Designing Institutions for Social Dilemmas

Bettina Rockenbach and Irenaeus Wolff

University of Cologne, Thurgau Institute of Economics/University of Konstanz

Abstract. *Considerable experimental evidence has been collected on rules enhancing contributions in public goods dilemmas. These studies either confront subjects with pre-specified rules or have subjects choose between different rule environments. In this paper, we completely endogenize the institution design process by asking subjects to design and repeatedly improve rule sets for a public goods problem in order to investigate which rules social planners facing a social dilemma ‘invent’ and how these rules develop over time. We make several noteworthy observations, in particular the strong and successful use of framing, the concealment of individual contribution information and the decreasing use of punishment.*

Keywords: Public good; strategy method; experiment; public choice.

1. INTRODUCTION

Understanding the determinants of human cooperation is one of the most challenging questions in economics. Human cooperation in social dilemmas is particularly puzzling because the conflict between individual and collective interests creates the well-known free-rider problem.¹ Repeated interactions in controlled social dilemma experiments show that initially positive cooperation levels decline over time and thus nurture Hardin’s pessimistic vision.²

Yet, in order to overcome the breakdown of cooperation and to disentangle different motives for cooperation and defection, researchers have incorporated a huge variety of institutional regulations into the workhorse for studying social dilemmas, the experimental public goods provision game. Among the most prominent regulations examined are punishment opportunities,³ communication,⁴ leadership,⁵ reputation opportunities,⁶ and ostracism.⁷ In this ‘first generation’

1. Hardin (1968) and Ostrom (1998).
2. Davis and Holt (1993), Fischbacher and Gächter (2010), Ledyard (1995) and Ostrom (2000).
3. For example Falk *et al.* (2005), Fehr and Gächter (2000), Masclet *et al.* (2003), Nikiforakis (2008), Nikiforakis and Normann (2008), or Yamagishi (1986).
4. For example Bochet *et al.* (2006), Brosig *et al.* (2003), Cason and Khan (1999), Isaac and Walker (1988), or Ostrom *et al.* (1994).
5. For example Arbak and Villeval (2013), Güth *et al.* (2007), Potters *et al.* (2005), or Vesterlund (2003).
6. For example Milinski *et al.* (2006), or Sommerfeld *et al.* (2007).
7. For example Cinyabuguma *et al.* (2005), Güth *et al.* (2007), or Maier-Rigaud *et al.* (2005).

of experiments, subjects face experimenter-determined institutional rules. This yields valuable information on the performance of these rules for fostering cooperation, but cannot answer the question of whether subjects would actually select a certain rule when having the choice. Consequently, subsequent studies have addressed the question of institutional choice among (mostly two) prespecified rules, both by ballot voting⁸ and by ‘voting-with-one’s-feet’.⁹ These ‘second generation’ studies provide important insights into the institutional rules subjects choose initially as well as in the longer run when different (experimenter-given) rules are at choice. However, are these the rules social planners facing a social dilemma would develop themselves? Which are the rules planners would ‘invent’ and how would these rules develop over time in the light of their ability to overcome the social dilemma? An experimental answer to these questions requires to not only endogenize the choice process, but the entire institution formation process. This is exactly what this paper aims at.

In this paper, we introduce an experiment of a ‘third generation’ in which the institution design process is entirely endogenous. The subjects in our experiment act as rule makers empowered to shape the institutional environment of a public goods provision game. They are free in ‘inventing’ any institutional setting and not bound to any prespecified rules, as long as they do not change the social-dilemma incentive structure. The subjects’ incentive for acting as ‘good’ rule makers is that their payment strongly depends on the social benefits that game playing students accrue under the designed rule. Based on performance feedback, subjects may improve their institutional design. To equip the institution designers with a profound understanding and experience in the public goods provision game and an intense experience in the institution design process, we followed Reinhard Selten’s advice and conducted the experiment as a *strategy seminar* (Selten, 1967).

In his seminal paper, Selten (1967) introduced an experimental method, called the *strategy method* and applied it to oligopoly situations.¹⁰ The idea of the strategy method is that subjects, which are highly familiar with the game situation through repeated play of the game, develop strategies for acting in the game. In a series of tournaments, subjects learn about the performance of their strategy and have the possibility to improve the strategy based on this feedback. Usually the strategy method is applied within the framework of a ‘strategy seminar’, a students’ seminar over several months (a semester) to provide a sufficient time frame.¹¹ The obvious advantage of the strategy method is that highly experienced subjects develop strategies that they continuously improve based on performance feedback. The apparent downside is that conducting the experiment is very time consuming, so that the research focus is not on collecting a sufficient

8. For example, Guillen *et al.* (2007), Potters *et al.* (2005), Sutter *et al.* (2005).

9. For example Gürer *et al.* (2006, 2014), Kosfeld *et al.* (2009), Rockenbach and Milinski (2006).

10. The strategy method was applied by Selten *et al.* (1997) in the framework of a Cournot duopoly, by Keser and Gardner (1999) for a common pool resource game, and by Selten *et al.* (2003) for a 3×3 game. The famous study by Axelrod (1984) and the study by Keser (2000) are also in this spirit.

11. Meanwhile the term *strategy method* is commonly used for an experimental protocol in which game playing subjects provide a complete list of actions prior to knowing the actual decision(s) of the other subject(s). It is applied to acquire knowledge about the subject’s reaction to *all* possible actions of the other subject(s) and not only to the specifically chosen one(s).

number of independent observations for statistical analysis, but rather on improving our understanding on how highly experienced experts solve (complex) decision problems. Since the high level of complexity also limits us in gathering observations sufficient for sensible statistical analysis, our study shares the explorative nature of the previous applications of the strategy method.

Our experiment provides us with remarkable insights, which we extensively discuss in the final section. Most noteworthy seems to be: (1) institution designers make extensive but decreasing use of incorporating punishment possibilities, predominantly not in the form of peer-punishment, but centralized punishment or redistribution. (2) Rule designers strongly incorporate communicative elements, not in the form of peer-to-peer communication, but rather in appeals to moral sentiments or a framing of the game situation. (3) Rule designers prefer to conceal individual contributions and rather communicate aggregated contributions levels.

2. MODEL AND EXPERIMENTAL DESIGN

We recruited 24 students of economics at the University of Erfurt for two separate seminar-type courses (12 in each seminar). The seminars ran over 70 days from April to July 2007. Upon arrival, students were allocated to rule-development groups of four who stayed together for the entire seminar. The course of each seminar is summarized in Table 1.

During a preliminary meeting, potential seminar participants were introduced to the schedule. However, they were not informed on what game they would be going to face during the seminar.¹² The seminar was not accompanied by any lectures on theoretic or experimental investigations into social dilemmas, neither in this term nor in previous terms. On the first day of the seminar (Day 1 in Table 1), the participants experienced the basic game by playing it in a laboratory setting over 25 rounds in a partner design. The basic game was a standard public goods game with four players. Each player i contributed an integer amount x_i ($0 \leq x_i \leq 20$) from an endowment of 20 tokens to a common project and kept the remainder. The sum of all contributions was multiplied by 1.6 and divided evenly amongst the players, so that the public good exhibited a constant marginal per-capita return of 0.4. The resulting payoff function of subject i was therefore:

$$\Pi_i = 20 - x_i + 0.4 \sum x_j$$

Subjects were informed only about the sum of contributions, not about individual contributions and there was neither punishment nor any other additional non-standard rule feature.

2.1. Rule design

At the end of the first meeting, participants were randomly allocated to *design groups* of four students each who stayed together until the end of the seminar.

12. See Appendix S1 for the information passed to those present at the preliminary meeting, as well as to the instructions for the basic game.

Table 1 Schedule of the seminars

Time	Playing stages	Rule development
	<i>Preliminary meeting</i>	
Day 1	Play of the basic public goods game	<i>Development groups formed</i>
Days 2–6		Development of 1st set of rules
Day 7		Handing-in of rules
Days 8–13	Software implementation (experimenter)	
Day 14	Play under 1st rule sets	
Days 15–20		Development of 2nd set of rules
Day 21		Handing-in of rules
Day 22–27	Software implementation (experimenter)	
Day 28	Play under 2nd rule sets	
Days 29–34		Development of 3rd set of rules
Day 35		Handing-in of rules
Days 36–41	Software implementation (experimenter)	
Day 42	Play under 3rd rule sets (double weighting)	
Days 43–48		4th set of rules
Day 49		Handing-in of rules
Days 50–55	Software implementation (experimenter)	
Day 56	Play under 4th rule sets (triple weighting)	
Day 70	<i>Final meeting; discussion</i>	

The participants had one week to develop their own set of rules for the public goods game within these design groups. There was no predefined ‘menu’ of rules and the subjects were free to develop whatever rules they wanted (in the boundaries of standard rules of ethics). Each rule (component) is potentially attached to costs. The costs reflect the expenditures the implementation of such a rule would imply in a real-world setting. For some rule sets, costs were split into a fixed and a variable part: for example, setting up the infrastructure to make public announcements as a fixed cost and the variable costs of actually making announcements. The most severe and therefore most expensive interventions are restrictions of the players’ action space, for example to exclude complete free-riding or to even enforce full contribution. These kinds of coercion would not only require a lot of enforcement power, but severely change the nature of the game and therefore, the cost of full-contribution enforcement was set to the maximum gains to be achieved by such an intervention, i.e. 1,200 points. On the other end of the cost range are simple announcements or advertisements to the players.¹³ The cost scheme was hidden from the groups to avoid an ‘anchoring effect’, but any subject was given the possibility to ask for the costs of a specific rule set at all times.

2.2. Implementation, play, and feedback

At the end of the week, each design-group had to hand in a verbal description of their rule set which was subsequently implemented in the experimental software

13. A detailed listing of the introduced rules and the attached costs is provided in Table B1 in Appendix S2.

z-tree¹⁴ and translated into neutrally-worded instructions by the experimenter.¹⁵ After another week, we met again for the first round of play under those rule sets. Randomly allocated *play groups* of 4 played under the different rule sets developed by the rule design groups. To avoid rule designs tailored to a specific subject population, the play groups differed from the design groups. To provide incentives for creating efficient rules and improve them as best as possible, the designers' 'payment'¹⁶ was relative to the social benefits their institutional rule delivered to the play groups. More precisely, we measured efficiency as the sum of individual payoffs within the play group (who also had to bear the costs caused by the rule-set), as a fraction of the payoffs in the social optimum.

By the end of each round of play, seminar participants were provided with detailed feedback on the performance of all rules within their seminar group by round and group, comprising (i) individual contributions, (ii) efficiency, (iii) returns from the public good abstracting from any costs, and (iv) variable costs, as well as the instructions for all rule sets.¹⁷ Overall, there were five rounds of play, with rounds two to five played under rules developed by the design groups.¹⁸ In order to leave some room for initial experimentation, later rules were weighted higher than earlier rules. Specifically, round two and three are weighted by 1, while round four is weighted by 2 and the final fifth round is weighted by 3. In addition to their rule sets' performance, the profits gained in students' individual play also made up for the seminar mark, so that sabotage of alien rules would be costly to the saboteur. Note that individual performance in the different tournaments was weighted evenly.

We had the students clearly separated into two distinct groups that met at different times, asking them not to communicate with students from the other group on the topic of the seminar. This was done to see whether the rule sets would 'converge' to similar sets. Not only did rules not 'converge' to the same set over the seminars; they did not even converge within a seminar group.

2.3. Database

In our experiment, we obtain a database of 24 rule sets from two seminar groups composed of 12 participants each. The corresponding three rule groups in each seminar interacted in four tournaments. Out of these 24 rule sets, we excluded two for our analysis, leaving us with 22 data points.¹⁹

14. Fischbacher (2007).

15. Subjects were given the possibility to have loaded instructions distributed, incurring the same rule costs as on-screen announcements during the experiment.

16. As in the previous applications of the strategy seminar we also 'paid' seminar participants in grades correlating to achieved points.

17. Fixed rule costs were stated in the instructions and implied in the efficiency figures.

18. Recall that there was no rule performance to be measured in the first round of play, given in this round, subjects played the standard public goods game without any additional rules.

19. One group redistributed any contributed points surpassing the minimum contribution back to the contributing player before the sum of contributions was multiplied by the public-good factor. This changes the game from being a social dilemma to a coordination game. We excluded this rule set, which was applied by the group in tournament 2 and 3.

3. RULES USED AND DISREGARDED RULE FEATURES

The rules introduced by our subjects are summarized in Table B1 in Appendix S2. In that table, we briefly describe the rule sets, providing the contribution and efficiency level achieved as well as the costs they gave rise to, grouped by the seminar group and the tournament number. In order to have a better understanding of what the determinants of those rule sets were, we classified the rule components along different categories.

3.1. Punishment and redistribution

The role of punishment in social dilemma situations is widely and prominently discussed in the literature.²⁰ Although our subjects were not trained in this literature, many of the submitted rule sets made use of punishment or redistribution features. Interestingly however, only two of them employed a peer-to-peer mechanism. Most often, rules specified the deduction of points from the lowest-contributor and, in case of redistributive mechanisms, the reallocation of these points in favor of the highest-contributing player. In some of the cases, deducted points were transferred to an account that was to be redistributed at the end to the player having contributed the most. One rule set incorporated a ‘warning system’: those contributing less than the mean were asked to increase their contributions in the following period. In case this advice was not followed, they were punished by an amount conditional on the earlier deviation from the mean. Most of the sets, however, had a direct, automatic punishment or redistribution mechanism based on contribution ranks, with differentiated punishment of the lower-contributors. These observations give rise to four (not mutually exclusive) components capturing different punishment and redistribution techniques. The components are formulated in such a way that the question of whether a rule set satisfies it can unambiguously be answered by yes or no.

1a. Punishment and redistribution (pun): the rule set provides for either destruction of a part of a player i 's points (punishment), or a reallocation thereof in favor of at least one other player j (redistribution), conditional on i 's behavior.

This condition was fulfilled for 75% of all rule sets. Remarkably, in the first tournament, all rule sets include either punishment or redistribution, or both. However, only half the sets feature any of them in the final tournament, and those that do, use redistribution. In other words, ‘pure’ punishment is no longer observed in the final tournament, but ‘punishment’ through redistribution or rule cost assignment conditional on the player’s contribution is. While this could be attributed to details in the cost scheme implemented, it cannot be ignored that the frequency of use of these mechanisms steadily declines over time.

20. See for example Carpenter and Matthews (2009), Denant-Boemont *et al.* (2007), Fehr and Gächter (2000, 2002), Nikiforakis (2008), Ostrom *et al.* (1992), Yamagishi (1986), for experimental studies, and Boyd *et al.* (2003), Dreber *et al.* (2008), Hauert *et al.* (2007), Henrich and Boyd (2001), for theoretical approaches.

To obtain a better understanding of the *pun* rules, we introduce three further classifying variables, the first of which relates to the frequency with which the deduction regimes played a role in the game:

1b. Punishment or redistribution in every period (punEP): punishment or redistribution (as described in 1a.) takes place in every single round.

This category makes the distinction between rules with roundly punishment or redistribution and those implementing them only periodically or even only once. A rule set exhibiting *punEP* provides for roundly deduction of points, a characteristic that was displayed by 13 out of 24 rule sets. In terms of rule sets which exhibit characteristic *pun*, the fraction of rule sets with *punEP* decreases from six out of six (eight out of ten in tournaments one and two) to two out of three (five out of eight in tournaments three and four).

Another distinction can be made with regard to players' influence on the points to be deducted:

1c. Redistribution endogeneity (redEnd): punishment or redistribution (as described in 1a.) is administered by players themselves rather than automatically.

Most rule sets (16 out of 18 rule sets with *pun*) proposed by our rule groups deprived the players from any influence on the points to be deducted, once the contribution decisions had been taken. Only in one instance, punishment was in the form of peer-to-peer punishment as in the typical public goods experiments with punishment,²¹ and in the other rule set, players contributing more than 14 tokens were allowed to jointly decide (through a voting procedure) on the allocation of rule costs among the remaining players.

Finally, we distinguish rules that involve a single, concentrated reward payment at the end of a session as a special case of a redistributive rule²² :

1d. 'Big bonus' (BB): redistribution takes on the form of a fund paid into and one distributive action (the allocation of a 'jackpot') at the end of round 25.

There were three rule sets that involved roundly payments into a fund that was awarded to the highest-contributing player at the end of the respective session (with an equal-split rule in case of a tie). All of them were designed by the same rule group.

3.2. Feedback on individual behavior

Several studies (e.g. Milinski *et al.*, 2006; Sommerfeld *et al.*, 2007) have shown that reputation-building opportunities may be very effective in promoting cooperation. A necessary prerequisite for being able to build a reputation is having some kind of identity. Bohnet and Frey (1999) show that identification alone suf-

21. For example Fehr and Gächter (2000, 2002).

22. Recall that rule sets involving a 'big bonus', like all other rule sets, could not be designed in a way that would make the defecting equilibrium disappear.

fices to increase the degree of pro-social choices in prisoner's dilemma and dictator games. Rule sets could allow for identification, for example in terms of a fixed player ID to subjects, but in our seminars, this happened only in a minority of the cases.

2a. Feedback on individual behavior (ID): players are informed not only on the sum of contributions, but also on individual contributions.

In fact, subjects were provided with identity-based feedback only in seven rule sets. One set provided detailed feedback on contributions without identity numbers, and the remaining two thirds of the rule sets did not inform players on individual contributions. In two cases, players faced a rather special situation we also want to separate: players under these two rule-sets did not receive any feedback on contributions at all.

2b. No feedback on contributions (noFB): the sum of contributions is concealed from the players.

In both cases, this rule was coupled with a 'big malus', a severe punishment action of the lowest-overall-contributor at the end of the game, unless a prespecified cumulative contribution level had been achieved by the group.²³ In the first instance of the rule-set, a considerable amount of contributions was achieved (1,271 tokens, as compared to the 834 required to circumvent the 'big malus'). Subsequently, the group modified their rule-set, asking for a group-contribution of 1,667 tokens, which the play-group failed to comply with by a margin of 120. The application of the 'big malus' led to a disastrous result in terms of efficiency and the group abandoned this strategy.

3.3. Communication: subject-to-subject and 'lawmaker'-to-'people' (framing)

Another tool that has proven to be very effective in inducing cooperation in social-dilemma situations is communication (e.g. Brosig *et al.*, 2003; Cason and Khan, 1999; Isaac and Walker, 1988; Ostrom *et al.*, 1994). Indeed, our subjects also implemented forms of communication, but on a vertical rather than a horizontal level. This means that they mostly abstained from peer-to-peer-communication and instead used a communication of the lawmaker to the players.

3a. Player-to-player communication (CommP2P): there is some form of communication between players, e.g. in form of a signaling opportunity such as in cheap-talk agreements.

No group implemented an open-communication component in the form of a virtual chat-room. Only in the final tournament, one rule group in each seminar opted to allow players to engage in signaling behavior through a unanimity vote on a (non-binding) covenant, with roundly renewal (contingent on that there

23. We do not introduce an additional 'big malus' category, as the two rule-sets are already uniquely classified by the *noFB* category. No other rule-set employed a 'big malus', and thus, such a category would not add any information.

had not been more than two breaches in the past) and requiring subjects to contribute fully in one case, and periodical votes coupled with a proposed minimum of 15 tokens in the other.²⁴

On the other hand, communication by the lawmakers to the players was common. Communication of this form was implemented in terms of (moral) appeals or a framing of the public goods situation.²⁵ This is especially remarkable because all subjects were – by then – experienced players of laboratory public goods provision and this was common knowledge. Furthermore, when playing under those rules they knew that the framing was artificially imposed and had no relevance to the game actually played.

3b. Some frame (frame): the game is played after the issuing of an appeal (to moral sentiments, the advantageousness of the social optimum, or general social-welfare considerations), with a background story, or even with feedback conditional on behavior.

Less than half of all rule sets did without a framing (13 out of 24 had some framing). Rule components classified as falling into this category include priming attempts, such as the display of roundly changing statements similar to ‘One cannot live without trusting others’. Further, they include individual feedback conditional on behavior such as messages ‘To live means believing in something. In our case, in the community, thanks for that!’ in case of a full contribution. Or framing in the narrower sense, embedding the contribution decision in contexts like the building of a school in Afghanistan or public goods arising in a local neighborhood.

To further distinguish rules making use of a *frame* component, we separate those containing a one-time appeal at the beginning only from those envisaging repeated messages.

3c. Frame in every period (frameEP): an appeal, a background story, or behavior-conditional feedback is displayed in every single round.

Out of the 13 rule-sets containing a *frame*, six also had *frameEP* as a characteristic.

3.4. Participative elements

Experimental studies like Ostrom *et al.* (1992) and Tyran and Feld (2006) have shown that increased participation opportunities, such that players may influence the situation they are facing, lead to more cooperative outcomes. We were curious to find out whether our subjects would anticipate this and implement features of endogenous rule adaption or even rule change.

24. The latter rule group envisaged the possibility of changing the minimum requirement in case of general adherence to the covenant; however, their covenant did not succeed in inducing the expected cooperation.

25. See Cookson (2000) and Ross and Ward (1996) on positive effects of framing, for studies not finding a framing effect, see Rege and Telle (2004) or Rutte *et al.* (1987).

4. *Endogeneity (end)*: the rule-set provides for institutional change as a consequence of either player behavior or a vote.

Indeed, such features are used increasingly towards the end. However, only one fifth of all rule sets contain them and in one of the seminar groups, such elements are introduced only in the last tournament. Examples of endogenous rule features are a vote on removing any features that go beyond the basic game after the eighth round or the introduction of peer punishment contingent on players paying for a monitoring institution.

3.5. Omitted rule elements

The list of rule components that were discussed within the rule design phase but not implemented contains quite invasive rules. Groups considered, for example, to enforce a certain contribution on subjects or to completely exclude non-cooperators from the game. Such harsh rules, however, never made it from the group discussion to the actual implementation, although they were not prohibitively expensive.

The incidence of the rule components summarized under the four component categories traced over the course of the seminar is depicted in Figure 1. As mentioned before, we can see a clear trend away from *pun* rules, as well as from individual feedback, which was not a widespread rule component even in the beginning. Communication in the form of frames or appeals as well as of signaling opportunities, on the other hand, was used increasingly, as was endogenous rule change.²⁶

In a second step, we looked at typical combinations of rule components – or at combinations that typically would not be combined by our subjects.²⁷ The main result to be taken from this analysis is that all combinations but one are implemented more than once; the only combination that has not been implemented was that of communication (*comm*) coupled with an individual-feedback (*ID*). In other words, those engaging in the psychological techniques of framing, priming, and appealing explicitly choose to render others' individual behavior opaque.

A table classifying the different rule sets according to our characteristics can be found in Appendix S2 (see Table B2). The classification is also used for the typicity analysis reported in section 5.²⁸ Before we do that analysis, we explore the performance of rules exhibiting a certain component in section 4, in order to assess whether the rule-adjustment process went 'in the right direction'.

26. In Figure 1, as in the ensuing analysis, we re-aggregate the subcomponents to main components, as this is enough to convey the important insights. The same steps of analysis using the subcomponents were taken for an earlier working-paper version of this paper and can be found in Appendix S3. For the feedback component we use 'no feedback on individual behavior (*noIndFB*)' as the main category, while for the communication component, we define every rule set to fall into the main category 'communication (*comm*)' that has either *commP2P* or *frame*.

27. Tables summarizing this analysis can be found in Appendix S3.

28. To run a typicity analysis, characteristics displayed by less than half of all strategies have to be reverted (such that, e.g., *ID* becomes *noID*, i.e., 'subjects are *not* informed on individual contributions but only on their sum'), for technical reasons outlined in Kuon (1993).

Figure 1 Frequency of components by rule set

4. RULE PERFORMANCE AND RULE ADAPTATION

In the preceding section, we categorized subjects' rule sets and established the frequencies with which rule components were used in the different tournaments. In this section, we set out to evaluate those components' contribution to the performance of the rule sets they are incorporated in.

Figure 2 gives a first hint at an answer, displaying boxplots of the efficiency levels achieved by all rule-sets employing a certain main component. Judging by Figure 2, one may argue that having characteristics like endogeneity (*end*), no individual feedback (*noIndFB*) or communication (*comm*) in one's rule set seems to be a good idea, while employing punishment (*pun*) does not seem to lead to a good performance. To examine this conjecture more closely, we analyze our rule-sets using a simple distribution-free measure of success. We rank the rule sets by efficiency (with the highest rank corresponding to the highest efficiency) and calculate the average rank of all rule sets displaying a certain characteristic. Comparing this to the median rank of all rule sets provides us with an index for the discriminatory power of the component. Division by $(n-1)/2$, where n is the total number of rule sets, finally normalizes the index to lie within the interval $[-1,1]$. This index measures a component's contribution to the efficiency gained by the rule set. This is why we term the index shortly as the *component-efficiency index*. A component-efficiency index of 1 implies that only the best-performing rule set displays this component, while a component-efficiency index of -1 means that the characteristic is only made use of by the worst-performing rule set. A characteristic that is employed in every strategy would lead to a component-efficiency index of 0, as would any component for which the sum of rank deviations from the median above the median is equal to the sum of those below

Figure 2 Boxplots of achieved efficiency levels, by rule components

it. If R_i denotes the rank of rule set i , and Y_j is the set of all rule sets that exhibit characteristic j , the component-efficiency index ρ_j can be expressed as follows:

$$\rho_j = \begin{cases} \frac{2}{n-1} \left(\sum_{y \in Y_j} \frac{R_y}{|Y_j|} - \frac{n+1}{2} \right), & |Y_j| > 0 \\ 0, & \text{otherwise.} \end{cases}$$

This formulation of a component-efficiency index is a rather simple approach assuming a linear separability of the performance of rule sets, neglecting any interaction effects between rule components. Note further that the index is based on a relative ranking of marginal component effects, and can therefore only explain deviations from the median efficiency rank. Nevertheless, a virtual rank ordering of rule sets by simply adding up the components' ρ_j yields a surprisingly good prediction for the rank ordering by efficiency: the Spearman correlation coefficient between the two rankings is $r_s = 0.49$ ($p = 0.0214$). If we take into account that not all components contribute to efficiency to the same extend, this is a surprisingly high correlation.²⁹ In Figure 3, we depict the component-efficiency indices for our rule characteristics, grouped by tournaments. The two components that seem to boast efficiency the most seem to be the use of communication (*comm*) and giving players a way to influence the rules they are playing under (*end*), while *not* providing subjects with feedback on individual contributions (*noIndFB*) seems to enhance

29. Running a mixed-effects estimate of normalized rank deviations from the median rank on the indices (with random effects for rounds, the interaction of rounds and seminar group, and rule groups), the rank correlation between the ensuing fitted ranks and the true ranks can be driven up further ($r_s = 0.88$, $p < 0.001$).

Figure 3 Component-efficiency indices: This index measures a component's contribution to the efficiency gained by the rule set. Panel (a): component-efficiency indices for rule sets 1–3. Panel (b): component-efficiency indices for rule set 4. Panel (c): component-efficiency indices for all rule sets. The numbers of individual rule-sets having a component are given in parentheses.

efficiency only slightly. In contrast, *pun* features do not seem to make up a clear picture across tournaments: while they tend to be more prevalent among worse-performing rule sets, this is notably different in the final tournament. It seems that our subjects have learnt to design efficient *pun* mechanisms by the end of the experiment. On the other hand, we see that the seemingly successful mechanisms of *comm* and *end* (Figure 3c) do not perform as well compared to final-set *pun* rules (Figure 3b). When compared to rule-sets from other rounds, however, they still score above average.³⁰

Based on our indices as a – however imperfect – measure of success, let us take a look at the rule changes over the tournaments as depicted in Figure 1. We observe that (i) punishment and redistribution systems lead to low efficiency (and only slightly higher contributions), which corresponds with the trend away from such mechanisms. The fact that the decline in the use of *pun* is slow and only partial is surprising especially in seminar I, where there was one rule group that managed to achieve the highest degree of efficiency within their seminar in three out of four tournaments – using rules that did not contain either punishment or redistribution. In fact, this was the only rule group to renounce the use of the *pun* component within this seminar group. In other words, participants in

30. Contrasting the component-efficiency indices to indices calculated in the same way but taking contribution levels instead of efficiency levels as the underlying metric, we find very similar results. We take this as an indication for the robustness of our rule-cost scheme.

the remaining rule groups fail to recognize that not making use of punishment may actually be a competitive advantage: even in view of the superior results achieved by their competitor group, they fail to mirror that group's superior strategy.³¹ In seminar II, rule sets abstaining from a 'monetary' sanctioning of behavior do not perform as well; in two out of three instances, they achieve only the second highest efficiency, the third set following one of the two just mentioned on third place.

In the end, is the use of punishment a bad idea? Maybe not: *pun* mechanisms as designed in the last round (i.e., redistribution of rule costs as the only such mechanism) increase efficiency.³² From Figure 1, we also know that the use of ID-based feedback tends to decline over time. Figure 3 shows that this is not as surprising as it may seem: the efficiency index for the first three rounds is $\rho_{noIndFB,1-3} = 0.09$. Hence, ID-based feedback does not seem to have a positive effect on cooperation as long as there is no explicit award or commendation, or the possibility to be rewarded in an ancillary game as for example in Milinski *et al.* (2006). Correspondingly, we observe the stated trend away from ID-based feedback, with three rule sets having the *noIndFB* component in the first and five sets in the final tournament.

Turning to components with some communicative element, be it between subjects or between the 'lawmakers' and their 'people', we observe a steady trend towards *comm* components. This corresponds well with the positive effect such components have in the first three rounds ($\rho_{comm,1-3} = 0.45$). Finally, having the 'people' participate in the shaping of their world through endogenous features tends to foster efficiency ($\rho_{end,1-3} = 0.24$).

4.1. Do 'lawmakers' change their rules for the better?

The final tournament was the most decisive for the students' success in the seminar and the intermediate tournaments were meant to improve the rule sets and their lower weights in the final grading enable experimentation. Participants have successfully used this opportunity. The median efficiency from the third to the fourth (final) tournament increases significantly ($p_{34} = 0.0625$, pair-wise two-tailed Wilcoxon signed-ranks test, where p_{xy} denotes the p -value for the test between tournament x and y), whereas it does not between the first and the second and between the second and the third, respectively ($p_{12} = 0.3125$ and $p_{23} = 0.625$). What did the rule groups do to achieve this improvement? To find an answer to this question and possibly shed some light on typical mistakes and typical improvements, we performed a typicity analysis based on the characteristics introduced in section 3.

31. This reluctance to follow the good example of their competitor group could, of course, easily be explained by postulating a psychological bias to focus on contributions. However, from our discussions with the subjects we got the impression that they were well-aware of the fact that high(er) contribution averages do not necessarily lead to higher efficiency, but have to be measured against potentially increasing rule costs.
32. Unfortunately, a meaningful statistical comparison using e.g. conditional average payoffs is not possible, given rules containing a feature and those not containing this feature within a seminar group are not independent.

5. TYPICITY ANALYSIS

The preceding analyses considered the rule components in isolation and estimated their contribution to the rules' efficiency. In the following, we analyze the interdependencies of the rule sets' characteristics with a typicality analysis. This type of analysis, introduced by Kuon (1993) and successfully implemented for example by Selten *et al.* (1997), is used to determine what a 'typical strategy' or, in our case, a 'typical rule set' is and what the typical characteristics of these rule sets are. The two measures are interdependently intertwined: the typical rule sets are the ones that carry the typical characteristics and the typical characteristics are the ones that occur in typical strategies. Technically the typicality analysis is the solution of an Eigen value problem, providing weights for the rule sets and the characteristics. Typicalities of characteristics add up to one and are all equal to one divided by the number of characteristics if the number of strategies containing each of these features is the same, such that the reciprocal value of the number of characteristics forms the natural reference point against which typicalities may be evaluated. On the other hand, the more typical a rule set's features are, the more typical is the rule set itself. More precisely, the typicality of a rule set is the sum of its characteristics' typicalities.³³ Once the typicalities of rule sets and characteristics are established, we can look for whether typicality is in any way correlated with efficiency and thus assess how well-targeted our subjects were in terms of their search for better solutions. Given we have found rule sets in the final tournament to perform better than those in the preceding tournaments, we run two distinct analyses for these two sets and compare their typicalities to find out how the typical rule sets differed. Table 2 lists the corresponding component typicalities.

What we see in Table 2 is that the change in rule component typicalities largely corresponds to the frequency changes we know from Figure 1. Also from Table 2, we observe that a positive component performance in rule sets 1 to 3 (i.e., $\rho_{j,1-3} > 0$) tends to lead to an increase in the component's typicality. Furthermore,

Table 2 Performance indices, as introduced in section 4 (lines 1 and 2); rule component typicalities over the first three and the final rule sets (lines 3 and 4). Note that the *end* characteristic has been reverted for the purpose of the typicality analysis, because – for technical reasons – the characteristics have to be formulated such that the majority of rule sets carry it. The performance indices reported in lines 1 and 2 correspond to the changed component formulation and may hence be different from those reported in section 4

	pun	noIndFB	comm	noEnd
$\rho_{j,1-3}$	−0.086	0.091	0.448	−0.019
$\rho_{j,4}$	0.333	−0.040	0.000	0.067
Rule sets 1–3	0.277	0.255	0.168	0.300
Rule set 4	0.176	0.333	0.285	0.206

33. For a thorough discussion of the mathematical properties of this method, cf. Kuon (1993).

we find an increasing diversity in rule sets: average rule-set typicality declines from 0.888 in the first rule-set round to 0.773 in the final one even if we base these averages upon final-rule-set typicalities. This is a clear sign of a ‘divergence’ of rule sets.

A correlation analysis reveals that there is no systematic relation between rule typicality on the one hand and either contributions or efficiency, on the other. This lack of a systematic correlation between rule typicality and efficiency may be due to the presence of initial ‘typical mistakes’ that are only hesitatingly done away with. ‘Typical improvements’ fall into two categories: on the one hand, *pun* rules are improved, decreasing their wastefulness through a switch from pure punishment to redistribution rules; on the other, we observe the increasing typicality of communication and participative features, as well as the decline in the typicality of ID-based feedback.

6. DISCUSSION AND IMPLICATIONS

This paper reports an experiment in which subjects act as rule makers for a social dilemma game. The novelty of the approach lies in endogenizing the institutional design process by transferring it to subjects who are free in ‘inventing’ any institutional setting and not bound to prespecified rules. The designers’ incentive for acting as ‘good’ rule makers is that their payment strongly depends on the social benefits that accrue from play under the designed rule. This experimental approach allows studying the rules that social planners facing a social dilemma ‘invent’ and the development of these rules over time. It expands our view on rules fostering cooperation and guides to important new avenues for future research.

We can summarize our findings as follows: (1) Institution designers combine two or more rule components instead of relying on single-component rule sets. (2) Designers regularly include appeals to moral sentiments or a framing of the game situation. (3) Designers predominantly choose rules that render information on individual contributions opaque but often provide aggregate contribution information instead. (4) Designers make extensive but decreasing use of punishment possibilities, predominantly not in the form of peer-punishment, but in the form of centralized punishment or redistribution. (5) Designers do not incorporate the possibility of assigning leaders, or ostracism.

The finding that institution designers frequently and successfully use framing is particularly noteworthy. The insight that the framing of the game may influence subjects’ behavior is not new. Examples are Sonnemans *et al.* (1998) who compare a ‘public goods’ game with an equivalent ‘public bad’ game, or Andreoni (1995) comparing a positive-frame investment in the public good and a negative-frame purchase of the private good, or framing the identical prisoner’s dilemma games as either the ‘Wall Street Game’ or the ‘Community Game’ (Lieberman *et al.*, 2004). However, in our experiment, the subjects playing under these rule sets were experienced players, equipped with a deep understanding of the logic of the game and completely aware that the chosen frames have no connection to reality. Nonetheless, framing was increasingly chosen and a successful means in achieving efficiency. While one of the groups opted to tell its subjects

the public good was a school in Afghanistan, another group had its subjects play in a virtual neighborhood. Yet other rule sets displayed a ‘slogan of the round’ such as: ‘Only within the community, people are beings conscious of their strength’, or moral appeals directly asking subjects to contribute. The success of these design components shows that attempts to activate social norms through appeals or even general moral statements tend to be more than ‘just words’: they seem to foster cooperation even amongst casehardened addressees. On a more general level this calls for further research to better understand the advantages and downsides of frames and appeals. A recent example is the study by Cohn *et al.* (2014), who show that a significant proportion of bank employees become dishonest when their identity as bankers is rendered salient.

A second noteworthy finding is that designers render information on the other players’ provisions rather opaque – and that this tends to be a successful strategy. Conditionally cooperative subjects align their provision with what they believe others will contribute. Providing them with detailed information on past contributions of their peers may yield the most precise basis for the calculations. Yet, institution designers increasingly preferred to establish a certain degree of opaqueness by just providing the average contribution, as this turned out to be successful in enhancing cooperation. This is well in line with the observations of Hoffmann *et al.* (2013), which show that not providing the above average contributors with feedback about the low contributions of others stabilized contributions to a public good on a high level. Nonetheless, hiding detailed contribution information and traceable identifications excludes reputation concerns that have been shown to be powerful in increasing cooperation (Milinski *et al.*, 2002). The interaction of these effects is not yet fully understood and calls for future research.

A third noteworthy observation is that punishment, in various disguises, was a focal design element in almost all initial rule sets. This finding underlines the importance of punishment rules in public goods settings. Remarkably, however, punishment mechanisms were not designed in the form of peer punishment, but rather in the form of prespecified rules of deduction and/or redistribution contingent on complying with provision targets. These provision targets were either fixed levels (e.g. full provision) or contingent on the other group members (e.g. not being the lowest-contributing player). This third observation parallels a scholarly interest in centralized punishment mechanisms that has gained momentum only very recently (with the two main lines of research sparked off by Putterman *et al.*, 2011 and Sigmund *et al.*, 2010). This research suggests that whether centralized punishment is preferred over peer-punishment may be influenced strongly by the details of the punishment schemes (e.g. whether semi-centralized ‘pool punishment’ includes second-order punishment of non-punishing contributors, cf. Zhang *et al.*, 2014), and boosts the idea that its effectiveness is mediated strongly by whether it is implemented by a vote (e.g. Markussen *et al.*, 2014). Yet other important determinants of a preference for centralized punishment that may play out in our design are the higher degree of control the social planner has over the ‘executive authority’, or a belief that clearly specified rules will foster cooperation better than (arbitrary) peer punishment. To determine the relative importance of each of these factors is an important question to be addressed by future research.

ACKNOWLEDGEMENTS

We are deeply indebted to Reinhard Selten for proposing this research and for our intensive discussions throughout the project. Special thanks go to our seminar participants for their enthusiasm while taking part in this study. We also thank participants at seminars in Bonn and Erfurt for their inspiring comments, as well as our colleagues at the chair of microeconomics for their tireless help in testing the experimental software programs, as well as for the fruitful discussions along the way.

Address for correspondence: Bettina Rockenbach, University of Cologne, Albertus-Magnus-Platz, 50923 Cologne, Germany. Tel./fax: +49 221 470 866-4/-8; e-mail: bettina.rockenbach@uni-koeln.de

REFERENCES

- Andreoni, J. (1995), 'Warm-Glow versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments', *Quarterly Journal of Economics* 110, 1–21.
- Arbak, E. and M.-C. Villeval (2013), 'Voluntary Leadership: Motivation and Influence', *Social Choice and Welfare* 40, 635–662.
- Axelrod, R. (1984), *The Evolution of Cooperation*, Basic Books, New York.
- Bochet, O., T. Page and L. Putterman (2006), 'Communication and Punishment in Voluntary Contribution Experiments', *Journal of Economic Behavior & Organization* 60, 11–26.
- Bohnet, I. and B. Frey (1999), 'The Sound of Silence in Prisoner's Dilemma and Dictator Games', *Journal of Economic Behavior and Organization* 38, 43–57.
- Boyd, R., H. Gintis, S. Bowles and P. Richerson (2003), 'The Evolution of Altruistic Punishment', *Proceedings of the National Academy of Sciences of the United States of America* 100, 3531–3535.
- Brosig, J., J. Weimann and A. Ockenfels (2003), 'The Effect of Communication Media on Cooperation', *German Economic Review* 4, 217–243.
- Carpenter, J. and P. Matthews (2009), 'What Norms Trigger Punishment?', *Experimental Economics* 12, 272–288.
- Cason, T. N. and F. U. Khan (1999), 'A Laboratory Study of Voluntary Public Goods Provision with Imperfect Monitoring and Communication', *Journal of Development Economics* 58, 533–552.
- Cinyabuguma, M., T. Page and L. Putterman (2005), 'Cooperation under the Threat of Expulsion in a Public Goods Experiment', *Journal of Public Economics* 89, 1421–1435.
- Cohn, A., E. Fehr and M. A. Maréchal (2014), 'Business Culture and Dishonesty in the Banking Industry', *Nature* 516, 86–89.
- Cookson, R. (2000), 'Framing Effects in Public Goods Experiments', *Experimental Economics* 3, 55–79.
- Davis, D. D. and C. A. Holt (1993), *Experimental Economics*, Princeton University Press, Princeton.
- Denant-Boemont, L., D. Masclet and C. Noussair (2007), 'Punishment, Counterpunishment and Sanction Enforcement in a Social Dilemma Experiment', *Economic Theory* 33, 145–167.
- Dreber, A., D. G. Rand, D. Fudenberg and M. A. Nowak (2008), 'Winners Don't Punish', *Nature* 452, 348–351.
- Falk, A., E. Fehr and U. Fischbacher (2005), 'Driving Forces behind Informal Sanctions', *Econometrica* 73, 2017–2030.

- Fehr, E. and S. Gächter (2000), 'Cooperation and Punishment in Public Goods Experiments', *American Economic Review* 90, 980–994.
- Fehr, E. and S. Gächter (2002), 'Altruistic Punishment in Humans', *Nature* 415, 137–150.
- Fischbacher, U. (2007), 'Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments', *Experimental Economics* 10, 171–178.
- Fischbacher, U. and S. Gächter (2010), 'Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments', *American Economic Review* 100, 541–556.
- Guillen, P., C. Schwieren and G. Staffiero (2007), 'Why Feed the Leviathan?', *Public Choice* 130, 115–128.
- Gürerk, Ö., B. Irlenbusch and B. Rockenbach (2006), 'The Competitive Advantage of Sanctioning Institutions', *Science* 312, 108–111.
- Gürerk, Ö., B. Irlenbusch and B. Rockenbach (2014), 'On Cooperation in Open Communities', *Journal of Public Economic* 120, 220–230.
- Güth, W., M. V. Levati, M. Sutter and E. van der Heijden (2007), 'Leading by Example with and without Exclusion Power in Voluntary Contribution Experiments', *Journal of Public Economics* 91, 1023–1042.
- Hardin, G. (1968), 'The Tragedy of the Commons', *Science* 162, 1243–1248.
- Hauert, C., A. Traulsen, H. Brandt, M. A. Nowak and K. Sigmund (2007), 'Via Freedom to Coercion: The Emergence of Costly Punishment', *Science* 316, 1905–1907.
- Henrich, J. and R. Boyd (2001), 'Why People Punish Defectors: Weak Conformist Transmission can Stabilize Costly Enforcement of Norms in Cooperative Dilemmas', *Journal of Theoretical Biology* 208, 79–89.
- Hoffmann, M., T. Lauer and B. Rockenbach (2013), 'The Royal Lie', *Journal of Economic Behavior and Organization* 93, 305–313.
- Isaac, M. R. and J. M. Walker (1988), 'Communication and Free-Riding Behavior: The Voluntary Contributions Mechanism', *Economic Inquiry* 26, 585–608.
- Keser, C. (2000). 'Strategically Planned Behavior in Public Good Experiments', CIRANO Scientific Series No. 2000s-35.
- Keser, C. and R. Gardner (1999), 'Strategic Behavior of Experienced Subjects in a Common Pool Resource Game', *International Journal of Game Theory* 28, 241–252.
- Kosfeld, M., A. Okada and A. Riedl (2009), 'Institution Formation in Public Goods Games', *American Economic Review* 99, 1335–1355.
- Kuon, B. (1993), 'Measuring the Typicalness of Behavior', *Mathematical Social Sciences* 26, 35–49.
- Ledyard, J. O. (1995), 'Public Goods: A Survey of Experimental Research', in: J. Kagel and A. Roth (eds.), *Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ, pp. 111–194.
- Liberman, V., S. Samuels and L. Ross (2004), 'The Name of the Game: Predictive Power of Reputations versus Situational Labels in Determining Prisoner's Dilemma Game Moves', *Personality and Social Psychology Bulletin* 30, 1175–1185.
- Maier-Rigaud, F. P., P. Martinsson and G. Staffiero (2005), 'Ostracism and the Provision of a Public Good – Experimental Evidence', *Journal of Economic Behavior and Organization* 73, 387–395.
- Markussen, T., L. Putterman and J.-R. Tyran (2014), 'Self-Organization for Collective Action: An Experimental Study of Voting on Sanction Regimes', *Review of Economic Studies* 81, 301–324.
- Masclét, D., C. Noussair, S. Tucker and M.-C. Villeval (2003), 'Monetary and Nonmonetary Punishment in the Voluntary Contributions Mechanism', *American Economic Review* 93, 366–380.
- Milinski, M., D. Semmann and H.-J. Krambeck (2002), 'Reputation Helps Solve the "Tragedy of the Commons"', *Nature* 415, 424–426.

- Milinski, M., D. Semmann, H.-J. Krambeck and J. Marotzke (2006), 'Stabilizing the Earth's Climate Is Not a Losing Game: Supporting Evidence from Public Goods Experiments', *Proceedings of the National Academy of Sciences* 103, 3994–3998.
- 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, 358–369.
- Ostrom, E. (1998), 'A Behavioral Approach to the Rational Choice Theory of Collective Action: Presidential Address, American Political Science Association, 1997', *The American Political Science Review* 92, 1–22.
- Ostrom, E. (2000), 'Collective Action and the Evolution of Social Norms', *Journal of Economic Perspectives* 14, 137–158.
- Ostrom, E., J. M. Walker and R. Gardner (1992), 'Covenants with and without a Sword: Self-Governance Is Possible', *The American Political Science Review* 86, 404–417.
- Ostrom, E., R. Gardner and J. A. Walker (1994), *Rules, Games, and Common-Pool Resources*, The University of Michigan Press, Ann Arbor.
- Potters, J., M. Sefton and L. Vesterlund (2005), 'After You-Endogenous Sequencing in Voluntary Contribution Games', *Journal of Public Economics* 89, 1399–1419.
- Putterman, L., J.-R. Tyran and K. Kamei (2011), 'Public Goods and Voting on Formal Sanction Schemes', *Journal of Public Economics* 95, 1213–1222.
- Rege, M. and K. Telle (2004), 'The Impact of Social Approval and Framing on Cooperation in Public Good Situations', *Journal of Public Economics* 88, 1625–1644.
- Rockenbach, B. and M. Milinski (2006), 'The Efficient Interaction of Indirect Reciprocity and Costly Punishment', *Nature* 444, 718–723.
- Ross, L. and A. Ward (1996), 'Naïve Realism in Everyday Life: Implications for Social Conflict and Misunderstanding', in: E. S. Reed, E. Turiel and T. Brown (eds.), *Values and Knowledge*, Lawrence Erlbaum, Mahwah, NJ, pp. 103–135.
- Rutte, C. G., H. A. M. Wilke and D. M. Messick (1987), 'The Effects of Framing Social Dilemmas as Give-Some or Take-Some Games', *British Journal of Social Psychology* 26, 103–108.
- Selten, R. (1967), 'Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopol-experiments', in: H. Sauermann (ed.), *Beiträge zur experimentellen Wirtschaftsforschung*, J.C.B. Mohr (Paul Siebeck), Tübingen, pp. 136–168.
- Selten, R., M. Mitzkewitz and G. R. Uhlich (1997), 'Duopoly Strategies Programmed by Experienced Players', *Econometrica* 65, 517–555.
- Selten, R., K. Abbink, J. Buchta and A. Sadrieh (2003), 'How to Play 3×3 Games – A Strategy Method Experiment', *Games and Economic Behavior* 45, 19–37.
- Sigmund, K., H. De Silva, A. Traulsen and C. Hauert (2010), 'Social Learning Promotes Institutions for Growing the Commons', *Nature* 466, 861–863.
- Sommerfeld, R. D., H.-J. Krambeck, D. Semmann and M. Milinski (2007), 'Gossip as an Alternative for Direct Observation in Games of Indirect Reciprocity', *Proceedings of the National Academy of Sciences* 104, 17435–17440.
- Sonnemans, J., A. Schram and T. Offerman (1998), 'Public Good Provision and Public Bad Prevention: The Effect of Framing', *Journal of Economic Behavior & Organization* 34, 143–161.
- Sutter, M., S. Haigner and M. G. Kocher (2005), 'Choosing the Stick or the Carrot? Endogenous Institutional Choice in Social Dilemma Situations', *Review of Economic Studies* 77, 1540–1566.
- Tyran, J.-R. and L. P. Feld (2006), 'Achieving Compliance When Legal Sanctions Are Non-Deterrent', *Scandinavian Journal of Economics* 108, 135–156.
- Vesterlund, L. (2003), 'The Informational Value of Sequential Fundraising', *Journal of Public Economics* 87, 627–657.

- Yamagishi, T. (1986), 'The Provision of a Sanctioning System as a Public Good', *Journal of Personality and Social Psychology Review* 51, 110–116.
- Zhang, B., C. Li, H. De Silva, P. Bednarik and K. Sigmund (2014), 'The Evolution of Sanctioning Institutions: An Experimental Approach to the Social Contract', *Experimental Economics* 17, 285–303.

Appendix S1. Information on the course of the seminar.

Appendix S2. Instructions for the basic game.

Appendix S3. Figures and Tables for the analysis on the basis of sub-components.