure 5, this is likely due to the much larger range [-1.33, 2.67] and observed variance (0.37) of the contribution index in the CPR environment because of its small difference (3) between the Nash equilibrium and socially optimal allocations. By contrast, the range for this index in the linear VCM is [0, 1] with an observed variance of 0.10, and the range for this index in the nonlinear VCM is [-0.43, 1.29] with an observed variance of 0.18.

4.5 Profitable Deviations from Cooperative Benchmarks

As discussed in Section 3, the nonlinear environments increase complexity, since free riding is more difficult to identify, and Table 3 indicated that lower levels of cooperation were punished less severely in these environments. Compared to the two VCM settings, we also noted that the CPR environment has substantially stronger potential gains of defection from the socially-optimal level of contribution/extraction. As indicated in our final result, the observed punishment levels can eliminate most or even all of the gains from defection based on the level of cooperation that emerged in all environments.

Result 5: (**Profitable Deviations**) When including the average punishment costs received, deviations from the social optimum are potentially profitable only in the CPR environment. However coordinating at an extraction level that marginally exceeds the social optimum allows subjects to essentially eliminate these positive returns from defection.

Support: For the two VCM environments, as shown earlier in Figure 1 the greatest potential gain from an individual deviation from the social optimum is 6 or 7 experimental dollars. Subjects must contribute considerably less than others to maximize their own (one-period static) payoff when deviating from the social optimum. Such large deviations are punished severely, however, on average by a payoff reduction of about 20 experimental dollars. Therefore, as illustrated in Figure 6 on average they obtain a post-punishment *reduction* in net payoff of 10 or more experimental dollars when deviating to full free riding. The figure indicates that all deviations from the social optimum are on average not profitable in the two VCM environments after subtracting received punishment costs.

Gains from defection are larger in the CPR, and the punishment assigned to free-riders is not sufficient to offset the payoffs from deviating from the social optimum. For example, consider a subject who increases extraction from the cooperative level of 5 to 10. The figure indicates that a subject could increase his net payment by over 10 experimental dollars on average. This makes cooperation to restrict extraction to 5 units per subject difficult to sustain. Hence we would expect subjects to deviate from the social optimum.

As shown earlier in Figure 2 (Panel C), however, mean extraction levels exceed 5 even when punishment is possible. In particular, for the CPR treatment with punishment, only 10 percent of the individual extraction levels were at the social optimum of 5. The most frequent extraction was 6, chosen 35 percent of the time, and the modal extraction was 6 in every period. The data do suggest some failed attempts to coordinate on the social optimum of 5, since the fraction of choices on 5 falls in half from periods 1-4 to periods 5-10. Although cooperation on an extraction level of 6 rather than 5 is suboptimal, it permits subjects to earn nearly the same level of profit (50.4) as the 52 earned at the social optimum. Importantly, this suboptimal extraction level of 6 substantially reduces the incentives to deviate, typically by 10 to 12 experimental dollars. Figure 6 shows that the gains from deviating from this suboptimal cooperation level of 6 are barely positive.

5. Conclusion

This paper explores the effectiveness of peer punishment in promoting cooperation in social dilemmas with nonlinear environments. While a large literature exists on peer punishment in public good games, nearly all of these studies consider the special context of linear VCMs. Nonlinear environments are at least if not more relevant in field applications. The interior solutions in nonlinear environments make it more difficult for agents to identify the equilibria and the gains to cooperation through peer punishment. By studying the impact of peer punishment in this largely neglected class of social dilemma games, our paper significantly strengthens the confidence in key conclusions drawn from this literature and provides a novel contribution that helps enhance the external validity of experimental results on informal peer sanctions.

We can highlight four key results from the experiment reported here. First, the availability of punishment opportunities does increase cooperation in nonlinear VCM and CPR environments. This finding for the CPR stands in contrast to Ostrom et al. (1992), which concluded that communication between subjects was also necessary to improve cooperation. That early study employed a considerably more restrictive form of peer punishment relative to the subsequent literature. Punishment effectiveness is low in our data, however, and outcomes are closer to the Nash equilibrium than the social optimum even when agents are allowed to punish their peers. In addition, in the nonlinear VCM, contributions do not increase significantly when considering all time periods. The impact of peer punishment takes longer to appear and is only observed in later periods, even though nonlinearity was introduced in the simplest possible way. This suggests that long-run efficiency benefits of punishment (Gaechter et al., 2008) may be delayed or never realized with certain nonlinearities. The nonlinear VCM is the only treatment with nonlinear private returns, indicating that nonlinear private returns may make the benefits of cooperation more difficult to understand.[11] In such cases, social norms promoting cooperation may need to be supplemented by external private incentives. In the academic refereeing example above, this could correspond to monetary incentives provided by some journals or excellence in refereeing awards to encourage timely and good quality reports.

[11] In a related study that compares the relative efficacy of cheap talk communication and peer punishment, in Cason and Gangadharan (2013) we conduct a new experiment to study a more complex environment, also with nonlinear private returns, and find that punishment opportunities in isolation have no significant impact on cooperation levels.

Second, similar to the linear VCM treatment, in our new nonlinear treatments the increase in cooperation does not generate a high enough return to compensate for the costs of punishment. Efficiency is therefore not improved by the availability of punishment in all environments studied here for this ten-period horizon. Third, as is also seen in the linear VCM, in the new environments individuals punish strong free riders, however the intensity of punishment assigned is significantly lower in the nonlinear treatments compared to the linear VCM. This suggests that punishment is affected by the increased complexity due to moving cooperative and equilibrium actions to the interior of the choice space. Finally, gains from defecting are higher in the CPR treatment and punishment levels observed are not sufficient to support the socially optimal level of extraction. Subjects often implement suboptimal levels of extraction, however, which reduce the gains from defection so that cooperation can be enforced effectively through peer punishment. This significantly reduces a key structural difference (payoffs from deviation) between the rival CPR and the nonrival VCM environments.

Overall, these results demonstrate that while peer punishment has similar influence on outcomes in the nonlinear environments, its impact is reduced and takes longer to emerge. This suggests that the effectiveness of peer punishment could vary in different environments, particularly in situations with more complexity that make it more difficult to identify defectors from the social norm.

References

Bracha, A., M. Menietti and L. Vesterlund. 2011. "Seeds to Succeed? Sequential Giving to Public Projects." *Journal of Public Economics*, 95, 416-427.

Casari, M. and C. Plott. 2003. "Decentralized Management of Common Property Resources: Experiments with a Centuries-Old Institution." *Journal of Economic Behavior and Organization*, 51, 217-247.

Cason, T. and L. Gangadharan. 2013. "Swords without Covenants do not lead to Self-Governance." Purdue and Monash Universities working paper.

Chaudhuri, A. 2011. "Sustaining Cooperation in Laboratory Public Goods Experiments: A Selective Survey of the Literature." *Experimental Economics*, 14, 47-83.

Fischbacher, U. 2007. "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics*, 10, 171-178.

Fehr, E. and Gaechter, S. 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, 90, 980-994.

Gaechter, S., E. Renner and M. Sefton. 2008. "The Long-Run Benefits of Punishment." *Science*, 5 December, 322, 1510.

Greiner, B. 2004. "An Online Recruitment System for Economic Experiments." in K. Kremer and V. Macho, eds., *Forschung und wissenschaftliches Rechnen 2003*. Göttingen: Ges. Für Wiss. Datenverarbeitung. 79-93.

Isaac, R.M. and J. Walker. 1988. "Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism." *Quarterly Journal of Economics*, 103, 179-200.

Isaac, R.M. and J. Walker. 1998. "Nash as an Organizing Principle in the Voluntary Provision of Public Goods: Experimental Evidence." *Experimental Economics*, 1, 191-206.

Janssen, M., R. Holahan, A. Lee and E. Ostrom. 2010. "Lab Experiments for the Study of Social-Ecological Systems." *Science*, 30 April, 328, 613-617.

Maier-Rigaud, F. and J. Apesteguia, 2006. "The Role of Rivalry Public Goods versus Common-Pool Resources", *Journal of Conflict Resolution*, 50(5), 646-663.

Masclet, D., C. Noussair, S. Tucker and M.-C. Villeval. 2003. "Monetary and Non-Monetary Punishment in the Voluntary Contributions Mechanism." *American Economic Review*, 93, 366-380.

Murphy, R., A. Rapoport and J. Parco. 2006. "The Breakdown of Cooperation in Iterative Real-Time Trust Dilemmas." *Experimental Economics*, 9, 147-166.

Nikiforakis, N., C. Noussair, and T. Wilkening. 2012. "Normative Conflict & Feuds: The Limits of Self-Enforcement." *Journal of Public Economics*, 96, 797-807.

Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action.* New York. Cambridge University Press.

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.

Ostrom, E., R. Gardner and J. Walker. 1994. *Rules, Games, and Common-Pool Resources*. Ann Arbor, MI: University of Michigan Press.

Reuben, E. and A. Riedl. 2013. "Enforcement of Contribution Norms in Public Good Games with Heterogeneous Populations." *Games and Economics Behavior*, 77, 122-137.

Sefton, M. and R. Steinberg. 1996. "Reward Structures in Public Good Experiments." *Journal of Public Economics*, 61, 263-287.

Smith, V. L. 1989. "Theory, Experiment and Economics, *Journal of Economic Perspectives*, 3, 151-169.

van Soest, D. and J. Vyrastekova. 2006. "Peer Enforcement in CPR Experiments: The Relative Effectiveness of Sanctions and Transfer Rewards, and the Role of Behavioural Types." In J. List (ed.), *Using Experimental Methods in Environmental and Resource Economics*. Cheltenham, UK: Edward Elgar, 113-136.

Xiao, E. and H. Kunreuther. 2013. "Punishment and Cooperation in Stochastic Social Dilemmas". NBER working paper 18458.

Yamagishi, T. 1986. "The Provision of a Sanctioning System as a Public Good". *Journal of Personality and Social Psychology*, 51(1), 110-116.

Table 1: Features of the Experimental Environments

	Linear VCM	Nonlinear VCM	Common Pool Resource
Private Return (α)	1 per unit	5 for first 2 2 for 3 to 9 0.8 for 10 to 12	2 per unit
Group Return (β)	0.5	1	$(g_i/G)(18G\text{-}0.4G^2)$
Extra Payment (C_i)	20	0	-12
Nash & Social Optimum	Boundary	Interior	Interior
Nash choice g_i^*	0	3	8
(Payoff)	(32)	(36)	(37.6)
Social Optimum $g_i^{s.o.}$	12	10	5
(Payoff)	(44)	(50)	(52)
% gains to cooperation	37.5%	38.9%	38.3%
Maximum Payoff Gain from Defecting from Social Optimum	6	7	22.4
Groups/Subjects	20/80	24/96	24/96

Common features: n=4 individuals per group; 10 periods of repetition in fixed groups with known endpoint; strategy space tokens to allocate $g_i \in \{0, 1, \ldots, 11, 12\}$; all 3 others in group could be punished, assigning up to 5 deduction points to each; each point cost the assigner 1 E$ and reduced the recipient's payoff by 3 E$.

Table 2: Average Contributions to Group Account

	Linear VCM		Nonlinear VCM		CPR	
	Periods 1-10	Periods 6-10	Periods 1-10	Periods 6-10	Periods 1-10	Periods 6-10
No Punishment	5.23 (0.53)	4.87 (0.80)	5.69 (0.44)	4.93 (0.45)	7.29 (0.21)	7.59 (0.18)
Punishment	7.38 (0.47)	8.73 (0.71)	6.47 (0.54)	6.77 (0.67)	6.80 (0.14)	6.97 (0.20)
Wilcoxon rank-sum p-value	0.023	0.007	0.184	0.046	0.046	0.057

Note: Standard errors shown in parentheses. Wilcoxon test p-values are for two-tailed tests.

Table 3: Tobit Model of Punishment Points Received

	Linear VCM
Constant	-2.58**
	(0.91)
Others average contribution index	-5.73**
	(1.47)
Negative (absolute) deviation from average contribution index X Linear VCM indicator	32.37** (3.55)
Negative (absolute) deviation from average contribution index X Nonlinear VCM indicator	24.11** (4.13)
Negative (absolute) deviation from average contribution index X CPR indicator	9.80** (1.71)
Number of Observations	1360
Pseudo R^2	0.06
F-test (2, 1356 d.f.) for treatment differences	19.36**

Notes: Clustered standard errors are robust to unspecified correlation within subjects and are shown in parentheses. ** denotes significantly different from zero at 1 percent level (all two-tailed tests). X denotes that the two variables are interacted.

Table 4: Determinants of Changes in Contribution Index

Dependent Variable: contribution index$_{it+1}$ – contribution index$_{it}$

	Low Contributors (below average)	High Contributors (above average)
Constant	0.011	0.020
	(0.029)	(0.024)
Punishment points received in period *t* X Linear VCM indicator	0.011** (0.003)	-0.003 (0.004)
Punishment points received in period *t* X Nonlinear VCM indicator	0.012** (0.004)	-0.002 (0.006)
Punishment points received in period *t* X CPR indicator	0.038** (0.009)	0.009 (0.012)
Deviation from average group contribution index in period *t*	-0.303** (0.090)	-0.734** (0.137)
Number of Observations	537	687
R^2	0.20	0.17
F-test for treatment differences	5.10**	0.61

Notes: Clustered standard errors are robust to unspecified correlation within subjects and are shown in parentheses. ** denotes significantly different from zero at 1 percent level (all two-tailed tests). X denotes that the two variables are interacted.

Figure 1: Payoff Gain from Defecting from Social Optimum

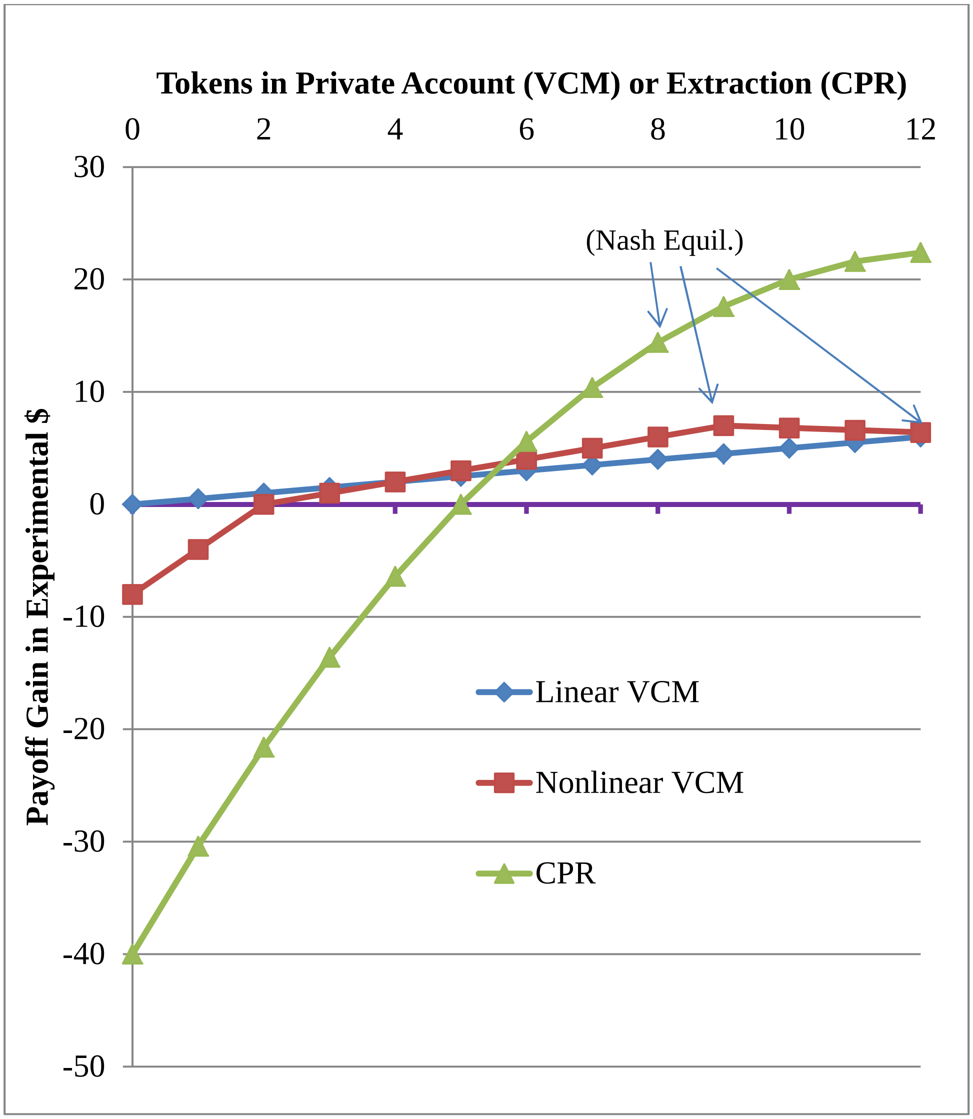

Figure 2: Average Contributions to the Group Account in Each Treatment

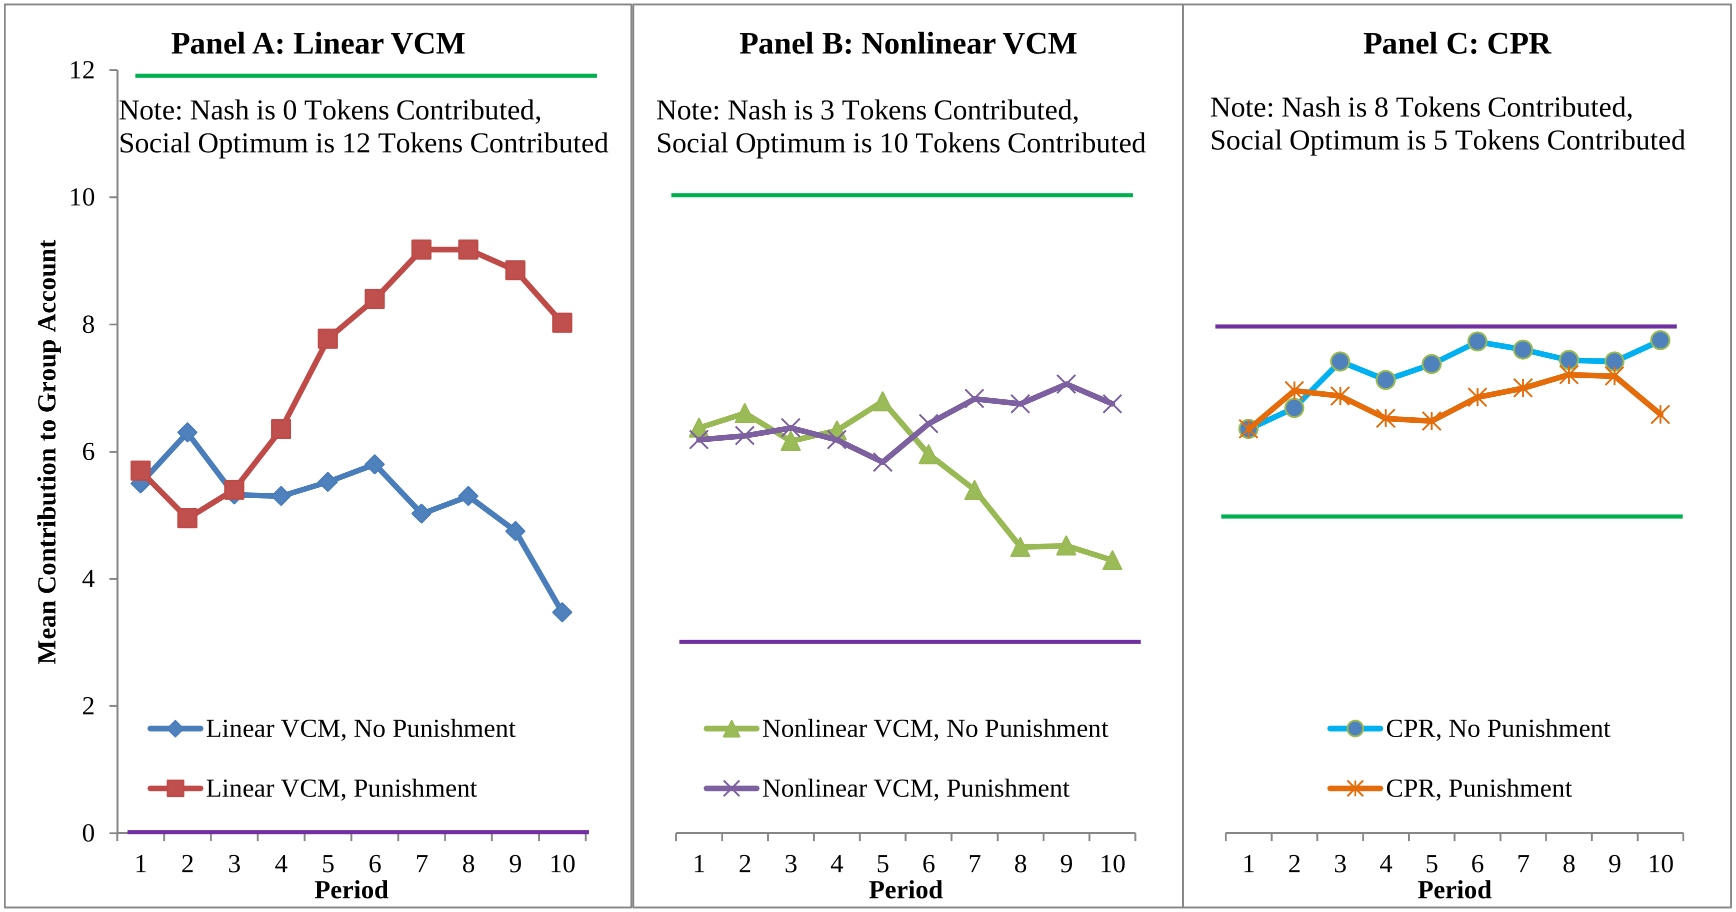

Figure 3: Average Efficiency Gains Relative to the Nash Equilibrium in Each Treatment

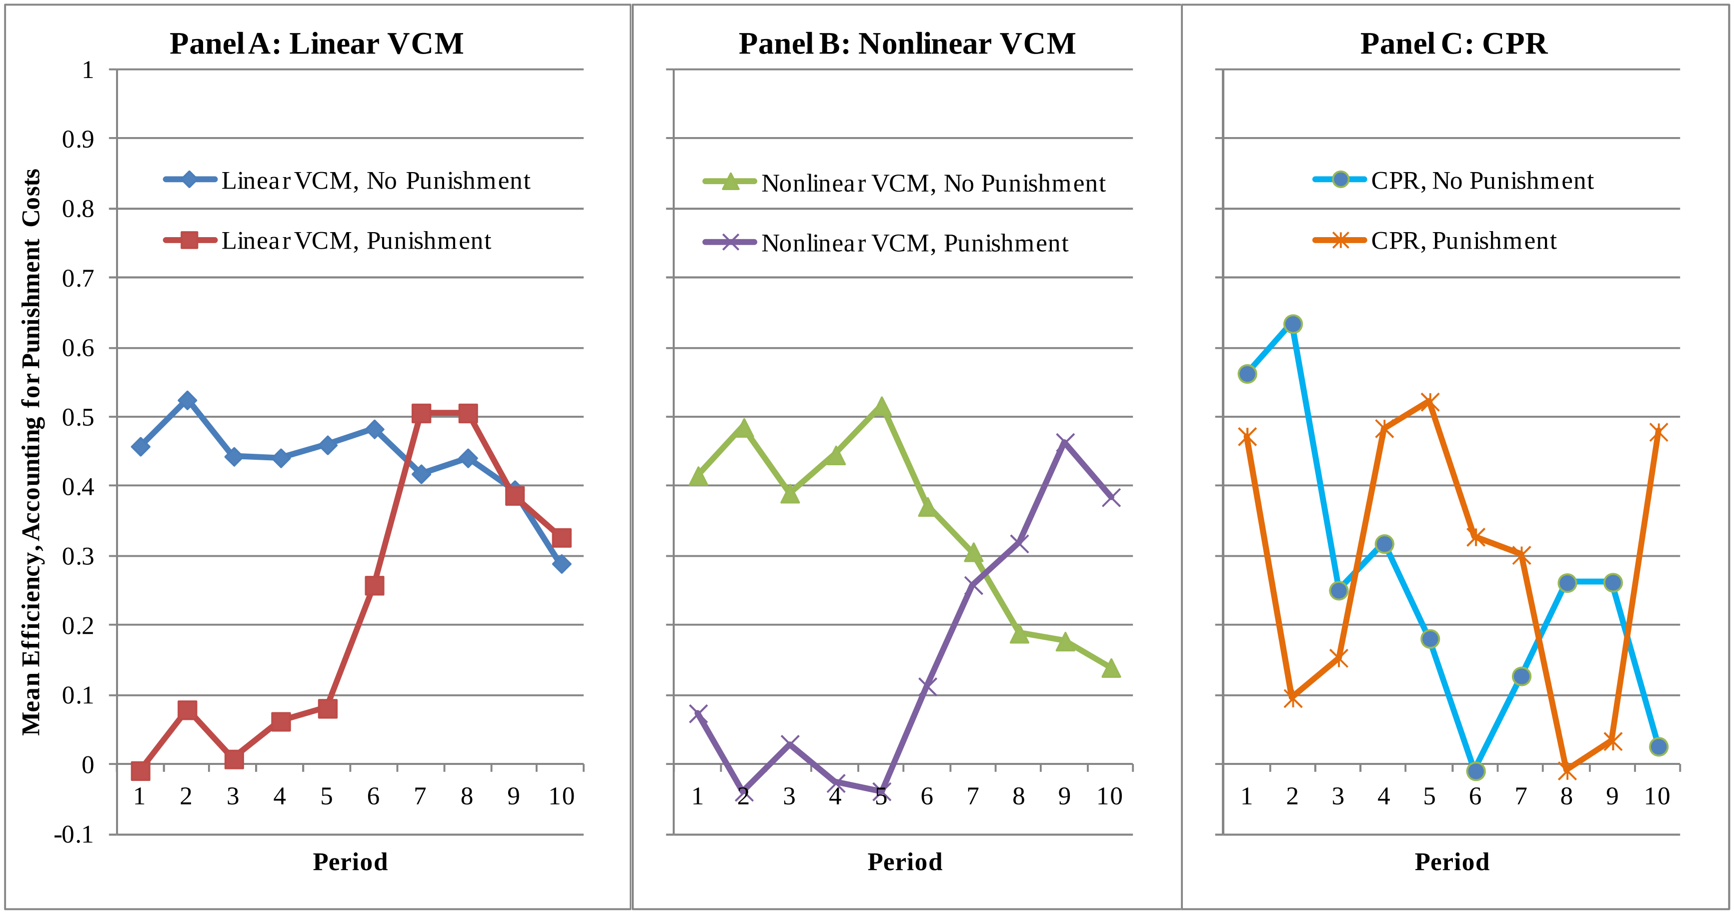

Figure 4: Average Punishment Points Assigned per Subject

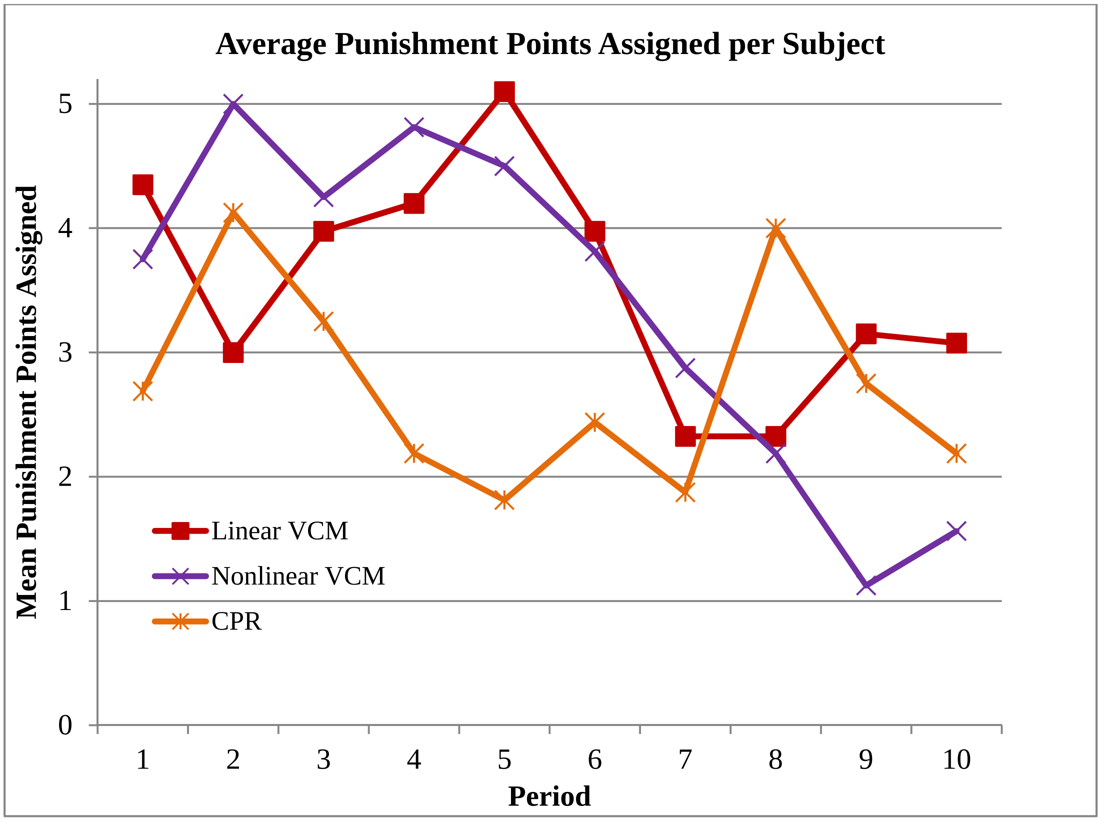

Figure 5: Punishment Points Received and Average Reaction

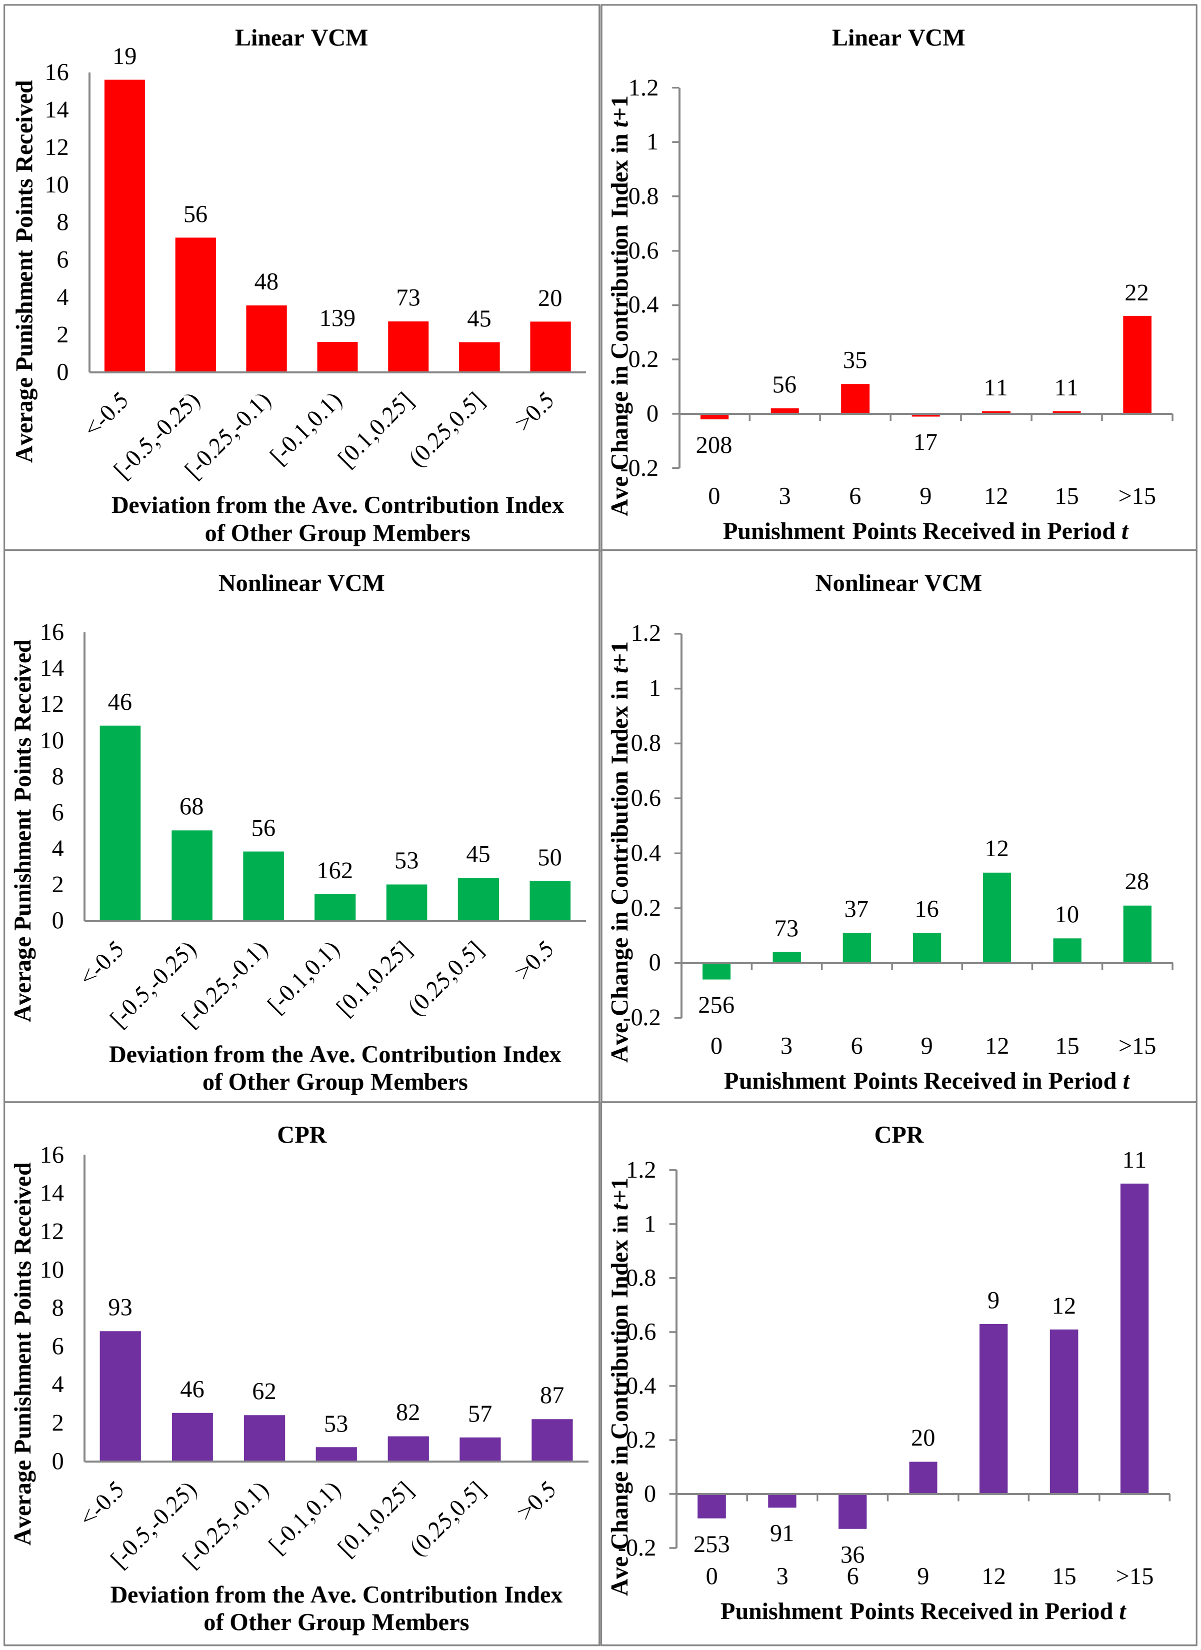

Figure 6: Payoff Gain from Defecting from Social Optimum, Including Average Punishment Received

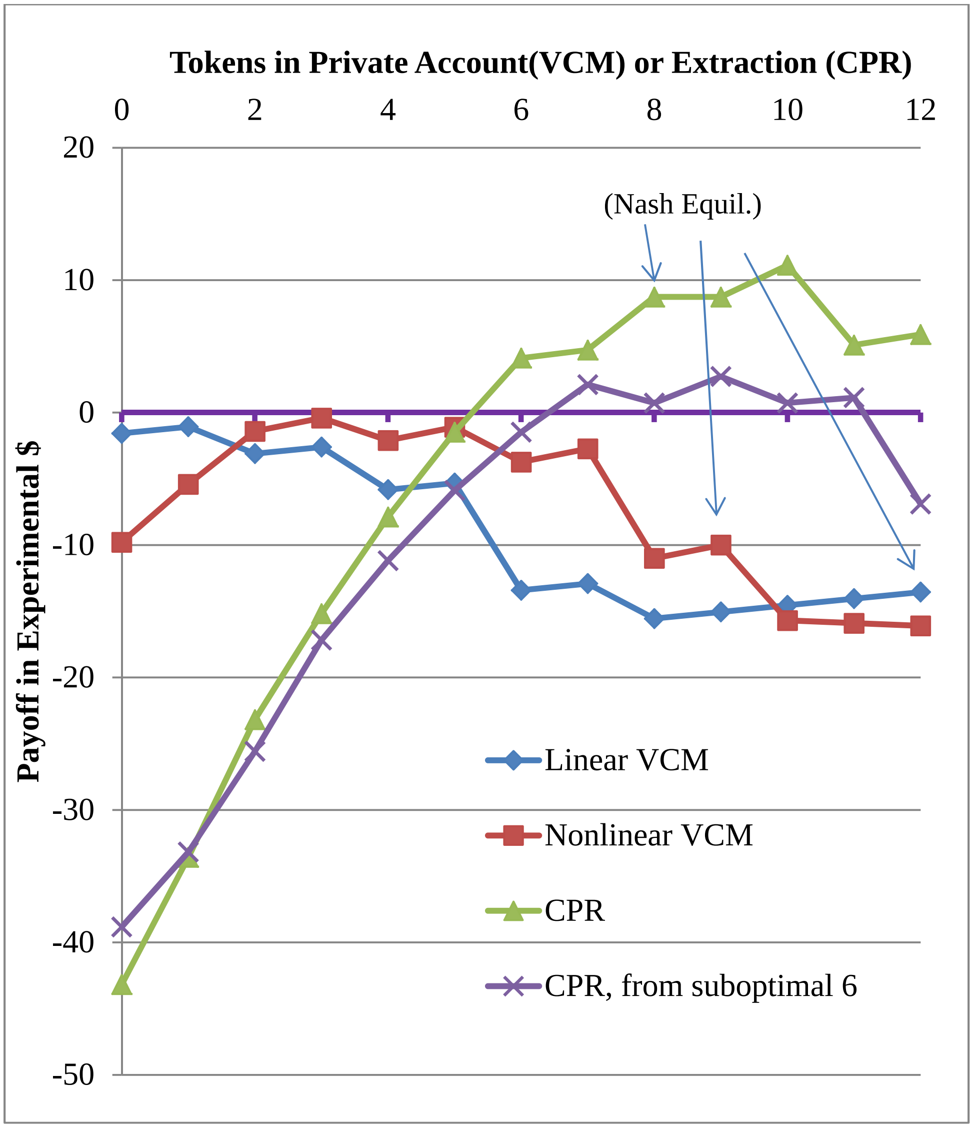

Appendix: EXPERIMENT INSTRUCTIONS (CPR with Punishment)

This is an experiment on decision making. If you read the following instructions carefully, you can, depending on your decisions, earn a considerable amount of money. All earnings on your computer screens are in Experimental Dollars. These Experimental Dollars will be converted to real Dollars at the end of the experiment, at a rate of _____ Experimental Dollars = 1 real Dollar.

Today's session will be conducted using the computer network located here in this laboratory. It will be divided into 10 different periods. Attached to these instructions you will find a sheet labeled Personal Record Sheet, which will help you keep track of your earnings based on the decisions you might make. You are not to reveal this information to anyone. It is your own private information.

You have been assigned to a group of four (yourself and three other) participants. This will be your group for the entire session.

The First Stage

At the beginning of each period each participant receives an endowment of 12 tokens. In **Stage 1** each period you (and the others in your group) must decide how many tokens to place into either or both of 2 accounts: a private account and a group account. All tokens must be placed in one account or the other. Each token you place in the private account generates a return to you (independent of what anyone else does), and each token you place in the group account generates a return that depends on how many tokens that others in your group place in the group account. Your earnings in a period are the sum of your earnings from the private account and your earnings from the group account, minus 12. Returns to the two accounts are listed on your input screen as shown on the next page. Everybody has the same returns.

You and all the other members of your group will each get a share of the total group earnings that depends on your token placements. If a total of X tokens are placed in the group account by group members, then the total group payoff is $18X-0.4X^2$ Experimental Dollars. For example, if a total of 10 tokens are placed in the group account, the total group payoff is $(18\times10)-(0.4\times100)=180-40=140$. This amount, along with every other possible total token placement in the group account are shown on your input screen below.

Your share of this total group payoff equals your number of tokens placed in the group account as a fraction of the total tokens you and the others in your group place in the group account. For example, if a total of 10 tokens are placed in the group account, and you placed 2 of these 10 tokens, then you receive 2/10=0.2 of the 140 total group payoff, or 0.2×140=28.

Your private account generates a return to you that depends only on your tokens placed in the private account. In particular, you will receive 2 Experimental Dollars for each token that you place in your private account.

You will indicate your decisions on the input-screen for the first stage:

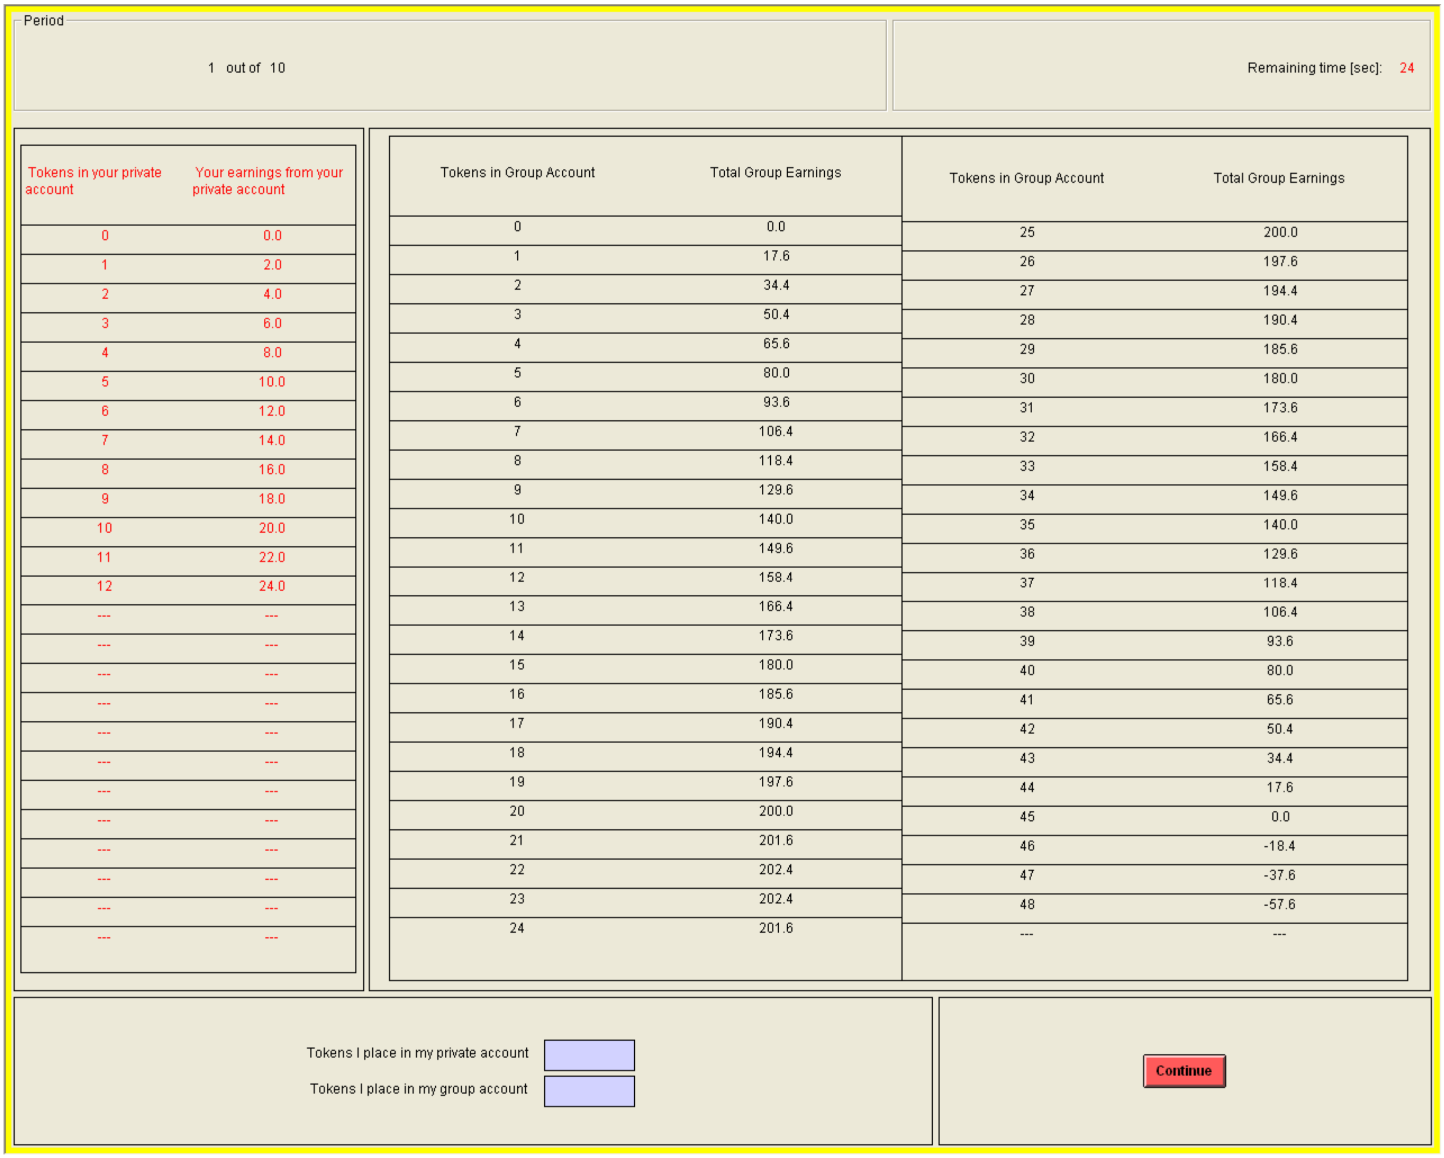

Period

1 out of 10

Remaining time [sec]: 24

Tokens in your private account	Your earnings from your private account
0	0.0
1	2.0
2	4.0
3	6.0
4	8.0
5	10.0
6	12.0
7	14.0
8	16.0
9	18.0
10	20.0
11	22.0
12	24.0
---	---
---	---
---	---
---	---
---	---
---	---
---	---
---	---
---	---
---	---
---	---
---	---

Tokens in Group Account	Total Group Earnings	Tokens in Group Account	Total Group Earnings
0	0.0	25	200.0
1	17.6	26	197.6
2	34.4	27	194.4
3	50.4	28	190.4
4	65.6	29	185.6
5	80.0	30	180.0
6	93.6	31	173.6
7	106.4	32	166.4
8	118.4	33	158.4
9	129.6	34	149.6
10	140.0	35	140.0
11	149.6	36	129.6
12	158.4	37	118.4
13	166.4	38	106.4
14	173.6	39	93.6
15	180.0	40	80.0
16	185.6	41	65.6
17	190.4	42	50.4
18	194.4	43	34.4
19	197.6	44	17.6
20	200.0	45	0.0
21	201.6	46	-18.4
22	202.4	47	-37.6
23	202.4	48	-57.6
24	201.6	---	---

Tokens I place in my private account

Tokens I place in my group account

Continue

Since your endowment each period is 12 tokens, the two numbers you indicate on your input-screen must be whole numbers between 0 and 12 and must sum to 12. After entering your decision you must press the **Continue** button. Once you have done this, your decision has been made and cannot be changed.

After all participants in your group have made their decisions the following income screen will show you the total amount of tokens placed in the group account by all four participants of your group (including you). Also this screen shows your earnings for the first stage of the period. Your earnings are the sum of your earnings from the private account and your earnings from the group account. During the experiment you will record this information on your hardcopy record sheet and then click the **Continue** button.

Detail of Results Screen for First Stage

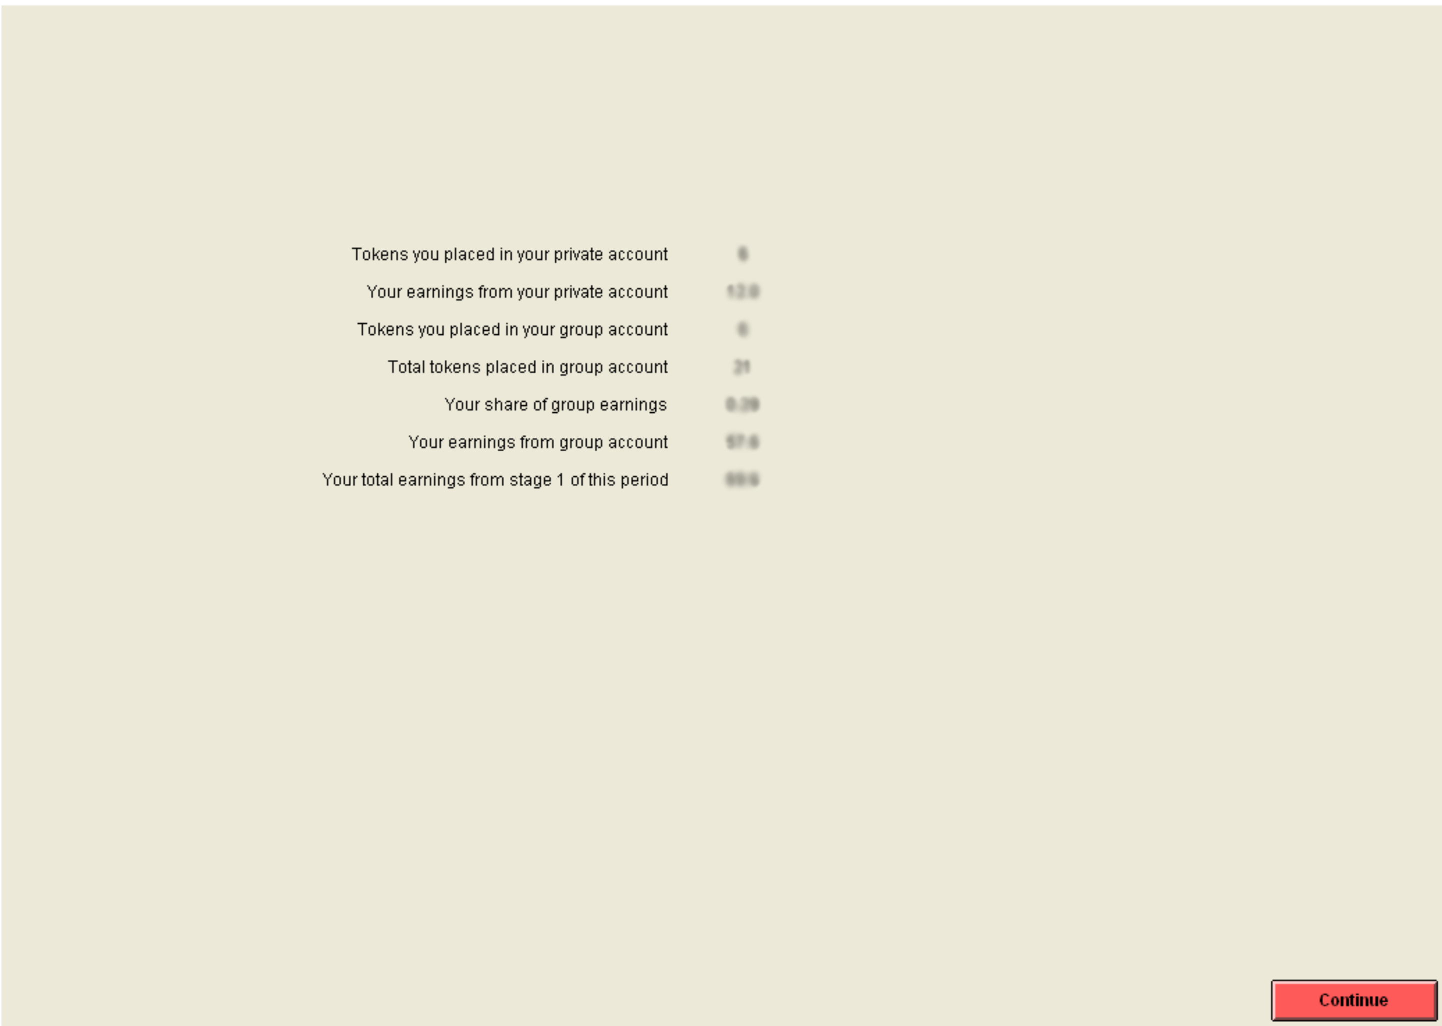

The Second Stage each Period

In the second stage you will see the amount of tokens placed in the group account by all four participants of your group. Moreover, in this stage you can decide whether to **decrease** the earnings of these others in your group by assigning **deduction points**. These other participants can also decrease your earnings if they wish to. This is apparent from the input screen at the second stage, shown below.

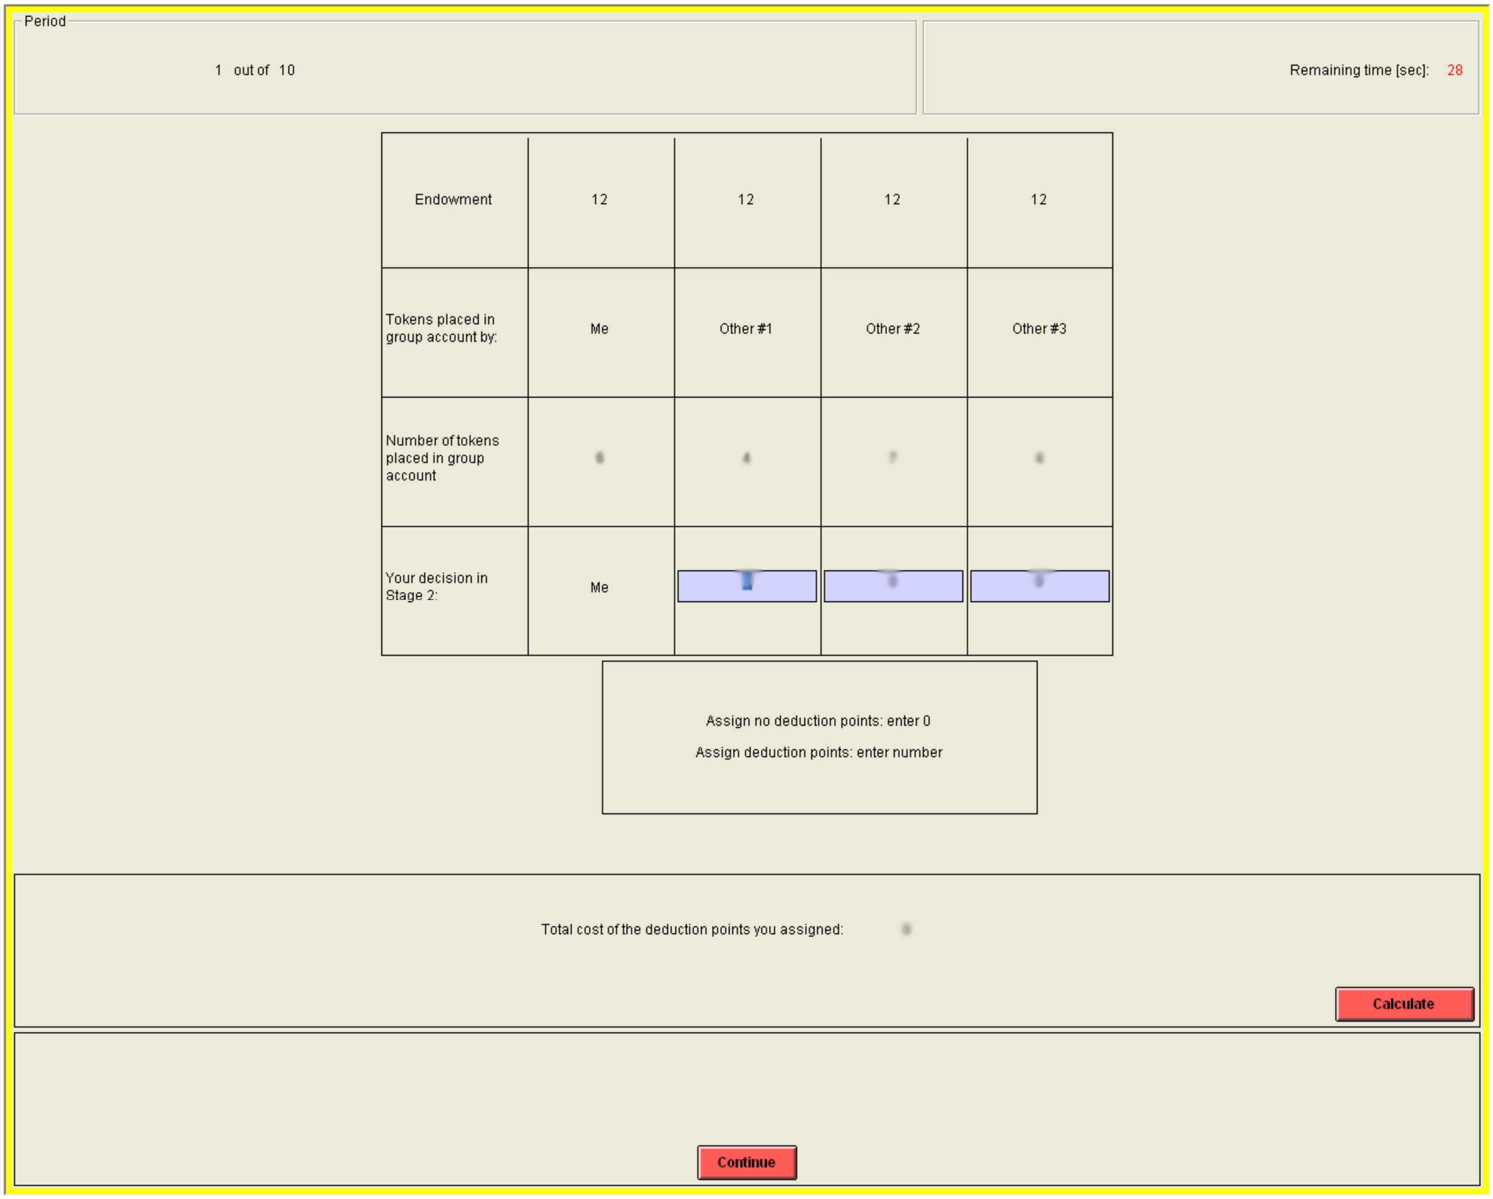

Your allocation to the group account is displayed **in the first column**, while the allocations to the group account **by the other people** are shown in the **remaining three columns**. Note that the order in which others' allocations are displayed will be determined at random in every period. The allocation in the second column, for example, could represent a different person in different periods. The same holds true for the other two columns.

You will have to decide how many deduction points to assign to **each** of these other three participants in your group. You must enter a number for each of them. If you do not wish to change the earnings of a person in your group then you must enter 0. You can **assign up to 5 points to each participant**.

You will incur costs from assigning deduction points. Every deduction point you assign costs you 1 Experimental Dollar. For example, if you assign 2 deduction points to one person, this costs you 2 Experimental Dollars; if, in addition, you assign 4 deduction points to another person this costs you an additional 4 Experimental Dollars. In total for this example you will have assigned 6 points and your **total costs** therefore amount to 6 Experimental Dollars.

After you have assigned points to each of the other three participants you can click the button "**Calculate**" (see the second stage input screen). On the screen you will then see the total costs of your assigned points. As long as you have not yet clicked the **Continue** button, you can still change your decision. To recalculate the costs after a change of your assigned points, simply press the "Calculate" button again.

If you assign 0 deduction points to a particular participant (i.e., enter "0"), you will not alter his or her earnings. However, if you assign **one deduction point** to a participant you will **decrease** his or her earnings by **3 Experimental Dollars**. If you assign a participant **2 deduction points** you will **decrease** his or her earnings by **6 Experimental Dollars**, and so on. Each deduction point that you assign to another person will reduce his or her earnings by 3 Experimental Dollars. Similarly, each deduction point assigned to you by another participant will reduce your first stage earnings by 3 Experimental Dollars:

Costs of received deduction points = 3 × Sum of received deduction points.

How much the earnings at the second stage are decreased depends on the sum of deduction points received. For instance, if somebody receives **a total** of **3 deduction points** (from all other participants in this period), his or her earnings would be decreased by **9 Experimental Dollars**. If somebody receives a total of **4 deduction points,** his or her earnings are reduced by **12 Experimental Dollars**.

Your total earnings from the two stages are therefore calculated as follows:

Total earnings at the end of the second stage = period earnings =
= Earnings from the first stage – 3 × (sum of received deduction points)
– (sum of deduction points you have assigned) – 12

Note that everyone has a fixed amount of 12 Experimental Dollars subtracted every period.

After all participants have made their decision, your earnings from the period will be displayed on a screen such as the one shown below. After you have viewed the earnings screen the period is over and the next period commences. Recall that 10 periods will be conducted.

Recording Rules

During every period you should write down the information shown on your results screens on your Personal Record Sheet. The Stage 1 results screen shows the group account allocations you should record, and the final earnings screen shows your Deduction Points and period earnings. Be sure to record your total earnings for each period in the rightmost column.

Summary

1. All subjects use the same Earnings Tables.
2. In each period, you and every other participant will each have 12 tokens to allocate.
3. In each period, you will decide how many tokens to place in your private account and how many to place in your group account. You must allocate all 12 tokens each period.
4. Your earnings from the private account depend only on your decision about how many tokens to place in this account.
5. Your earnings from the group account depend upon how many tokens you and the other three participants of your group place in this account. You receive a share of group earnings that depends the fraction of the total tokens in the group account that you placed in the group account.
6. You may assign up to 5 Deduction Points to each of the other individuals in your group. Each Point you assign costs you 1 Experimental Dollars.
7. These other individuals can assign Deduction Points to you. Each Point assigned to you reduces your earnings by 3 Experimental Dollars.
8. You will interact with the same 3 other individuals for all decision periods.
9. Results and earnings should be recorded on your Record Sheet at the end of each period.

Detail of Earnings Screen at the end of the Second Stage:

Personal Record Sheet

Period Number	My Tokens in Private Account	My Earnings from Private Account	My Tokens in Group Account	Total Tokens in Group Account	My Share of Group Earnings	My Earnings from Group Account	My Total Earnings from Stage 1	Amount of Received Deduction Points	Payoff Reduction through Deduction Points	Cost of Assigning Deduction Points	Payoff for the Period	Cumulative Payoff so far
1												
2												
3												
4												
5												
6												
7												
8												
9												
10												