

The Dose Does it: Punishment and Cooperation in Dynamic Public-Good Games

Bettina Rockenbach¹ and Irenaeus Wolff^{2*}

¹*University of Cologne, Albertus Magnus Platz, Cologne, Germany*

²*University of Konstanz/TWI, Kreuzlingen, Switzerland*

ABSTRACT

We experimentally study the role of punishment for cooperation in dynamic public-good problems where past payoffs determine present contribution capabilities. The beneficial role of punishment possibilities for cooperation is fragile: successful cooperation hinges on the presence of a common understanding of how punishment should be used. If high-contributors punish too readily, the group likely gets on a wasteful path of punishment and retaliation. If punishment is administered more patiently, even initially uncooperative groups thrive. Hence, when today's punishment also determines tomorrow's cooperation abilities, it seems crucial that groups agree on the right 'dose' of sanctions for punishment to support cooperation.

Keywords: Cooperation, Dynamic game, Punishment, Retaliation, Endowment endogeneity, Experiment

1 Introduction

An extensive literature has shown that opportunities for peer-punishment boost cooperation in social-dilemma settings (e.g., Fehr *et al.*, 2002; Gintis *et al.*, 2005;

*We would like to thank the research group led by Urs Fischbacher, seminar audiences at the University of Erfurt, the 2010 annual meetings of the Association for Public Economic Theory (Istanbul) and the Economic Science Association (Copenhagen) as well as three anonymous referees for their helpful comments. We thank Özgür Gürerk for his contribution at an early stage of the project, in particular in the design phase. Irenaeus Wolff did most of the data analysis and the derivation of the equilibria under social preferences. In all other parts, the authors contributed equally. We have no competing interests. We gratefully acknowledge the financial support by the DFG through research grant RO 3071.

Henrich *et al.*, 2004; Sigmund, 2007; Mathew and Boyd, 2011). The underlying (theoretical and experimental) models typically assume (repeated) one-shot interactions which implies that present actions impact neither future cooperation capabilities nor future punishment capabilities. Yet, in many instances, both the present cooperation level and the present punishment expenses will have repercussions on the future. Obvious examples come from ‘primitive’ societies: cooperatively building a boat will increase the future catch potential. Likewise, if a hunter in a hunter-gatherer society takes a stick and ambushes one of the fellow-hunters to punish him for his lack of effort during that day’s battle with the neighboring horde, this will lower the non-cooperator’s energy level available for the next-day hunt as well as the punisher’s. If, then, the next day’s prey is smaller because of these repercussions, and if the resulting lower food intake further weakens the hunters, what looked like a short-time punishment action may have a serious negative long-term effect on fitness.¹ If, finally, as in the typical public-good experiment with punishment opportunities, punishment is particularly heavy during early interaction periods, it could well be that these dynamic effects annihilate any positive effects of longer time horizons pointed out by Gächter *et al.* (2008). On the other hand, if the hunters shy away from using punishment completely to avoid the negative dynamic effects, cooperation may also break down in the same way as it does in the vast experimental literature on repeated public goods without further means for disciplining free-riders. Although this temporal interdependence in cooperation and punishment seems of crucial importance for understanding the evolution of human cooperation, it has been largely ignored in the literature so far.

In this study, we aim at closing this gap by examining a dynamic public-good problem allowing for peer punishment. To do so, we compare behavior in a dynamic game allowing for peer punishment (DYNPUN) to a control treatment that does not allow for punishment (DYNNOPUN). Both treatments implement a dynamic version of a typical 4-player public-good game played over 20 rounds. The crucial difference to the usual repeated one-shot public-good game is that present contribution capabilities are not constant but depend on earlier play. Broadly speaking, participants’ earnings at the end of a round are their endowment in the next round. In the DYNPUN game, a punishment stage using a convex punishment technology follows each contribution stage.²

¹Another example stems from the political arena: when countries cooperate to ban nuclear weapons, the absence of nuclear weapons and the savings on military expenditures typically benefit all of the parties, thus increasing present as well as future resources. Likewise, if one of the countries faces economic sanctions for trying to develop nuclear capabilities, the sanctions will typically lead to lower available resources for future cooperative endeavors in the sanctioning countries but in particular in the sanctioned country. So again, short-time punishment actions may have negative long-term effects, for example, if cooperative acts, like e.g., environmental-protection efforts, are cancelled because of a lack of available resources.

²We chose a convex punishment technology to reflect a certain power-asymmetry under heterogeneous wealth (or fitness) levels: in real-life environments, it is increasingly difficult

We show that only few groups succeed in achieving long-term cooperation. In DYNPUN, the key for success seems to lie in applying the right dose of punishment. Successful groups are characterized by administering punishment cautiously, and by punished players not reacting with counter-punishment. When low contributions are heavily punished in the first round, punishment is often retaliated against straight away and long-term performance is low. Yet, when the available punishment option is not used at all to stop non-cooperation, free-riding spreads and long-term performance is also low. Remarkably, in DYNPUN, success or failure of a group cannot be predicted by the initial contribution average or contribution variance. What seems to make the difference is that in those DYNPUN groups who succeed, right from the start there seems to be a common understanding (or “social norm”) of when and how punishment should be used (see Nicklisch and Wolff, 2011; Reuben and Riedl, 2013). Only when punishment is not possible (in DYNNOPUN), initial contribution levels have strong predictive power for overall success.

1.1 Literature on Dynamic Public-Goods

There is a substantial theoretic literature on dynamic social-dilemma games with earlier play influencing later distributions of different (player) types in evolutionary settings.³ There also is a huge experimental literature on repeated public-good games.⁴ However, until recently, experimental studies on social dilemmas focusing on truly dynamic games were surprisingly limited.⁵

Noussair and Soo (2008) study public-good provision when the group’s past cooperation level influences each member’s current marginal per-capita return of provision. This resembles a situation in which players’ abilities to contribute to a public good are unrelated to the payoff stemming from it, but the more cooperative the group has been in the past the higher is the return from future cooperation. In their setting, contribution levels generally do not exhibit the usual falling trend except for a minority of the groups. Sadrieh and Verbon (2006) consider a situation in which a group member’s benefit from the

(or risky) for a weak individual to punish an experience-hardened, strong, and powerful individual. In using a convex punishment technology, we follow the example of studies like Fehr and Gächter (2000), Denant-Boemont *et al.* (2007), or Nikiforakis (2008). In a recent study, Gächter *et al.* (2017) use a setup similar to ours but with a linear punishment technology and a shorter time-horizon, obtaining similar results.

³For examples of the evolutionary settings, cf. e.g. Richerson and Boyd (2005) and the many references cited therein. For a game-theoretic treatment of a differential-game dynamic public good, see Fershtman and Nitzan (1991). Admati and Perry (1991) analyze a two-player step-level public good with alternating contribution stages.

⁴See, e.g., the literature reviewed in Ledyard (1995) and Chaudhuri (2011).

⁵An early exception is Rapoport (1988) who examines how different information conditions affect behavior in a common-pool-resource game with depletion. He observes little cooperation, even though communication of an optimal strategy did have a positive effect on harvesting behavior.

public good depends on the player’s current wealth. This setup is well-tailored to their focus on inequality and situations prone to the accentuation of this inequality. Subjects’ propensity to cooperate is not affected by the degree of inequality induced. In contrast, in a control treatment that does not involve a dynamic component, induced inequality has a positive effect on cooperation. They conclude that participants’ fairness concerns seem to be ‘crowded out’ by the introduction of the dynamics.

Grosse (2011) studies dynamic public-good provision in a setting similar to our DYNNOPUN treatment. He analyzes the role of the dynamics and separates the effects of inequality, rising endowments, and endogeneity of evolving wealth levels. Unlike in the study of Noussair and Soo (2008), groups in the main treatment of the study by Grosse (2011) tend to do worse than those in any of their ‘non-dynamic’ control treatments. In particular, this holds for groups in which the endowment history was induced corresponding to the history of a randomly chosen ‘twin group’ in the main treatment. This latter finding seems to be in line with the earlier findings of Sadrieh and Verbon (2006) reported above.

Cadigan *et al.* (2011) investigate carry-over effects in a two-stage public-goods game. In one of their experimental settings, the payoffs in stage one determine the endowments in stage two and participants are paid only after stage two. Similar to the results of Grosse (2011), they find cooperation-inhibiting effects of introducing dynamically evolving contribution capabilities in all of their treatments.

Battaglini *et al.* (2016) examine behavior in a durable-public-good setting, focusing on whether agents play trigger strategies, and on how behavior is influenced by whether decisions are made independently or by some political mechanism such as majority voting. They find that even though there are equilibria sustaining high levels of cooperation in their game, the decentralized institution yields little cooperation as predicted by the Markov-equilibrium steady state. At the same time, a majority rule substantially increases the amount of cooperation, albeit not enough to achieve an efficient outcome.

Recently, Gächter *et al.* (2017) published a study that used a shorter time-horizon in conjunction with linear punishment technology, as well as a linear point-to-money ratio. While their set-up is similar to ours, our research focus is different. While Gächter *et al.* (2017) focus on the role of inequality and anti-social punishment for long-term group performance, we focus on the effects of different *pro*-social punishment behaviors.

1.2 *Punishment and the Evolution of Cooperation*

The idea that informal punishment by peers is likely to have played a crucial role in enabling large-scale human cooperation has been contested on three accounts. First, it has been questioned whether punishment plays a role in those societies

that arguably are closest to our evolutionary past (Guala, 2012). Bowles *et al.* (2012) and Casari (2012) point to several papers documenting informal punishment by peers such as, for example, Mathew and Boyd (2011), who provide clear evidence of informal punishment for the Turkana, an egalitarian nomadic pastoral society in northwest Kenya.

Second, researchers like Dreber *et al.* (2009) have pointed out that in most studies, the costs of punishment lead to average earnings that are even lower than in the corresponding treatments where punishment is not available. Yet, Gächter *et al.* (2008) show that when the time horizon is long enough, this is no longer true.

Third, studies like Rand *et al.* (2010) or Nikiforakis (2008) have argued that the optimistic findings of studies like Gächter *et al.* (2008) are due to the fact that the latter have defined away the potential problem of retaliation. On the other hand, Nicklisch and Wolff (2011) who explicitly allow for retaliation observe average payoffs that stabilize clearly above the Nash prediction (at about a third of the way between Nash prediction and social optimum), which is higher than what is achieved towards the end in the typical public-good experiment without punishment.

2 Methods

In this section, we first describe the games underlying each treatment in due detail. Then, we describe the experimental procedures. In the DYNNOPUN game, each round t , $t = 1, \dots, 20$, has exactly one stage in which a standard public-good game with $n = 4$ players is played. In the first round, each player is endowed with an identical amount of $E_i^1 = 20$ tokens. The contribution capability (or current wealth) E_i^t , of each subsequent round corresponds to that player's wealth at the end of the preceding round, Ω_i^{t-1} , plus a small 'recovery surplus' of $m = 2$, so that $E_i^t = \Omega_i^{t-1} + 2$ for $t > 1$. In every round, each player i may contribute x_i^t tokens from her current wealth to a common project and keeps the remainder on a private account. The total contributions are multiplied by $n\mu = 1.6$ and divided evenly amongst the players in the group, so that each unit contributed yields $\mu = 0.4$ units for every player in the group. Thus, player i 's wealth Ω_i^t at the end of round t is:

$$\Omega_i^t = E_i^t - x_i^t + 0.4 \sum_j x_j^t, \quad t = 1, \dots, T$$

In the DYNPUN game, a second stage is added. After the first stage, which is identical to that of the DYNNOPUN game, players are informed about all players' contribution decisions and may then assign punishment points to the other players in their group. By assigning p_{ij}^t points to player j , player i can

reduce the round- t wealth of player j by p_{ij}^t . Punishment is not only costly for the punished, but also for the punisher. The assignment of p_{ij}^t points inflicts costs of $c(p_{ij}^t)$ on player i , where $c(p_{ij}^t) = p_{ij}^t/3 + p_{ij}^3/2000$. This cost function is a convex function that is positive for all positive values of p_{ij}^t , monotonically increasing, and approximates the typical 1 : 3 punishment technology for small punishments. We set two further constraints on punishment: players cannot assign values of p_{ij}^t that would drive their own current account below zero, and they cannot drive other players' current account at the end of the round below zero. If they assign more points than necessary to eliminate another player's positive earnings, they nevertheless have to bear the full costs of their choice. Player i 's current wealth Ω_i^t at the end of round t hence is given by:

$$\Omega_i^t = E_i^t - x_i^t + 0.4 \sum_j x_j^t - h \left(\sum_j p_{ji}^t \right) - \sum_j c(p_{ij}^t) \quad t = 1, \dots, T$$

with

$$\sum_j c(p_{ij}^t) \leq E_i^t - x_i^t + 0.4 \sum_j x_j^t$$

and

$$h \left(\sum_j p_{ji}^t \right) = \min \left\{ \sum_j p_{ji}^t, E_i^t - x_i^t + 0.4 \sum_j x_j^t \right\}$$

To stress the dynamic aspect of the environment, participants' earnings at the end of the experiment were determined by their final wealth level Ω_i^{20} .

Our experiment was programmed in z-Tree (Fischbacher, 2007) and run at the Erfurt Laboratory for Experimental Economics (eLab). We ran 4 sessions, 2 for each of our treatments, collecting the data of 9 groups per treatment. A total of 72 participants (47 female, 25 male) were recruited using ORSEE (Greiner, 2015). The instructions (see Appendix A) were handed out in written form before being read aloud by the experimenter. After this, participants were given the opportunity to go over the instructions again and ask any questions they might have. The experiment started only after the experimenter answered all questions individually.

At the beginning of the experiment, each participant was assigned an identification letter (R, S, T, or U) that was kept constant over the course of the experiment. Assignment to groups was random and groups did not change during the entire session. We obtained 9 independent observations, i.e., groups, for each treatment. Participants were paid by their individual performance according to the following formula:

$$\text{Payment in Euros} = (\text{Number of experimental tokens accumulated})^{2/7}$$

This translated into possible payments between 0 and 40 euros. This formula satisfies the “precept” of *saliency* as formulated by Smith (1982, p. 930f) and extensively discussed in Bardsley *et al.* (2010). However, it does so in a non-standard way. The reason for this choice is the following: a linear token-to-Euro ratio leads to one of two problems. A high linear exchange rate would make the experiment prohibitively costly, bearing in mind that the maximum number of tokens that could be achieved in the experiment exceeded 400’000. In contrast, too low a rate may violate the *dominance* “precept”: incentives would be too low to make monetary payoffs an important concern.⁶ To fulfill *dominance*, we made a clear reference to the maximum potential earnings of 40 Euros in the instructions.

The sessions lasted approximately three quarters of an hour, average payments being 8.30 Euros. Payments were settled individually to ensure players’ anonymity. Also, no other information was given to the participants that would enable them to connect the players in the game with the respective participants in the session.

3 Hypotheses

Under rational selfish payoff-maximization and common knowledge of preferences and rationality, no player will make positive contributions in either treatment, nor punish other players in case of the DYNPUN game, following directly from the typical backward-induction argument. Yet, if either common-knowledge assumption fails, Kreps *et al.*’s (1982) argument – that it may pay for a selfish payoff-maximizer to invest in the possibility of the other player being either non-selfish or non-rational – gains additional strength in our setup. This is because (early) contributions increase future contribution capabilities. If others can be expected to contribute in later stages – be it because they follow a tit-for-tat strategy or simply because they are unconditionally altruistic – increasing their contribution capabilities at the same time will increase the source on which to free-ride later on. The important difference between

⁶Imagine, for a moment, the case of a 1000:1 ratio, corresponding to a maximum potential gain of 400 Euros (compared to an average hourly wage of a student research assistant of 7.50 Euros). Recall that participants are endowed with 20 tokens, corresponding to 2 Cents in our example. This could lead to a false perception by the participants that their decisions throughout the experiment must result in a very small final payment. Additionally, for a participant to obtain the hourly wage of a student assistant under this exchange rate, it would need 13 (out of 20) rounds of full cooperation if all players would cooperate right from the start (for comparison: under our payoff scheme, 20 tokens translate to 2.35 Euros, and a student assistant’s hourly wage was equivalent to 8 rounds of full cooperation). In other words, for the monetary payoff from the experiment to be considered meaningful, we would have needed to increase a linear token-to-ratio even further. For groups under a linear exchange rate to reach the hourly-wage equivalent as under the scheme actually used, the ratio would have needed to be as high as 125:1, with a maximum gain of 3’200 Euros per participant.

the typical repeated one-shot game and our dynamic game in this respect is that contributing in early rounds is a best-response even to *unconditional* contributors.

In addition to the strengthened incentive for early contributions when players cannot be sure of the other players' strategies, there is another reason to believe that contributions will be higher than "usual" in the DYNNOPUN treatment. As we show exemplarily for players with inequality-averse preferences à la Fehr and Schmidt (1999), there are additional positive-contributions equilibria on top of the very unlikely case of all players being strongly inequality-averse. In these equilibria, money-maximizing players always contribute their full endowment up to the final periods, while inequality-averse players play suspicious-tit-for-tat: they start out contributing nothing and then always copy the action of the next-richest player from the preceding round (in these equilibria, the money-maximizing players). In the final periods, money-maximizing players free-ride completely, and inequality-averse players contribute until payoffs are equalized with the money-maximizers.⁷ In other words, we expect relatively high rates of cooperation in our DYNNOPUN treatment.

Do we expect to observe the common cooperation-enhancing effect of punishment in our DYNPUN treatment? Rather not, for the reasons outlined already. In repeated public-good experiments with punishment, punishment levels tend to be high in early periods.⁸ In a situation in which contribution capabilities do not depend on earlier play, high early-round punishment levels have been shown to lead to stable or growing contribution levels and earnings surpassing those from the comparable game without sanctions.⁹ In a situation with dynamically evolving contribution capabilities, however, a "conditionally cooperative enforcer" has to strike a balance in the following trade-off: punishing a low-contributing player may induce higher future cooperation levels, but at the same time, it destroys parts of the future contribution capabilities of both the punisher and the punished player. This difficulty, together with the relatively high expected cooperation gains in the control treatment without punishment suggest that earnings in DYNPUN will not surpass those from DYNNOPUN.

⁷See Appendix B; for the parameters used in our experiment, the likelihood of the preconditions for this additional type of equilibrium to be met amounts to roughly 35%, according to the type distribution suggested by Fehr and Schmidt (1999). In fact, this class of equilibria is more general than proposed here. Money-maximizers' equilibrium strategy could prescribe to contribute any arbitrary fraction of their wealth, as long as it is symmetric, and to stop contributing in period $T - t'$. The conditionally cooperative players would mirror money-maximizers' contributions in the respective subsequent period and refrain from contributing positive amounts in all periods $t > T - t' + 1$. However, the most efficient of these equilibria is the one with full money-maximizer contributions and $t' = 1$. Hence, this equilibria would be chosen by the same equilibrium refinement argument Fehr and Schmidt (1999) employ to choose the full-contribution equilibrium for the public-good game with punishment.

⁸For an overview, cf., e.g., Gächter and Herrmann (2009).

⁹Cf., e.g., Nikiforakis and Normann (2008).

4 Results

4.1 Main Results

The upper two panels of Figure 1 show the trajectories of average group wealth for all 18 groups, by experimental treatment (DYNPUN in the left panel and DYNNOPUN in the right panel). Treatment averages are drawn in bold, red, and with x-symbols.

Result 1. *Final wealth levels in the dynamic public-good experiment with punishment opportunities are neither substantially nor significantly higher than in the experiment without punishment opportunities.*

As Figure 1 shows, there is no clear difference in final-wealth averages between the two treatments: only the best-performing DYNPUN group seems to clearly outperform the best-performing DYNNOPUN group; apart from those two groups, the distributions are remarkably similar.¹⁰ Looking at average contributions per group in the lower part of Figure 1, there seems to be slightly more of a difference. However, most of this difference is (again) due to the most cooperative group in either treatment and to a stronger end-game effect in the DYNNOPUN treatment.¹¹ On average, punishment expenditures are offset by slightly higher contributions relative to current wealth levels.

Figure 1 also shows that there is considerable heterogeneity between groups. To understand this heterogeneity, we focus on early-round behavior, as a proxy for the norms, values, or social heuristics that our participants bring to the laboratory. To do so, we look at first-round average contributions to assess a group’s “intrinsic cooperativeness”. In the treatment with punishment opportunities, we also separate the groups into (i) “groups with heavy first-round punishment” vs “groups without heavy first-round punishment”, depending on whether all high-contributing participants punished in the first round,¹² and into (ii) “retaliative groups” and “non-retaliative groups”, depending on whether any punishment acts classifiable as retaliation happen in rounds 2 or 3. To use the broadest definition of what may be classified as “retaliation”, a punishment act was classified as “retaliation” if player j had punished player i in period $t - 1$ and player i punished j in period t .

Result 2. *In the dynamic public-good experiment without punishment opportunities, initial contribution levels predict final wealth levels, i.e. initially more*

¹⁰A Wilcoxon Mann-Whitney test for a difference in final-wealth averages yields $p = 0.730$ (the test includes the best-performing groups).

¹¹A Wilcoxon Mann-Whitney test yields $p = 0.133$ for final average contributions but only $p = 0.340$ for contributions in periods 14–19.

¹²There are no groups in which only a subset of the high-contributing participants punished. As high-contributing participants, we categorize those who contributed at least the average contribution of the other players in the group.

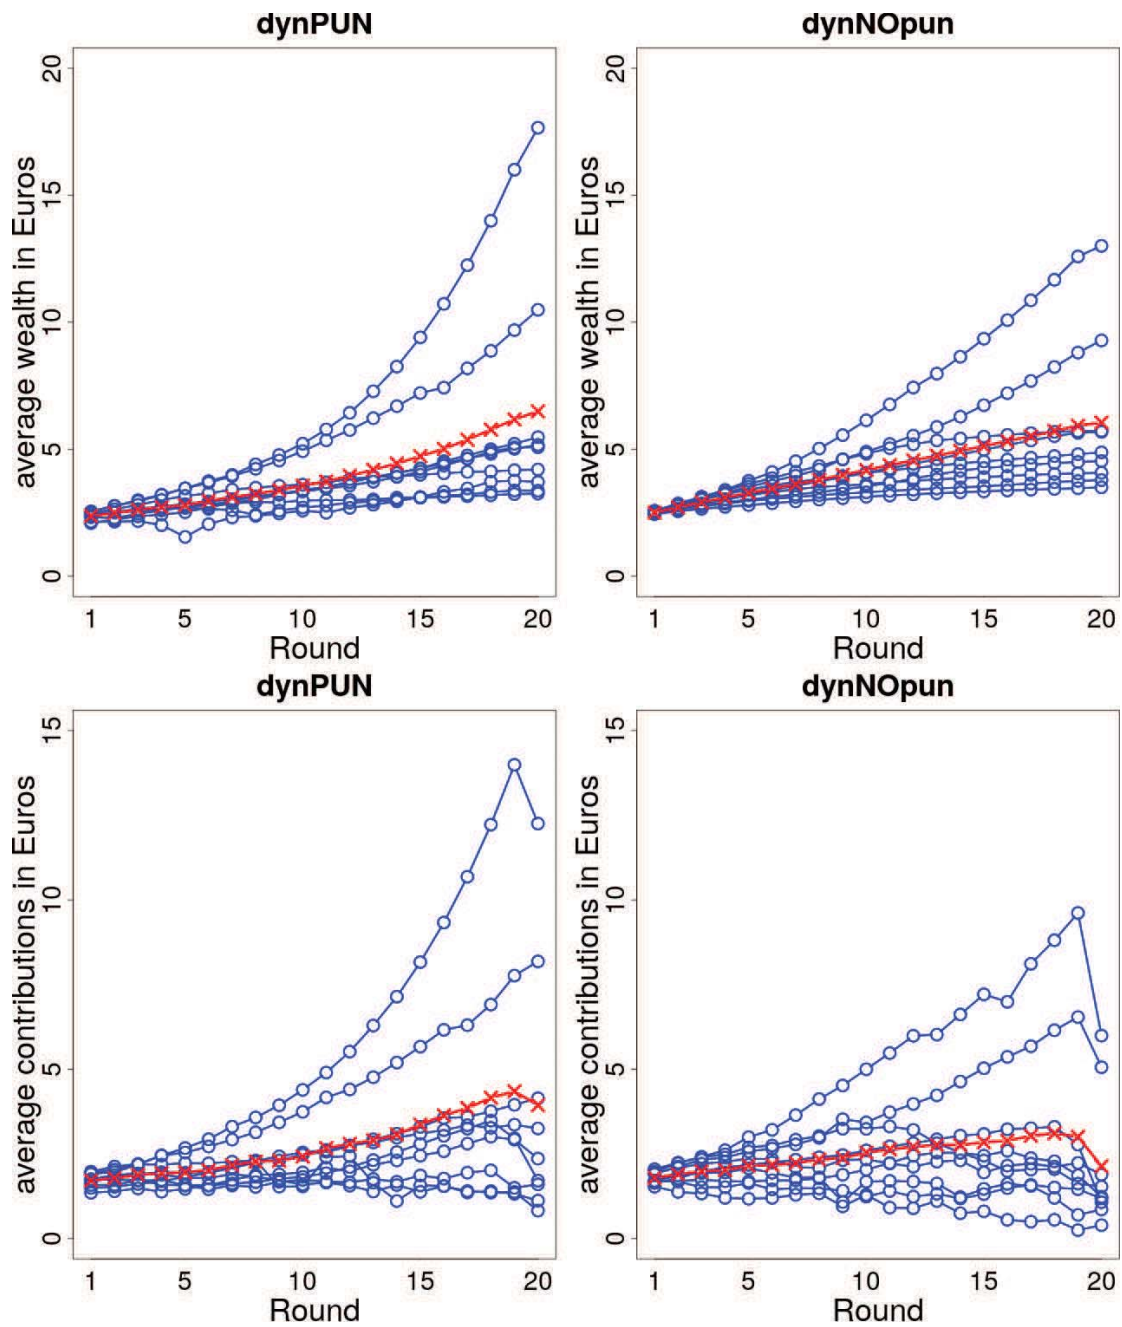


Figure 1: Wealth levels (above) and contributions (below) in Euros (group averages), from DYNPUN (left) and DYNNOPUN (right). Treatment averages are drawn in bold, red, and with x-symbols.

cooperative groups achieve higher final wealth. In contrast, in the experiment with punishment opportunities, final wealth levels are determined by early punishment behavior and not by early contribution levels: first-round punishment is strongly associated with early retaliation, and groups with early retaliation achieve lower final-wealth levels. At the same time, initial contributions or their variance do not explain early punishment behavior.

In DYNNOPUN, the disposition towards cooperation as measured by first-round contribution levels clearly predicts final-round wealth levels (Spearman’s $\rho = 0.68$, $p = 0.045$). This relationship is not significant in DYNPUN ($\rho = 0.29$, $p = 0.444$). Instead, it is the separation into “retaliative” and “non-retaliative groups” that predicts final wealth ($p = 0.063$).¹³ At the same time, there is a strong correlation between “groups with heavy first-round punishment” and “retaliative groups” (Spearman’s $\rho = 0.68$, $p = 0.045$). However, in our data the difference between “groups with heavy first-round punishment” and “groups without heavy first-round punishment” in terms of final wealth levels does not reach significance ($p = 0.191$). This suggests that being punished already in the very first round very often sparks retaliation; whenever it does, this leads to detrimentally high frequencies of punishment (and counter-punishment; $\rho = 0.87$, $p = 0.003$). We plot both relationships in Figure 2. The Figure shows a clear separation of final group earnings by whether a group is “retaliative” or not, and a somewhat weaker separation when conditioning on whether all the group’s high-contributors punish in period 1. The fact that both “first-round punishment” and “early retaliation” are not related to the contribution level and variance (all $p \geq 0.242$; also in linear regressions including both, all coefficients have $p \geq 0.289$) seems to suggest that it is not the degree of participants’ cooperative disposition that matters, but really their ideas of how and when punishment should be administered.

4.2 Individual Behavior

So far, we have looked at the aggregate determinants of final wealth levels in the two treatments. We have seen that the dynamics differ markedly between groups in the DYNNOPUN treatment, groups with heavy first-round punishment, and groups without first-round punishment in the DYNPUN treatment.¹⁴ To explore how these group differences are reflected in individual behavior, we examine the factors that influence individual contribution decisions by the regression analysis reported in Table 1. For the analysis, we used mixed-effects models with (nested) random effects for groups and individuals. We regress a player’s period-to-period change in relative contributions on a number of lagged variables that may be expected to influence the player’s decision. We use contributions relative to current wealth levels to make decisions comparable between rounds (and consequently, between endowments; we normalize most

¹³If we exclude from the analysis the group abstaining from punishment altogether, the p -value changes only slightly, $p = 0.071$. None of these figures change if we broaden our definition to the first three periods in which retaliation could happen. Only when we restrict it to the very first period, we do not have enough observations for “retaliative groups”.

¹⁴We use “heavy first-round punishment” as the focal variable rather than “early retaliation” because first-round punishment precedes retaliation. If we use “early retaliation” as the explanatory variable in the following regression analysis, we obtain virtually identical estimates.

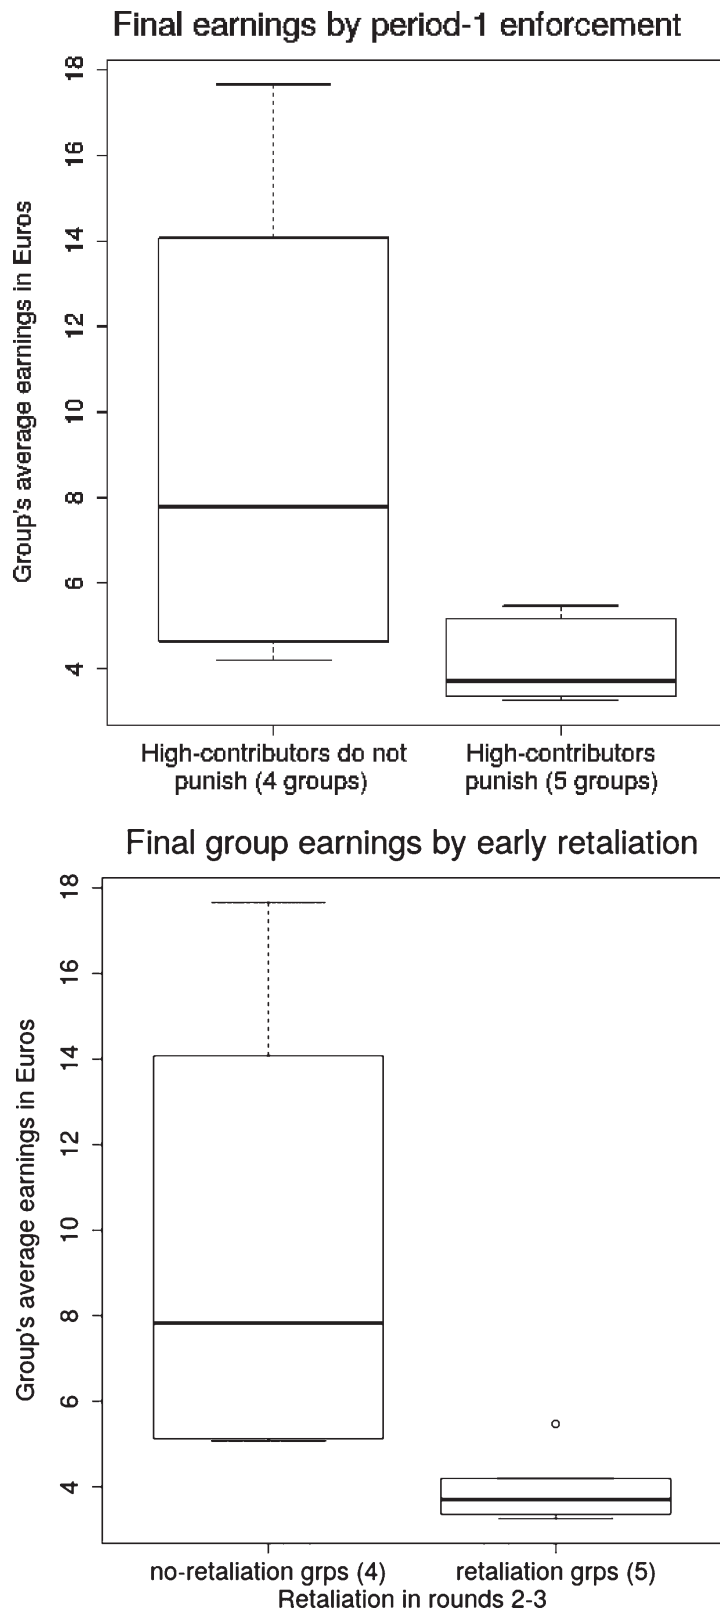


Figure 2: Final wealth levels in Euros (group averages) when there are no vs. when there are punishment actions in period 1 by high-contributors (above), and by whether there are no vs. when there are retaliative actions in rounds 2 or 3.

of the explanatory variables for the same reason). As explanatory variables we use first-round contributions relative to the first-round endowment (which is identical for all players; EV1) and three variables that measure deviations from the group average, each split up into a positive-deviation (the individual's contribution is above the group average) and a negative-deviation (the individual's contribution is below the group average) variable, to account for the fact that reactions may differ depending on the player's position within the group. The variables EV2 and EV5 capture conditionally cooperative behavior: they refer to the deviation of the player's contribution from the group members' average contribution, in absolute or relative terms. We still normalize absolute deviations (by the group's sum of contributions) to make the decisions from different rounds comparable. The variables EV6 and EV7 measure the player's wealth standing, to account for the player's 'historical' relative wealth level within society. Together with the variation coefficient of the group's current contribution capabilities, EV10, this variable gives an indication of how the endogenously arising inequality influences behavior. While all three deviation measures are related, their combination allows for a more subtle picture of players' reaction to their peers' past contribution behavior. Additionally, we include the logarithm of the average capability, EV12, to account for the current group level of prosperity, and the period, EV11, to allow for potential time trends. In the first model pertaining to the data from our DYNPUN treatment, we add a dummy variable indicating whether the player had been sanctioned in the preceding round, EV8, as well as the fraction of the player's current wealth destroyed by others' assignments, EV9. Given our interest in the difference in dynamics between groups with heavy first-round punishment and those without, we also add a dummy variable indicating whether the player is in a group with heavy first-round punishment, EV13, and interact all other explanatory variables with it.

Mirroring the observation that initial contributions predict final outcomes very well in DYNNOPUN groups, we see that higher initial contributions are associated with higher increases of the relative contribution round after round. In addition, behavior in DYNNOPUN groups shows a high degree of conformism: when participants have made higher relative contributions than others in $t - 1$, they significantly reduce their relative contributions in t . Also, higher relative contributions of others were followed by higher relative contributions of the respective participant. Finally, having been a (relative) free-rider in the past (as indicated by a positive deviation of the participant's $t - 1$ capability from the average capability in $t - 1$) also has a slight contribution-increasing effect in this treatment, and conversely, having been taken advantage of tends to have a contribution-decreasing effect. At the same time, there is a slight – counteracting – tendency to reduce the following relative contribution when a participant's absolute contribution was lower than the contribution of others. The fact that the group's average contribution capability enters

Table 1: Mixed-effects regression of period-to-period changes in relative contributions with nested random effects on the individual and group levels.

	DYNPUN	DYNNOPUN
<i>(Intercept)</i>	0.068 (0.031)	-0.011 (0.034)
EV1: Relative contribution in $t = 1$	0.065 (0.041)	0.373 (0.061)
EV2: Positive deviation from the average relative contribution in $t - 1$	-0.374 (0.085)	-0.648 (0.061)
EV3: Negative deviation from the average relative contribution in $t - 1$	-0.036 (0.106)	0.253 (0.065)
EV4: Positive deviation from the average absolute contribution in $t - 1$	-0.013 (0.014)	0.008 (0.007)
EV5: Negative deviation from the average absolute contribution in $t - 1$	-0.078 (0.040)	-0.035 (0.019)
EV6: Positive deviation from the average capability in $t - 1$	-0.036 (0.071)	0.102 (0.045)
EV7: Negative deviation from the average capability in $t - 1$	0.215 (0.088)	-0.112 (0.064)
EV8: Dummy: having been punished in $t - 1$	0.047 (0.022)	
EV9: Received punishment as a fraction of the current wealth level in $t - 1$	0.606 (0.289)	
EV10: Variation coefficient of the group's current contribution capabilities	-0.079 (0.091)	-0.113 (0.081)
EV11: Period	0.002 (0.003)	0.003 (0.002)
EV12: Logarithm of the group's average contribution capability	-0.019 (0.010)	-0.031 (0.009)
EV13: Heavy Punishment in $t = 1$ (HPin1)	0.016 (0.071)	
EV14: HPin1*Relative contribution in $t = 1$	-0.070 (0.082)	
EV15: HPin1*Positive deviation from the average relative contribution in $t - 1$	0.339 (0.123)	
EV16: HPin1*Negative deviation from the average relative contribution in $t - 1$	0.132 (0.163)	
EV17: HPin1*Positive deviation from the average absolute contribution in $t - 1$	0.014 (0.024)	
EV18: HPin1*Negative deviation from the average absolute contribution in $t - 1$	-0.034 (0.054)	
EV19: HPin1*Positive deviation from the average capability in $t - 1$	0.015 (0.093)	
EV20: HPin1*Negative deviation from the average capability in $t - 1$	-0.180 (0.117)	
EV21: HPin1*Dummy: having been punished in $t - 1$	-0.056 (0.027)	
EV22: HPin1*Received punishment as a fraction of the current wealth level in $t - 1$	-0.396 (0.297)	
EV23: HPin1*Variation coefficient of the group's current contribution capabilities	0.081 (0.116)	
EV24: HPin1*Period	0.001 (0.003)	
EV25: HPin1*Logarithm of the group's average contribution capability	-0.009 (0.020)	
R ² /Number of obs./groups/individuals	0.22/684/9/36	0.57/684/9/36

Notes: Deviations are normalized by the average contribution capability and average sum of absolute contributions, respectively. Standard errors are reported in parentheses. The coefficients in italics are not significant if "Heavy Punishment in $t = 1$ " is the baseline category. Significance Levels: ***, **, *, +, -.

negatively is likely to be explained by the fact that the observations with high contribution capabilities come from the second half of the experiment in the better-performing groups. In these groups and at that stage of the experiment, relative contributions will have converged, so that they no longer exhibit increases. Hence, higher contribution capabilities will be associated with lower increases.

In DYNPUN groups without heavy first-round punishment, behavior seems to follow similar principles as in DYNNOPUN groups, with two exceptions: initial contributions do not have a clear influence, but “earlier” contributions do, as the significant coefficient for negative deviations from the average capability in $t - 1$ shows (participants who have been exploited more in the past continue to increase their contributions more than others). Positive deviations from the average relative contribution are followed by lower relative contributions in the next round – albeit less strongly so than in DYNNOPUN – and again, there is a slight tendency for negative deviations from the average absolute contribution to be followed by a further reduction. As in the DYNNOPUN case, this tendency is likely to be overpowered by other effects, in this case, by the contribution-enhancing effects of being punished. Interestingly, in those DYNPUN groups displaying heavy first-round punishment, hardly any of the effects survive: in these groups, only a negative deviation from the average absolute contribution continues to bring about lower increases in the next round while received punishment as a fraction of the current wealth level still increases relative contributions. However, the effect is reduced to one third of the effect in those groups without first-round punishment, and the dummy for being punished is no longer predictive of behavior at all.

The above analysis nicely mirrors the main findings from the analysis of the aggregate data: the pre-disposition to cooperate has strong predictive power in the DYNNOPUN treatment, but in DYNPUN groups a cooperative pre-disposition seems to be irrelevant. Rather, the occurrence of heavy first-round punishment foreshadows the dynamics of the game: in those groups with heavy first-round punishment, positive examples by high-contributors do not help (having contributed less does not lead to an increase in contributions as in DYNPUN, EV3), and punishment has a rather small effect. In groups without heavy first-round punishment, punishment seems to be playing a beneficial role, as described in the literature: having been punished is followed by an increased contribution, and the heavier this punishment is, the larger the increase.

5 Discussion

We have studied a dynamic public-good game in which players’ contribution capabilities are endogenously determined by players’ earlier choices. We find

that dynamics by themselves do not solve the dilemma, replicating the earlier findings of Grosse (2011) and Cadigan *et al.* (2011). In this kind of environment, we analyzed the effects of peer-punishment opportunities to understand better how these opportunities will have affected groups in our evolutionary past. When choices have dynamic effects, punishment opportunities may give rise to two critical issues that have not been taken into account in the literature: in one extreme scenario, the level of punishment could be so high that endowments shrink over time and contributions are lower than when there are no punishment opportunities, in spite of significantly higher relative contribution levels. At the other extreme, we could observe no punishment whatsoever, as potential punishers would be too concerned about maintaining future contribution capabilities. Both scenarios are relevant, but neither dominates the results: While we do not observe any group in which wealth levels actually decrease, there was one group in DYNPUN in which all individual relative contributions are well above the average relative contribution from the DYNNOPUN treatment for most of the time – and yet, this group’s wealth levels stay as low as in the second-worst performing DYNNOPUN group. On the other extreme, we have a group in which punishment was virtually never used before the kicking in of the end-game effect in round 19.¹⁵ This group’s performance corresponds to the median group from the treatment without punishment opportunities.

Guala (2012), Bowles *et al.* (2012), and Casari (2012), argue that uncoordinated peer-punishment is not the kind of informal sanctioning that is typically observed in acephalous societies. Our results may provide a rationale for why this is the case. We suggest that in dynamic settings a successful group needs something in addition to the punishment opportunities, be it a reputation system based on punishment (Dos Santos and Wedekind, 2015) or a common understanding – a social norm – of when punishment is acceptable, who is allowed to punish, and of how retaliative behavior can be prevented (cf., e.g., Bowles *et al.*, 2012). The implications of our results suggest that peer punishment can make the societies less vulnerable to free-riders if such a reputation system or social norm is present: if free-riders are part of the group, initial cooperation levels tend to be low. When there are no punishment opportunities, Result 2 shows that final earnings then will also be low. However, when there are punishment opportunities and a reputation system or social norm prevents excessive punishment, the group will thrive in spite of its uncooperative starting point.

¹⁵As a matter of fact, there was a single assignment of 1 punishment point in period 16. In the final two periods, there were 2 (4) assignments, destroying 34 (50) out of 1216 (1259) points in period 19 (20; punisher costs included).

References

- Admati, A. R. and M. Perry. 1991. "Joint Projects without Commitment". *Review of Economic Studies*. 58(194): 259–276.
- Bardsley, N., R. Cubitt, G. Loomes, G. Moffatt, C. Starmer, and R. Sugden. 2010. *Experimental Economics*. Princeton and Oxford: Princeton University Press.
- Battaglini, M., S. Nunnari, and T. R. Palfrey. 2016. "The Dynamic Free Rider Problem: A Laboratory Study". *American Economic Journal: Microeconomics*. 8(4): 268–308.
- Bowles, S., R. Boyd, S. Mathew, and P. J. Richerson. 2012. "The Punishment that Sustains Cooperation is Often Coordinated and Costly". *Brain and Behavioral Sciences*. 35(1): 20–21.
- Cadigan, J., P. T. Wayland, P. Schmitt, and K. Swope. 2011. "An Experimental Dynamic Public Goods Game with Carryover". *Journal of Economic Behavior and Organization*. 80(3): 523–531.
- Casari, M. 2012. "Weak reciprocity alone cannot explain peer punishment". *Behavioral and Brain Sciences*. 35(1): 21–22.
- Chaudhuri, A. 2011. "Sustaining Cooperation in Laboratory Public Goods Experiments: A Selective Survey of the Literature". *Experimental Economics*. 14: 47–83.
- Denant-Boemont, L., D. Masclet, and C. Noussair. 2007. "Punishment, Counterpunishment and Sanction Enforcement in a Social Dilemma Experiment". *Economic Theory*. 33: 145–167.
- Dos Santos, M. and C. Wedekind. 2015. "Reputation Based on Punishment Rather than Generosity Allows for Evolution of Cooperation in Sizable Groups". *Evolution and Human Behavior*. 36: 59–64.
- Dreber, A., D. G. Rand, D. Fudenberg, and M. A. Nowak. 2009. "Winners Don't Punish". *Nature*. 452: 348–351.
- Fehr, E., U. Fischbacher, and S. Gächter. 2002. "Strong Reciprocity, Human Cooperation, and the Enforcement of Social Norms". *Human Nature*. 13(1): 1–25.
- Fehr, E. and S. Gächter. 2000. "Cooperation and punishment in public goods experiments". *American Economic Review*. 90(4): 980–994.
- Fehr, E. and K. M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation". *Quarterly Journal of Economics*. 114(3): 817–868.
- Fershtman, C. and S. Nitzan. 1991. "Dynamic Voluntary Provision of Public Goods". *European Economic Review*. 35(5): 1057–1067.
- Fischbacher, U. 2007. "z-tree: Zurich Toolbox for Ready-Made Economic Experiments". *Experimental Economics*. 10(2): 171–178.
- Gächter, S. and B. Herrmann. 2009. "Reciprocity, Culture and Human Cooperation: Previous insights and a new cross-cultural experiment". *Philosophical Transactions of the Royal Society B: Biological Sciences*. 364: 791–806.

- Gächter, S., F. Mengel, E. Tsakas, and A. Vostroknutov. 2017. "Growth and Inequality in Public Good Provision". *Journal of Public Economics*. 150: 1–13.
- Gächter, S., E. Renner, and M. Sefton. 2008. "The Long-Run Benefits of Punishment". *Science*. 322(5907): 1510.
- Gintis, H., S. Bowles, R. Boyd, and E. Fehr. 2005. "Moral Sentiments and Material Interests: Origins, Evidence, and Consequences". In: *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*. Ed. by H. Gintis, S. Bowles, R. Boyd, and E. Fehr. Cambridge and London: The MIT Press. 3–40.
- Greiner, B. 2015. "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE". *Journal of the Economic Science Association*. 1(1): 114–125.
- Grosse, S. 2011. "Experiments on the Fragility of Cooperation and Mechanisms to Overcome this Problem". In: Cuvillier Verlag, Göttingen. Chap. 1. 1–25.
- Guala, F. 2012. "Reciprocity: weak or strong? What Punishment Experiments Do (and Do Not) Demonstrate". *Behavioral and Brain Sciences*. 35(1): 1–15.
- Henrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, and H. Gintis. 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. New York: Oxford University Press.
- Kreps, D. M., P. Milgrom, J. Roberts, and R. Wilson. 1982. "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma". *Journal of Economic Theory*. 27(2): 245–252.
- Ledyard, J. 1995. "Public Goods: A Survey of Experimental Research". In: *Handbook of Experimental Economics*. Ed. by A. E. Roth and J. Kagel. Princeton: Princeton University Press. 111–194.
- Mathew, S. and R. Boyd. 2011. "Punishment Sustains Large-Scale Cooperation in Prestate Warfare". *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*. 108(28): 11375–11380.
- Nicklisch, A. and I. Wolff. 2011. "Cooperation Norms in Multiple-Stage Punishment". *Journal of Public Economic Theory*. 13(5): 791–827.
- Nikiforakis, N. 2008. "Punishment and Counter-Punishment in Public Good Games: Can we really Govern Ourselves?" *Journal of Public Economics*. 92: 91–112.
- Nikiforakis, N. and H.-T. Normann. 2008. "A Comparative Statics Analysis of Punishment in public-good experiments". *Experimental Economics*. 11(4): 358–369.
- Noussair, C. and C. Soo. 2008. "Voluntary contributions to a dynamic public good: Experimental evidence". *Economics Letters*. 98(1): 71–77.

- Rand, D. G., J. J. Armao, M. Nakamaru, and H. Ohtsuki. 2010. "Anti-Social Punishment Can Prevent the Co-evolution of Punishment and Cooperation". *Journal of Theoretical Biology*. 265: 624–632.
- Rapoport, A. 1988. "Experiments with N-Person Social Traps II: Tragedy of the Commons". *Journal of Conflict Resolution*. 32: 473–488.
- Reuben, E. and A. Riedl. 2013. "Enforcement of Contribution Norms in Public Good Games with Heterogeneous Populations". *Games and Economic Behavior*. 77(1): 122–137.
- Richerson, P. J. and R. Boyd. 2005. *Not by Genes Alone*. Chicago and London: The University of Chicago Press.
- Sadrieh, A. and H. A. A. Verbon. 2006. "Inequality, Cooperation, and Growth: An Experimental Study". *European Economic Review*. 50: 1197–1222.
- Sigmund, K. 2007. "Punish or Perish? Retaliation and Collaboration Among Humans". *Trends in Ecology and Evolution*. 22(11): 593–600.
- Smith, V. L. 1982. "Microeconomic Systems as an Experimental Science". *American Economic Review*. 45(4): 671–690.