

Fostering cooperation through the enhancement of own vulnerability

Anita Kopányi-Peuker, Theo Offerman and Randolph Sloof*

November 27, 2012

Abstract

We consider the possibility that cooperation in a prisoner's dilemma is fostered by people's voluntarily enhancement of their own vulnerability. The vulnerability of a player determines the effectiveness of possible punishment by the other. In the "Gradual" mechanism, players may condition their incremental enhancements of their vulnerability on the other's choices. In the "Leap" mechanism, they unconditionally choose their vulnerability. In our experiment, subjects only learn to cooperate when either one of these mechanisms is allowed. In agreement with theory, subjects aiming for cooperation choose higher vulnerability levels in Gradual than in Leap, which maps into higher mutual cooperation levels.

JEL classification: D03, D81, D83

Keywords: prisoner's dilemma, cooperation, endogenous punishment.

Acknowledgements: This paper benefited from the suggestions of audiences at the University of Zurich, University of California Santa Barbara, Friedrich Schiller University Jena, the University of Paris 1, the ESA European Conference 2012 (University of Cologne), the 5th Maastricht Behavioral and Experimental Economics Symposium (Maastricht University), the ACLE Competition & Regulation Meeting 2012 (University of Amsterdam) and 2011 Workshop on Trust and Cultural Evolution (University of Valencia and ERI-CES). Financial help of the Research Priority Area Behavioral Economics of the University of Amsterdam is gratefully acknowledged.

*CREED, University of Amsterdam and Tinbergen Institute, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands. We are very grateful to CREED programmer Jos Theelen for programming the experiment. A.G.Kopanyi-Peuker@uva.nl; T.J.S.Offerman@uva.nl; R.Sloof@uva.nl

1 Introduction

In situations where private interests are at odds with the collective interest, the possibility to punish has proven to be an effective tool to support cooperation. In small groups, contributors are often willing to pay a small cost to punish free riders. This process helps to make free riders behave in agreement with the collective interest (see for instance Yamagishi, 1986; Ostrom et al., 1992; Fehr and Gächter, 2000; 2002). In practice, the possibility to punish free riders will often be limited though, because people are protected by property rights. For instance, a neighbor who refrains from contributing to the local public good can presumably be disciplined if his car is damaged, but the person who inflicts the damage has to face the possibility that she will be persecuted in court.

Here we consider a prisoner’s dilemma in which cooperation cannot be contracted and in which the possibility to punish may only endogenously become available. That is, players can only be punished if they voluntarily make themselves vulnerable in the first place. They may do so to signal that they are interested in pursuing mutual cooperation. If the other is willing to punish a free rider at a small cost, then the player’s signal to conditionally cooperate becomes credible. Theoretically and in an experiment, we investigate whether and how the possibility of enhancing one’s own vulnerability may foster cooperation in the prisoner’s dilemma. In particular, we want to know whether (i) people voluntarily make themselves vulnerable if they have the possibility to do so, (ii) whether cooperation is enhanced when players have made themselves vulnerable and (iii) whether it matters if the trust-building process occurs gradually or in one single step.

The trust-building process that we have in mind corresponds to how strangers are reported to build friendships. In his “*Moralia*”, Plutarch already described how a reciprocal exchange of secrets may lead to a relationship in which there is a fear of loss of trust (Gambetta, 2009, p. 66; Plutarch, 1992). Strangers who have exchanged secrets before they interact in a prisoner’s dilemma may refrain from free riding if they fear that this will trigger the other to publicly disclose the secret. Recent psychological research shows that feelings of intimacy develop in a dynamic process in which a person discloses personal information, thoughts and secrets and the partner responds in a likewise manner (Altman, 1973; Rotenberg, 1986; Dindia and Allen, 1992; Laurenceau et al., 1998). An important (but possibly unintended) by-product of such feelings may be that people have learned to trust each other when they subsequently interact in a prisoner’s dilemma.¹ Interestingly, Derlega et al. (1976) report that reciprocal self-disclosure is especially observed among strangers, as hypothesized by Altman (1973). Once relationships have been established and friends have learned

¹Gossip may also serve as a form of self-disclosure that promotes trust through the enhancement of vulnerability. If a worker communicates damaging information about a superior to a colleague, he faces the risk that the colleague reveals this to the superior. If instead the colleague reciprocates with another negative story, a bond may be formed which may help the workers to solve free rider incentives in the work place. See Sommerfeld et al. (2007) for a discussion of other functions of gossip.

to trust each other, other arguments may take over. For instance, in long term relations the behavior of friends may be disciplined by repeated game considerations.²

In this paper, we model this trust-building process in a stylized three stage game. In the first stage, the two players voluntarily decide upon the extent to which they make themselves vulnerable for punishment. In the second stage, after having been informed of each other's own possible punishment level, the two partners decide whether or not to cooperate in a prisoner's dilemma. Based on the observed outcome, each player decides in the third stage whether or not to punish the partner at a small cost. If a player decides to punish, the partner loses an amount equal to her chosen own possible punishment level in the first stage.

We consider two different trust-building mechanisms. In the "Gradual" variant, the players may build trust in small steps, while observing the partner's willingness to go along in this process. This variant has the advantage that a player can condition the own possible punishment level on the partner's possible punishment level. It agrees with the empirical observation that trust is often formed in small incremental steps. The possibility of "starting small" can be advantageous, as has been shown by Andreoni and Samuelson (2006) for a repeated prisoner's dilemma and by Weber (2006) in a team production game with Pareto ranked equilibria.³

In the "Leap" variant, players decide whether or not to take a leap in the dark by simultaneously choosing an own possible punishment level without the possibility of conditioning it on the partner's level. This variant may correspond to situations where players do not have the time to build trust in small steps, or where it is too costly to engage in a slow gradual process.

We first provide a theoretical analysis of the two trust-building mechanisms. We show that in either mechanism mutual defection as well as mutual cooperation can be supported in equilibrium. The cooperative equilibrium is selected in a process of iterated elimination of weakly dominated strategies in the Gradual mechanism. Such a process does not have a bite in the Leap mechanism though. Further, when weakly dominated strategies are eliminated iteratively, players choose higher own punishment levels in the Gradual variant than they would ever do in the Leap variant. In essence, the possibility to condition one's

²Other applications of the mechanisms studied in this paper include the use of hostages and trust-building among criminals. Since Roman times, hostages have been exchanged to enforce truces and treaties. In some cases hostages were voluntarily exchanged (Schelling, 1960, p. 135-137; Lee, 1991; Herrmann and Palmieri, 2005). Gambetta (2009) describes some examples where trust-building mechanisms are used to support criminal activities. For instance, mafia bosses have been reported to bring their wives to potentially explosive dinners to signal their own willingness not to start a shooting. Pedophiles are often asked to share compromising photos before they get access to a child-pornography website.

³The mechanism behind the result in those papers is quite different though. In Andreoni and Samuelson (2006), players differ in their taste for cooperation when they participate in a twice-played prisoner's dilemma. If the stakes are larger in the second stage game, players have a larger willingness to invest in the first stage to achieve cooperation in the second. In Weber (2006), coordination on the efficient equilibrium is facilitated if teams start small and new members observe the history of the team before they enter.

own vulnerability on the partner's vulnerability allows players to turn themselves into unconditional cooperators in the Gradual variant, provided that the partner does the same. Theoretically, such high levels of trust cannot be reached in the Leap variant.

We test the performance of the mechanisms in an experiment in which we use a harsh strangers environment. In a Control treatment where players do not have the possibility to make themselves vulnerable, cooperation is not sustained. In the treatments with a trust-building mechanism, after some initial aversion a significant fraction of the subjects actively employs the mechanism to achieve cooperative outcomes. The two trust-building treatments do not differ in the extent to which subjects use the mechanism. Instead, conditional on the mechanism being used, subjects choose higher own possible punishment levels in the Gradual variant than in the Leap variant. These higher own possible punishment levels subsequently map into a higher frequency of cooperative outcomes in the Gradual variant. At the same time, by helping to align the players' vulnerability levels, the Gradual variant diminishes the occurrence of miscoordination outcomes in which one player cooperates and the other defects.

In the second half of the experiment, cooperative outcomes are regularly observed with a mechanism but almost never in the Control treatment. After subjects have gained experience with the environment and the mechanism, subjects cooperate in only 4.3% of the cases in the Control treatment, while they cooperate in 26.5% of the cases in the Leap variant and in 36.3% of the cases in the Gradual variant. The Gradual mechanism performs well, even when it is put in the perspective of studies in which the punishment possibility was introduced exogenously. For instance, in Fehr and Gächter's (2000) strangers treatment, subjects' contributions increased from 18.5% to 57.5% when subjects were exogenously allowed to punish. So in absolute terms the effect of introducing the possibility to punish is somewhat bigger in Fehr and Gächter (2000), but not in relative terms.

The remainder of the paper is organized in the following way. Section 2 discusses the related literature. Section 3 presents the game and the theoretical analysis. In Section 4, we provide the experimental design and procedures. Section 5 presents the experimental results and Section 6 concludes.

2 Related Literature

Our study contributes to several strands of literature. One related literature considers how agents may achieve cooperation if they can write binding contracts. In the spirit of the work of Coase (1960), Varian (1994) proposes a simple two-stage compensation scheme that implements the efficient outcome as a subgame perfect equilibrium in a wide class of games. In a prisoner's dilemma, the mechanism amounts to the following. In the first stage, each player announces a non-negative amount that he will pay to the partner provided that the partner cooperates in the second stage. In the second stage, the players participate in a prisoner's dilemma. An announcement made in the first stage is binding, so

if the partner cooperates the contract is automatically carried out.

In an experiment, Andreoni and Varian (1999) show that the compensation mechanism performs fairly well. When the mechanism is introduced in the second part of the experiment, cooperation levels go up from 25.8% to 50.5%. So the data provide support for the mechanism, even though the mechanism does not completely weed out defective choices.

Charness et al. (2007) provide a more demanding test of the compensation mechanism. They extend Andreoni and Varian's experimental results in two ways. First, they consider variants of the prisoner's dilemma in which there is a substantial range of transfers that support cooperation in a subgame perfect equilibrium (in the prisoner's dilemma of Andreoni and Varian, there was essentially a unique subgame perfect equilibrium). Second, in their variants of the prisoner's dilemma subjects face a coordination game in the second stage of the experiment; both (C,C) and (D,D) can be supported as equilibria of the second subgame (in contrast, in Andreoni and Varian subjects have a dominant strategy to contribute provided that sufficient transfers in the first stage are made). The compensation mechanism even performs well in this harsher environment. Charness et al. find cooperation rates in the range of 43-68% when transfer payments are permitted compared to 11-18% in the control without the compensation scheme.

In a more complicated emissions trading game, Hamaguchi et al. (2003) investigate a punishment version of Varian's compensation scheme and find less support for the mechanism. In their game, there are many Nash equilibria, which may have inhibited coordination on the subgame perfect equilibrium. Chen and Gazzale (2004) show that supermodularity can be an important factor in the performance of a generalized version of the compensation scheme; they find that supermodular games converge significantly better than games far below the threshold of supermodularity. Falkinger (1996) introduces a related mechanism in which players are subsidized or taxed in accordance with how their contribution to the public good deviates from the mean contribution. In an experiment, Falkinger et al. (2000) confirm the potential of the mechanism to support cooperation in public goods.

Probably closest to us in terms of mechanism are Iossa and Spagnolo (2011). In a theoretical contribution, they discuss a contractual clause that may help players to achieve cooperation in a non-verifiable task. Prior to the interaction, the players sign a contract that allows them to punish each other to a certain extent at their own discretion. The punishment is without value to the punisher but harms the partner at whom it is aimed.⁴ The main difference between their

⁴The contractual clause studied by Iossa and Spagnolo (2011) thus has the defining characteristics of a hostage (cf. Schelling, 1960). Contractual hostages to facilitate trade were first studied theoretically in the economics literature by Williamson (1983). Inspired by his work, a small literature emerged in sociology that investigates so-called "hostage games". In essence, the hostage here is a contract that specifies a punishment that is automatically administered to the poster of the hostage if he acts opportunistically in the subsequent social dilemma. Raub and Keren (1993) show that posting of such hostages in a stage prior to a prisoners' dilemma may enhance cooperation. Likewise, Snijders and Buskens (2001) find that the postage of hostages may also help achieving cooperation in the trust game. Raub (2009)

and our approach lies in how the first stage is modeled. In their paper, players have the possibility to write a binding contract in the first stage, so that they can agree on a mutually beneficial arrangement without strategic uncertainty. In contrast, in our non-cooperative approach, players have to deal with the danger that their choice to enhance the own vulnerability is not matched by the partner. Among other things, this implies that in our approach non-use of the mechanism in combination with mutual defection cannot be excluded in equilibrium.

We contribute to this literature in two ways. First, an attractive feature of the mechanism that we study is that it does not require the possibility to write binding contracts. Second, we show that it may matter a lot for the outcome if players have the possibility to build trust in many small steps. Our conjecture is that this finding will generalize to Varian's compensation scheme, and that higher levels of cooperation may be achieved if the first contract writing stage proceeds gradually.

Our paper also contributes to a somewhat more remote strand of literature that studies how subjects choose among existing institutions that do or do not allow for punishment. In the public good game of Gürer et al. (2006), subjects can choose between a sanction-free institution and a sanctioning institution in which punishments and rewards are allowed. After an initial phase in which a small majority prefers the community without sanctioning, subjects massively migrate to the community in which sanctions are allowed and where high levels of cooperation are sustained throughout the experiment. In a follow-up experiment, Gürer et al. (2009) show that the sanctioning institution remains equally successful if punishments but no rewards are allowed, and they find that initial self-selection in the sanctioning institution is a key-factor explaining its success.

Kosfeld et al. (2009) investigate a social dilemma in which sanctions are imposed by a central authority. In their three stage game, players first decide whether or not to participate in the organization. Non-members of the organization can free ride on the members of the organization. In the second stage, players are informed of the number of participants in the organization and the organization is only actually formed if all members agree. In the third stage, all players decide about their contribution decision and members of the organization are automatically punished if they do not contribute their full endowment to the public good. In the experiment, the grand organization is usually formed in which everybody participates. This result is in line with the notion that social preferences matter and is inconsistent with standard theory that predicts the formation of incomplete organizations.

In the public good game studied by Sutter et al. (2010), subjects vote to play in a standard Voluntary Contribution Mechanism (VCM), a VCM that allows for rewards and a VCM that allows for punishments. When rewards and punishments are sufficiently effective, subjects prefer the VCM with rewards, even though the VCM with punishments is more effective. Sutter et al. find that endogenously chosen institutions trigger higher cooperation levels than if

surveys the relevant sociological literature.

the same institution is implemented exogenously.

These studies share the feature that once an institution is voluntarily implemented, its effectiveness is exogenously determined (by the experimenter).⁵ In contrast, in our study subjects themselves determine how effective a punishment by the partner will be.

3 The model and its predictions

3.1 Structure of the game

In this section we present a simple model on which our experiment is based. The model consists of a two-player game with three stages. In the first stage players unilaterally determine the level of possible punishment they might receive themselves later on in the game. We consider two different ways in which these possible punishments can be chosen. In the *Gradual mechanism*, players gradually increase their own possible punishment level as in an English auction while observing the other player’s process of increasing her possible punishment level. In particular, a player has the possibility to match the partner’s own punishment at any level where the partner decides to drop from the auction-like process. Under this mechanism, players can thus reciprocate small pieces of vulnerability step by step. In contrast, in the *Leap mechanism* both players simultaneously submit one final possible punishment level, just as in a sealed bid auction. Here a player thus takes a leap in the dark when choosing to make himself vulnerable.

After the first stage is completed, the resulting possible punishment levels – denoted $x_1, x_2 \geq 0$ – are revealed to both players. Players subsequently play the following prisoner’s dilemma game in the second stage:

Table 1: Payoff matrix of the prisoner’s dilemma

	C	D
C	c, c	s, t
D	t, s	d, d

where $t > c > d > s > 0$.⁶ Moreover, to facilitate the exposition of the equilibrium analysis, we also assume that $t - c \leq d - s$.

Having observed the behavior in the previous two stages players decide in the third stage whether they want to punish their partner or not. Punishment leads to a fixed cost of $p > 0$ for the punisher, which is assumed to be a small amount ($p < d$). Players can only punish their partner with the possible punishment level the partner determined herself in the first stage.

⁵Sutter et al. (2010) discuss more studies in which players choose between existing institutions.

⁶In our notation, c gives the payoffs when both players cooperate and d when both defect. Parameter t gives the ‘temptation’ payoff to defect when the other cooperates, while s reflects the ‘sucker’ payoff when a player is the only one cooperating.

3.2 Equilibrium predictions

We simplify the equilibrium analysis by analyzing the model in reduced form. In particular, we assume that players use the following punishment strategy: cooperators are never punished, while defectors are punished with probability one after outcome $C - D$ and with probability $\beta > 0$ after outcome $D - D$. It is also assumed that players punish only if the partner's punishment level is strictly higher than zero, i.e. players are unwilling to pay the punishment cost p if punishment does not cause any harm. Given assumed punishment behavior, players maximize their own payoffs.

Our ad hoc assumptions regarding punishment behavior are well in line with existing experimental evidence. People tend to punish (much) more if they cooperated but were betrayed, than if they both defected. Furthermore, people usually do not punish others who were nice to them and cooperated.⁷ Models of other regarding preferences, like e.g. Rabin (1993), Levine (1998), Fehr and Schmidt (1999) and Bolton and Ockenfels (2000), may provide a behavioral justification for our reduced form approach.⁸

Based on the assumed punishment behavior, we can collapse the second and the third stage into the following simultaneous move game:

Table 2: Expected payoff of the whole game given possible punishment levels

	C	D
C	c, c	$s - p \cdot I_{\{x_2 > 0\}}, t - x_2$
D	$t - x_1, s - p \cdot I_{\{x_1 > 0\}}$	$d - \beta(x_1 + p \cdot I_{\{x_2 > 0\}}), d - \beta(x_2 + p \cdot I_{\{x_1 > 0\}})$

Here $I_{\{x_i > 0\}}$ denotes the indicator function, equal to one if $x_i > 0$ and zero otherwise.

We first consider equilibrium behavior in the above collapsed (second stage) game. Let $\delta_i(x_i, x_j) = Pr(i \text{ chooses } D \mid (x_i, x_j))$ denote the probability with which player i defects, given possible punishment levels (x_i, x_j) chosen in the first stage. The following lemma characterizes players' equilibrium strategies in the second stage. All proofs are relegated to Appendix A.1.

Lemma 1 *If the other player j has chosen $x_j > 0$ in the first stage, the equilibrium strategy of player i (with $i \neq j$) in the second stage is as follows:*

- (i) *If $x_i < t - c \equiv \underline{x}$, then player i will be a defector, that is, $\delta_i^*(x_i, x_j) = 1$ for all $x_j > 0$;*

⁷These types of behavior have already been observed in other public good game experiments; see e.g. Fehr and Gächter (2000) or Sefton et al. (2007).

⁸Alternatively, a more direct approach as in Andreoni and Samuelson (2006) could be taken. Studying a prisoner's dilemma setting, they make direct assumptions on players' preferences for taking a given action (either cooperate or defect) that rationalize the common findings observed in experiments: viz. that people sometimes prefer to cooperate themselves, but differ in the strength of this preference, and value cooperation more when the other player cooperates as well.

- (ii) If $\underline{x} \leq x_i \leq \frac{d-s+(1-\beta)p}{\beta} \equiv \bar{x}$, then player i will be a conditional cooperator; that is, player i prefers to cooperate if the other player cooperates and prefers to defect if the other player defects;
- (iii) If $x_i > \bar{x}$, then player i will be an unconditional cooperator, that is, $\delta_i^*(x_i, x_j) = 0$ for all $x_j > 0$.

If the other player j has chosen $x_j = 0$ in the first stage, then for player i the same strategies hold with $\bar{x} = \frac{d-s}{\beta}$.

Lemma 1 shows that when a player chooses either a very low or a very high punishment level, he has a dominant strategy in the subsequent second stage game. For punishment levels in between \underline{x} and \bar{x} the player becomes a conditional cooperator, and prefers to coordinate to match the other player's choice.⁹

If there is at least one player with a dominant strategy, a unique equilibrium in the (collapsed) second stage subgame results. Only if both players turn out to be conditional cooperators, multiple equilibria exist: $D - D$, $C - C$, and a mixed one. In that case the two players face a coordination problem. Note, however, that it is unlikely that players will coordinate on $D - D$, given that they could have done so more cheaply by choosing $x_i = 0$ (and thus becoming a defector) in the first stage. They would then arrive at the same outcome but save at least $\beta\underline{x}$ in terms of expected punishments. As it appears, also mixing cannot occur on the equilibrium path. The intuition here runs as follows. The mixed equilibrium of the second stage subgame yields a player less than outcome $C - C$. A best response to the other player being a conditional cooperator is, therefore, to become an unconditional cooperator oneself by bidding above \bar{x} . In that case coordination on $C - C$ is secured and the mixed equilibrium is avoided.

The following proposition translates the above intuitions to the equilibria of the entire game.

Proposition 1 *In both versions of the mechanism there are multiple subgame-perfect equilibrium outcomes:*

- (i) $x_i = 0$, and $\delta_i^*(0, 0) = 1$ for $i = 1, 2$;
- (ii) $\underline{x} \leq x_i^* \leq \bar{x}$, and $\delta_i^*(x_i^*, x_j^*) = 0$ for $i = 1, 2$ (and $i \neq j$);

In the Gradual mechanism two additional sets of equilibrium outcomes exist:

- (iii) $\underline{x} \leq x_i^* \leq \bar{x}$ and $x_j^* > \bar{x}$, and $\delta_i^*(x_i^*, x_j^*) = 0$ for $i = 1, 2$ (and $i \neq j$);
- (iv) $x_i^* > \bar{x}$, and $\delta_i^*(x_i^*, x_j^*) = 0$ for $i = 1, 2$ (and $i \neq j$).

⁹Lemma 1 assumes that if player i is exactly on the border (i.e. $x_i = \underline{x}$ or $x_i = \bar{x}$), he behaves like a conditional cooperator. Obviously, ties between the three cases could have been broken differently, leading to essentially the same results.

The proposition only specifies the behavior on the equilibrium path; off path behavior is characterized in Appendix A.1. Note that there are infinitely many equilibria for both mechanisms. Yet the Gradual mechanism allows a larger set of equilibria, as it also includes equilibria in which either one or both of the players turn themselves into an unconditional cooperator in the first stage (cf. Proposition 1 (iii) and (iv)). This cannot occur in the Leap mechanism, because the best response to the other player choosing $x_{-i}^* > \bar{x}$ is to choose $x_i = 0$ oneself. However, in the Gradual mechanism it may happen, since there players can simultaneously increase their own possible punishment levels in small steps (and cannot go back). They thus can secure that they become an unconditional cooperator only if the other player becomes at least a conditional cooperator at the same time. This ensures that outcome $C - C$ is reached.

With multiple equilibria, the question of interest becomes in which equilibrium players will ultimately end up: either in the non-cooperative equilibrium of part (i), or in one of the cooperative equilibria of parts (ii) through (iv) of Proposition 1. Regarding the latter, in the cooperative equilibria of part (ii) the indeterminacy of the cutoff level for x_i^* to end up in the $C - C$ outcome is caused by the fact that if both players are conditional cooperators, multiple equilibria exist in the second stage subgame (as discussed above). Even though on the equilibrium path the players should always end up in the $C - C$ equilibrium, off the equilibrium path they may coordinate on one of the other second stage equilibria. Depending on this off path behavior, the equilibrium cutoff for x_i^* is higher or lower.¹⁰

In the Gradual mechanism players can increase their own punishment level step by step while observing whether the other player follows suit. One would therefore expect that coordination on the cooperative outcome should be easier. By using a forward induction like argument that refines the set of equilibria, the following proposition shows that this is indeed the case.

Proposition 2 *In the Gradual mechanism, only the equilibria described in part (iv) of Proposition 1 survive the iterated elimination of weakly dominated strategies.*

The intuition here is that in the Gradual mechanism players can profitably use a ‘wait and see’ strategy. Instead of jumping out first, a player can simply wait and see what the other player does. If the other player immediately drops out at zero, it is best to immediately follow suit. Otherwise, the players use the mechanism to achieve the cooperative outcome. In the proof of Proposition 2 it is shown that dropping out first below \bar{x} is (iteratively) weakly dominated. Equilibria in which both players stay in until at least \bar{x} are thus more focal.

For the Leap mechanism we cannot reduce the set of equilibria by eliminating

¹⁰The indeterminacy in parts (iii) and (iv) of Proposition 1 results from the fact that if both players always cooperate, punishments are never carried out. In that case players are indifferent between all possible punishment levels (within the given ranges). As a result, in e.g. part (iv) any x_i^* above \bar{x} can be supported in equilibrium.

weakly dominated strategies.¹¹ The reason is that there, unlike in the Gradual case, players cannot condition their own punishment level on the punishment level of their partner. They may therefore remain trapped in skeptical beliefs that the other will choose $x_{-i} = 0$, because there is nothing the other player can do to disprove such skepticism. To illustrate, suppose that player i wants to coordinate on a cooperative equilibrium if the partner also wants to coordinate on it, but prefers to choose zero if the partner does so. Skeptical beliefs then induce player i to choose a punishment level of zero, followed by defection in the second stage. If the other player thinks the same, players are stuck in the non-cooperative equilibrium even though both prefer a cooperative one. In contrast, in the Gradual mechanism players can freely disprove their partner’s skeptical beliefs, simply by waiting and not dropping out first at zero.

To sum up, in the Leap mechanism players face a larger amount of strategic uncertainty, as there appears to be no compelling argument that makes coordination on one of the cooperative equilibria focal. Players can more easily coordinate on cooperation in the Gradual case, by choosing higher possible punishment levels than in the Leap mechanism. The Gradual mechanism is therefore predicted to perform better in fostering cooperation.¹²

To conclude this section we briefly illustrate how models of other regarding preferences can rationalize the assumed punishment behavior. Take for instance the inequity aversion model developed by Fehr and Schmidt (1999). This model predicts that players who sufficiently dislike being behind, punish the defector after outcome $C - D$ because this reduces inequality. At the same time, the defector will typically prefer not to punish (assuming that the punishment he receives himself is not too high), as this increases inequality. If both players choose the same action in the prisoner’s dilemma, multiple punishment equilibria may exist. The intuition is that, depending on the choices made in the first two stages, players may want to mimic the other’s punishment behavior. They may then either coordinate on an equilibrium in which both punish, or another one in which they do not. After joint cooperation (outcome $C - C$) no punishment serves arguably as focal point; there is simply no compelling reason why players should punish. Yet after outcome $D - D$ it is much less clear what the focal punishment equilibrium is; one could either argue that punishment is appropriate given that the other defected, or make the case that punishment

¹¹Ben-Porath and Dekel’s (1992) result is very similar to ours, but obtained for other settings. They show that if players have the possibility to simultaneously burn money before playing a coordination game, the “worse” equilibrium cannot be eliminated by iterated elimination of weakly dominated strategies. Unlike their setup, in our model players do not literally burn money to show their intention for cooperation. Moreover, the strategic nature of our second stage game depends on the possible punishment levels chosen in the first stage: for low possible punishment levels it corresponds to a prisoner’s dilemma, while for e.g. in between possible punishment levels it becomes like a coordination game.

¹²The same general prediction holds when we assume instead that the incentive to defect is stronger if the other player cooperates than when the other player defects ($t - c > d - s$). In Appendix A.2 we provide the full equilibrium characterization for that case and show that coordination on a cooperative outcome then can be sustained as equilibrium outcome in the Leap mechanism only if β is sufficiently low. In contrast, equilibrium cooperation can still occur under the Gradual mechanism for all values of β .

is superfluous because both players did so. Our β parameter reflects this indeterminacy and can capture the average behavior of inequity averse players in practice. With Fehr-Schmidt style preferences, Lemma 1 and Proposition 1 remain qualitatively valid, except that the thresholds \underline{x} and \bar{x} will then depend on the inequity aversion parameters of the players involved. For example, players who dislike being ahead are more easily turned into conditional cooperators, i.e. will have a lower \underline{x} .

4 Experimental design and procedures

The experiment was conducted in the CREED-laboratory at the University of Amsterdam. In total, 144 subjects participated in 6 sessions. None of the subjects participated in more than one session. Subjects were mainly undergraduate students from various fields (e.g. economics, business, psychology, law). At the start of the computerized experiment, subjects received the instructions on their screen. Subjects read the instructions at their own pace and had to successfully answer some control questions testing their understanding before they could proceed with the experiment. The instructions for one of the treatments are included in Appendix B.¹³ Because the games in two of the three treatments are not simple, we used meaningful labels like punishment levels, cooperation and defection in the instructions of all treatments.

Subjects received a starting capital of 500 points at the beginning of the experiment, and they could earn additional points by their decisions. At the end of the experiment, their points were converted into real money; the conversion rate was such that 100 points corresponded to 1 euro. In an experiment that lasted between 1 to 2 hours, subjects earned on average 19.7 euros (with a minimum of 12.9 euro, and maximum of 29.6 euro).

There were 24 subjects in each session. In each of the 50 rounds, subjects were randomly rematched in pairs within their matching group of 8 subjects. Subjects were informed that they would never meet twice in a row with the same other subject. We organized two sessions for each of the three treatments. Thus, for each treatment we collected data of 6 independent matching groups.

We varied the trust-building mechanism between treatments. In the Control treatment, there was no mechanism, and the game consisted of only one stage in which the two subjects that formed a pair in a round played a prisoner's dilemma. Both subjects simultaneously chose between Cooperate (C) and Defect (D). At the end of the period, subjects were informed of each other's choices and received a payoff corresponding to the action pair. The payoffs in the prisoner's dilemma are listed in Table 3.

The other two treatments allowed for trust-building as in the model. In each of these treatments, each round had a three-stage structure. In the first stage, subjects determined their own possible punishment levels. After the first stage, subjects were informed about their partner's possible punishment level. In the second stage, subjects played the prisoner's dilemma game listed in Table 3. At

¹³The instructions for the other two treatments are similar and will be sent upon request.

Table 3: Payoffs in the prisoner’s dilemma in the experiment

	C	D
C	55; 55	5; 70
D	70; 5	25; 25

the end of this stage, they were informed about their partner’s action in the prisoner’s dilemma. In the third stage, subjects decided whether they wanted to punish their partner or not. The possibility to punish the partner was not limited to cases where the partner defected. If they wanted, they could also punish a partner who cooperated. If a subject decided to punish, the partner received a deduction in points equal to the his own possible punishment level of the first stage. At the same time, the subject who decided to punish incurred a cost of 4 points. After the third stage, punishment decisions of the own pair and earnings were revealed to the subjects. In each round, a subject’s earnings was equal to the payoff in the prisoner’s dilemma game, possibly diminished by the received punishment and/or the punishment cost (if punishment was received and/or given).

The two trust-building treatments differed in the way in which trust could be built in the first stage. In treatment “Gradual”, an automatized clock slowly raised the possible punishment levels, and subjects could stop the clock at the desired integer own possible punishment level. If one subject in a pair submitted a punishment level, the partner was immediately informed about this. In order to give the second subject the chance to submit the same punishment level as the first one, the clock was stopped for a few moments after the first subject had stopped. The clock only stopped for a given pair, not for everybody in the experiment. If the second subject decided not to stop at the same level, the automatized clock continued to rise until the second subject was satisfied and stopped the clock. To guarantee that 50 rounds could be run in a reasonable time frame, we imposed a maximum possible punishment level of 50 in this stage. When the clock reached this amount, subjects who still remained in the process automatically dropped out of the mechanism, and their possible punishment level was set equal to this maximum amount. Subjects were aware of this procedure.

In the first stage of treatment “Leap”, subjects essentially had to make a leap in the dark and choose an own possible punishment level without having any information about the partner’s own possible punishment level. Subjects simultaneously chose the own possible punishment level, a single integer number of at least zero and at most the maximum possible punishment level of 50. In the other stages, there were no further differences between Gradual and Leap.

With the parameters used in the experiment, players can turn themselves into a conditional cooperator if they choose an own possible punishment level of at least 15. The upper bound of the conditional cooperator range depends on β , and is strictly decreasing in it. This upper bound is at least $d - s = 20$, and it goes to $+\infty$ as β goes to 0. Above this threshold, players turn themselves

into unconditional cooperators.

In all three treatments, a social history screen with the 10 most recent completed rounds was always visible. In Gradual and Leap, this screen contained the possible punishment levels, the actions, and the punishment decisions of each pair in the subject’s matching group. The observations were ordered on the basis of the own possible punishment level of the first player in a pair.¹⁴ Appendix B provides an example of what a social history screen may look like. We decided to give this social history screen because the game is complex and we wanted to facilitate better understanding and speed up learning. In the social history screen of the control treatment we provided information about the actions of the pairs in a matching group.

5 Results

In Section 5.1 we present an overview of the experimental results. We start with the question how the achieved outcomes in the prisoner’s dilemma vary with the treatments. Then we analyze how individuals behave in the three stages of the treatments that use one of the mechanisms to build trust. First we deal with the extent to which subjects make themselves vulnerable in stage 1. Then we show how the possible punishment levels of stage 1 map into decisions in the prisoner’s dilemma in stage 2. Finally we discuss how actual punishment behavior in stage 3 depends on the behavior of stage 2. In Section 5.2 we zoom in on the dynamics in our data, and we provide an explanation of the results.

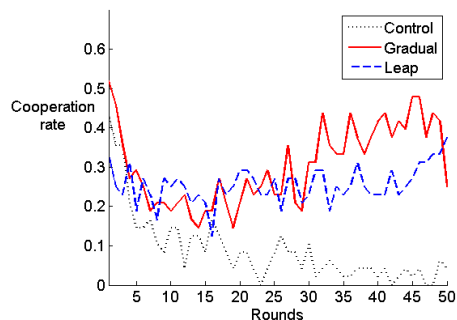
5.1 Overview of the experimental results

We start with the question whether the mechanism works as predicted. That is, do subjects cooperate more often when a mechanism is used and if so, is Gradual more successful than Leap in achieving this goal. Figure 1 displays the cooperation rates over time in the three treatments. In agreement with our conjecture that we employed an environment in which it is hard to sustain cooperation, cooperation levels fall dramatically in the Control treatment. In the second half of the experiment, cooperation rate is around 4%. Interestingly, in the Gradual treatment cooperation levels also fall rapidly in the beginning of the experiment. Apparently it takes time before the mechanisms become effective. After round 25 the cooperation rate grows to approximately 26.5% in Leap and 36.25% in Gradual, while it falls short of 4% in Control.

Table 4 presents the cooperation rates, and also the levels of the proportions of $C - C$ and $C - D$ outcomes and shows to what extent the differences between treatments are statistically meaningful. For the non-parametric test results reported in this paper, we employ a prudent testing procedure in which we

¹⁴If player 1’s own punishment levels were the same across pairs, the ordering was determined by the own possible punishment level of the player 2; if these numbers were also the same, then they were sorted on the basis of the cooperation decisions. It was randomly determined who in a pair was listed as player 1 and who was listed as player 2.

Figure 1: Cooperation rate over time



use average statistics per matching group as data points. Panel A of the table focuses on a comparison of the cooperation rates between treatments. The main message here is that the differences between each of the treatments that use a trust-building mechanism and the control treatment are significant in the second half of the experiment, while the difference between Gradual and Leap is never significant. The latter result is caused by the fact that not all matching groups in Gradual actively exploit the possibilities offered by the mechanism. The difference between Gradual and Leap lies more in the extent to which the mechanism is used if it is used. We come back to this in Section 5.2.

Although the increased cooperation rates already show that the mechanism is successful to foster cooperation, it is also important to examine how the mechanism leads to an increase in ending up in the cooperative outcome. To answer this question, panel B and C of Table 4 presents the proportion of the $C - C$ and $C - D$ outcomes. From panel B it can be seen that subjects ended up in the $C - C$ outcome significantly more often in both treatments with the mechanism compared to the Control treatment, both in the first and in the second half of the experiment, while again, the difference between the Gradual and Leap treatment is never significant. Panel C of the table presents the comparison of the proportions of $C - D$ outcomes. The proportion of $C - D$ outcomes diminishes over time in the three treatments. In the first half of the experiments, $C - D$ outcomes are observed approximately equally often in the three treatments, and all pairwise differences between the three treatments are not significant. In the second half of the experiment, the rate of $C - D$ outcomes drops substantially in Gradual and Control, but much less so in Leap. Remember that Gradual offers better opportunities than Leap to prevent $C - D$ outcomes, because in the former subjects may use the first stage to coordinate their possible punishment levels. In agreement with this intuition, subjects appear to be more successful in preventing $C - D$ outcomes in Gradual than in Leap. As a result, (weakly) significantly less $C - D$ outcomes are observed in Gradual than in Leap in the second half of the experiment.

Table 4: Percentage of cooperation, $C - C$ and $C - D$ outcomes

Treatment	First half (rounds 1-25)	Second half (rounds 26-50)
<i>Panel A. Cooperation rate</i>		
Control	13.67%	4.25%
Gradual	24.67%	36.25%
Leap	24.50%	26.50%
Control vs. Gradual	0.15	0.09*
Control vs. Leap	0.17	0.04**
Gradual vs. Leap	0.69	0.63
<i>Panel B. C-C outcome</i>		
Control	3.16%	0.50%
Gradual	13.50%	33.67%
Leap	13.67%	18.67%
Control vs. Gradual	0.04**	0.01**
Control vs. Leap	0.09*	0.05**
Gradual vs. Leap	0.63	0.26
<i>Panel C. C-D outcome</i>		
Control	21.00%	7.50%
Gradual	22.33%	5.17%
Leap	21.67%	15.67%
Control vs. Gradual	0.81	0.57
Control vs. Leap	0.94	0.11
Gradual vs. Leap	0.87	0.06*

Notes: **: significant at 5% level, *: significant at 10% level according to ranksum test with $n = 6$. The proportions of the given outcome and the p -values for tests for differences across treatments are displayed in the first and the second half of each panel, respectively.

We now turn to an analysis of individual behavior in the three stages of Gradual and Leap. We first deal with the extent to which subjects are willing to make themselves vulnerable in stage 1. Figure 2 illustrates how average own possible punishment levels change over time for each mechanism. Average possible punishment levels increase as time passes for both treatments. However, this increase is larger in Gradual than in Leap. Figure 1 and Figure 2 together suggest that subjects first try to achieve mutual cooperation without making use of the mechanism. Then when they find out in approximately round 10 that this does not work very well, they start using the mechanism which gradually enhances cooperation levels.

In the remainder of this section, we focus on the second half of the data. The reason is that we want to compare the treatments in the phase of the experiment in which subjects have learned about the potential usefulness of the mechanism in the game that they are playing. In Section 5.2 we come back to the data of the whole experiment when we describe the dynamics in the data. Figure 3 shows the distribution of possible punishment level pairs for each mechanism. The distributions of punishment level pairs differ substantially across treatments. As expected, possible punishment levels are much closer to each other in Gradual than in Leap; note that pairs of possible punishment levels are mainly along the diagonal in Gradual. There, subjects rarely choose a

Figure 2: Evolution of average punishment levels for each mechanism

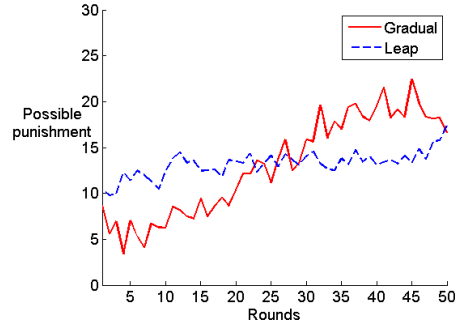
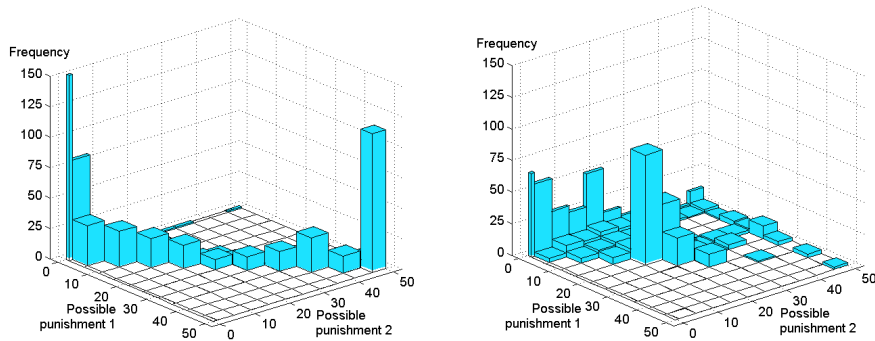


Figure 3: Distribution of punishment levels in the Gradual (left panel) and in the Leap treatments (right panel)

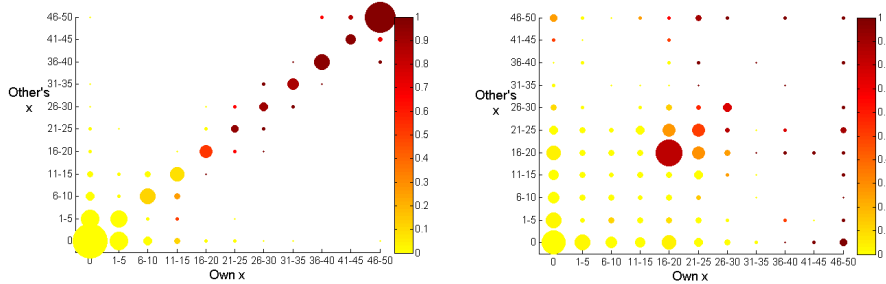


Notes: We define 11 categories for the 50 possible punishment levels. The first category equals the singleton set $\{0\}$, because many subjects choose that level. The other 10 categories cover the range from 1 to 50 with equal category length (that is, $[1;5]$, $[6;10]$, and so on). Each possible punishment level in each pair is assigned to a category, and pairs are plotted according to these categories. The figure is based on rounds 26-50. We plot the punishment levels of all pairs such that punishment level 1 \leq punishment level 2 (thus, by definition, there are no observations in the front part of the histogram).

higher possible punishment level after they have been informed of their partner’s possible punishment level. In Leap, subjects do not have the possibility to condition their possible punishment level on the level chosen by the partner, thus these levels correlate less. Overall, the correlation between the players’ possible punishment levels equal 0.89 in Gradual and 0.35 in Leap.

The figure shows that in both Gradual and Leap a substantial proportion of pairs of subjects coordinate on low levels of possible punishment (0, or in the interval $[1,5]$). Another interesting feature of the figure is that both distributions have, besides the origin, another but different spike. In Gradual, subjects often choose a very high possible punishment level (in the interval $[46;50]$), while

Figure 4: Cooperation probability for a given punishment level pair in the Gradual (left panel) and in the Leap treatments (right panel)



Notes: The 11 categories for the 50 possible punishment levels are the same as for Figure 3.

in Leap, many pairs manage to coordinate on a possible punishment level in [16;20]. The locations of these alternative spikes are in agreement with the theory presented in Section 2. If we condition on cases where the minimum possible punishment level of a pair is at least 15 (the theoretically threshold above which cooperation becomes possible), the average punishment level in Gradual (39.5) is almost twice as large as in Leap (20.6). The difference is significant according to a Mann-Whitney test ($p=0.03$). Thus, when pairs of subjects aim for cooperation, they choose substantially higher possible punishment levels in Gradual than in Leap.

Now we consider how the chosen possible punishment levels affect actions in the prisoner's dilemma. Figure 4 shows how often subjects cooperate for given combinations of possible punishment levels. The darkness of a circle corresponds to the extent to which subjects cooperate for a combination of the own and the partner's possible punishment level; the darker the circle is, the more likely it is that subjects choose to cooperate. The size of the circle indicates how often a pair of possible punishments is observed in a treatment. The following pattern emerges from the figures. Only when their own possible punishment level surpasses 15, subjects seriously consider to cooperate. In the quadrant where both players' own punishment levels exceed 15, subjects gradually cooperate more often when their own punishment level increases as well as when the other's own punishment level increases. The former result is in line with the notion that very high own possible punishment levels turn a player into an unconditional cooperator. The latter result agrees with the possibility that if a player chooses an own punishment level in the conditional cooperator range while the partner has become an unconditional cooperator, the player can become more confident that cooperation is the better choice. Interestingly, even in the quadrant where a player's own punishment level exceeds 15 while the partner's level falls short of 15, players tend to cooperate when they choose very high own punishment levels. This result supports the idea that with unilateral very high own punishment levels players face the danger of becoming an unconditional cooperator who can be exploited by the partner. This danger only materializes occasionally

Table 5: Logit model for the cooperation probability

<i>Dependent variable:</i> Cooperation	Model I	Model II	Model III
$I_{\{x_{own} \geq 15\}} * I_{\{x_{other} \geq 15\}}$	5.84 (1.31)***	2.11 (2.29)	-0.78 (2.77)
$x_{own} * I_{\{x_{own} \geq 15\}} * I_{\{x_{other} < 15\}}$	0.09 (0.03)***	0.12 (0.06)**	0.12 (0.05)**
$x_{own} * I_{\{x_{own} \geq 15\}} * I_{\{x_{other} \geq 15\}}$	-	-	0.15 (0.06)**
$x_{other} * I_{\{x_{own} \geq 15\}} * I_{\{x_{other} \geq 15\}}$	-	0.22 (0.06)***	0.20 (0.05)***
$Gr * I_{\{x_{own} \geq 15\}} * I_{\{x_{other} \geq 15\}}$	-0.42 (1.77)	-3.92 (2.95)	-1.15 (3.34)
$Gr * x_{other} * I_{\{x_{own} \geq 15\}} * I_{\{x_{other} \geq 15\}}$	-	0.07 (0.12)	-0.05 (0.14)
-Log likelihood	372.77	268.26	255.86
Number of panels (subjects)	69	69	69
Number of observations per panel	25	25	25

Notes: ***: significant at the 1% level, **: significant at 5% level, *: significant at 10% level. In Model I we regress cooperation decisions (cooperate: 1, defection: 0) on a dummy variable which is 1 if both players choose a possible punishment level higher than 15 ($I_{\{x_{own} \geq 15\}} * I_{\{x_{other} \geq 15\}}$), on the own possible punishment level provided that it is higher than 15, and the other's is lower than 15 ($x_{own} * I_{\{x_{own} \geq 15\}} * I_{\{x_{other} < 15\}}$), and on an interaction term with the first variable and a treatment dummy (Gradual: 1, Leap: 0). Model II includes the previous regressors, and the other's punishment level provided that both players chose a punishment level higher than 15 ($x_{other} * I_{\{x_{own} \geq 15\}} * I_{\{x_{other} \geq 15\}}$) and an interaction term between this variable and the treatment dummy. Model III includes Model II, and the own punishment level provided that both players chose a punishment level higher than 15 ($x_{own} * I_{\{x_{own} \geq 15\}} * I_{\{x_{other} \geq 15\}}$) as an additional regressor. Regression based on second half of the experiment (rounds 26-50). Std errors are in parentheses after the coefficients (Std errors are adjusted for 11 clusters).

in Leap.

In a logit model, we investigate more carefully how possible punishment levels map into decisions in the prisoner's dilemma. Table 5 reports the results of the logit regressions with fixed effects, in which we employ the player's id as panel variable. We cluster the standard errors by matching groups, because subjects' decisions are not independent from those with whom they are matched. Since we have a between subjects design and estimate a fixed effect model where the coefficients are only identified through the "within" dimension of the data (that is within subjects), we cannot directly estimate whether being in Gradual has a different effect on the decisions than being in Leap. Instead, to identify a possible treatment effect we include interaction terms between treatment and the main variables of interest. In all three models we regress the cooperation decision (cooperate: 1, defect: 0) on (i) a dummy variable which is 1 if both players choose a possible punishment level at least 15, otherwise 0, (ii) the own possible punishment level if it is at least 15 and the other player's is lower than 15, and (iii) the interaction between treatment and the first dummy variable.¹⁵

The first model does not include more variables. Its estimation result is that cooperation is significantly more likely if both players choose a punishment level of at least 15, that is, they both turn themselves into a conditional cooperator, as theory predicts. Further, when a subject's own punishment level exceeds the threshold of 15 while the partner's falls short of it, the subject is significantly more inclined to contribute when the own possible punishment level increases,

¹⁵We do not include an interaction term between treatment and the second variable because there are only a few observations in that region in the Gradual treatment.

in agreement with the idea that high own punishment levels may make players committed to cooperate. The second model additionally includes the other player's possible punishment level if both players choose a punishment level at least 15, and its interaction term with the treatment. When they have become conditional cooperators, subjects contribute significantly more often when the other chooses a higher own punishment level. Finally, in the third model we also add the own possible punishment level if both players choose a punishment level at least 15.¹⁶ The own punishment level also has a significant and positive effect on the cooperation probability if both punishment levels are at least 15. This again suggests that subjects may become more inclined to contribute when they have turned themselves into unconditional cooperators.¹⁷

It is important to note that the interaction terms with the treatment dummies are never significant which shows that in Gradual and Leap subjects respond rather similarly to given combinations of possible punishment levels.

Finally we discuss the question which choices in the prisoner's dilemma actually trigger punishments. The data are roughly in line with the assumptions of the theoretical model. For a start, own punishment levels of 0 rarely trigger punishments (< 1% of the cases). Conditional on strictly positive own punishment levels, Table 6 presents for each treatment how often subjects punish after a combination of actions in the prisoner's dilemma. Subjects primarily choose to punish after they cooperated while the partner did not (that is, in the $C - D$ outcome). These punishment frequencies are close to our theoretical assumption that the cooperator punishes with probability 1 if the other defects. In either treatment, cooperating subjects punish significantly more often when the partner defects than when the partner cooperates. If we combine all cases where the partner cooperates, subjects punish in only 8% of the cases in Gradual and in 5% of the cases in Leap. These numbers are not substantially higher than zero, providing support to our assumption that a cooperative partner is not punished. When both subjects defect, punishments occur, but only in a minority of the cases. Pooling the data of the two treatments, this occurs in approximately 13% of the cases. The frequency of punishment after $D - D$ increases to approximately 20% if we condition on cases where a subject's own punishment level is larger than 15.

Roughly the same punishment pattern emerges in the two treatments where punishments are allowed. In both cases, subjects most often punish if they cooperate and the partner defects, while they least often punish in cases where

¹⁶Here we do not add the interaction term with treatment, because this interaction term is highly correlated with the interaction term with the other's possible punishment level if both players choose a punishment level at least 15.

¹⁷In the reported regressions, we use the theoretical threshold of 15 as the lowerbound of the range in which players can turn themselves into conditional cooperators. We also ran the 3 regressions for any threshold level between 10 and 20. It turns out that the likelihood for model I is maximized for a threshold at 17 and for models II and III at threshold 16. Notice that with an own punishment level of 15 subjects are indifferent between cooperating and defecting if the partner cooperates. It turns out that subjects cannot turn themselves credibly into conditional cooperators with lower own punishment levels, which might have been the result if subjects dislike being ahead of the other.

Table 6: Punishment behavior

Own action	Treatment	Partner's action		Wilcoxon p
		Cooperate	Defect	
Cooperate	Gradual	0.06 (404)	0.93 (29)	0.04
	Leap	0.02 (220)	0.88 (78)	0.07
	p -value	0.62	0.13	
Defect	Gradual	0.37 (30)	0.14 (316)	0.89
	Leap	0.13 (88)	0.13 (472)	0.35
	p -value	0.45	0.52	

Notes: Percentage of given punishment (excluding cases with partner's punishment level zero) with the number of cases in parentheses. The p -value shows the test results of testing for the differences across treatments based on the Mann-Whitney test, whereas Wilcoxon p compares behavior within a treatment.

they and their partner cooperate. For none of the four possible action combinations in the prisoner's dilemma, the actual punishment probability differs significantly across treatments.

So far the following picture emerges from our data in Gradual and Leap. In both treatments, subjects either refrain from making themselves vulnerable or they make themselves vulnerable to a sizable extent. If they refrain from making themselves vulnerable, subjects primarily choose to defect. On the other hand, if they do make themselves vulnerable, they tend to do this to a substantially larger extent in Gradual than in Leap. The process of mutually monitoring allows subjects to coordinate on high own punishment levels in Gradual. Given a combination of own possible punishment levels, subjects choose similar actions in the prisoner's dilemma in both treatments. The higher the own possible punishment level is (given that it is already at least 15), the larger the probability of cooperation is. In addition, subjects' cooperation decisions respond positively to the other's possible punishment level. Finally, we do not observe systematic differences in the treatments in how combinations of choices in the prisoner's dilemma map into actual punishments. Therefore, the main difference between the treatments is that subject make themselves vulnerable to a larger extent in Gradual if they aim for cooperation. This feature of our data is in line with the theoretical analysis. Theoretically, mutual high own possible punishment levels can be supported in Gradual by a process of iteratively eliminating weakly dominated strategies. Such a process is impossible in Leap, where the best that subjects can do is to turn themselves into conditional cooperators. From this perspective, it makes a lot of sense that subjects choose higher own possible punishment levels in Gradual than in Leap.

5.2 Dynamics in the data and explanation of main result

In this section, we take a closer look at the dynamics in our data. In the previous section, it was shown that the positive effect of the trust-building mechanisms primarily emerges in the second half of the experiment. Even then, it is not the case that all subjects switch to using the mechanisms. To get a sharper

view on how and why play evolves as it does in our experiment, we classify each subject as user or non-user. We do this on the basis of a subject's behavior in the preceding 10 rounds.¹⁸ In Leap, we say that a subject is a user if the own possible punishment levels are 15 or higher in more than 5 times of the preceding 10 rounds. We employ 15 as a threshold to judge an own possible punishment level in a round because theoretically from this level players credibly signal that they are interested in achieving the cooperative outcome. In Gradual, the characterization is not as simple, because subjects who want to use the mechanism may revert to a low own possible punishment level after they find out that the partner chooses a low level. However, also in Gradual there are certain types of behavior that unambiguously characterize a user or a non-user. If a subject chooses an own possible punishment level of 15 or higher, the subject behaves as a user. If a subject chooses the own possible punishment level earlier than the partner, while the possible punishment level has not yet reached 15, the subject behaves as a non-user. Only when a subject submits a low punishment level (lower than 15) after the partner decided on the own punishment level, we cannot use the case to classify the subject because we do not know if the subject had been willing to stay in until a level of 15 or higher if the partner had done the same. When classifying a subject on the basis of the preceding 10 rounds, we simply ignore the latter cases and say that a subject is a user if and only if the subject behaves as a user in more than 50% of the remaining cases.¹⁹ Note that our procedure allows for the possibility that a subject's type changes over time.

Figure 5 illustrates how the relative frequency of types evolves over time in the two treatments where the mechanism can be used. The horizontal line at 24 corresponds to 50% of the subjects (in each treatment we have observations of 48 subjects). In Gradual the number of users increases considerably over time (starting from 11 users), and after round 37 there are more users than non-users. In contrast, in Leap the number of users starts at a higher level and increases less steeply. In either treatment the frequency of users seems to converge to a small majority of the subjects.

A trust-building mechanism is only useful if other subjects in a matching group employ it as well. In each round, we classify a matching group of 8 subjects as a user group if it has at least 4 users in the given round. In Leap, there are 4 user groups in more than 60% of the rounds (in the remaining rounds there are 3 user and 3 non-user matching groups). In contrast, there are 4 or 5 non-user matching groups till round 25 in the Gradual treatment, and there are almost always 3 user and 3 non-user groups in the second half of the experiment. This shows that the better performance of Gradual does not result from an increase in the proportion of subjects or the proportion of groups that actually adopt the mechanism.

Instead, it appears that user groups in Gradual are more successful than user

¹⁸It follows from this choice that we cannot characterize subjects in the first 10 rounds, but only from the 11th one onwards.

¹⁹If a subject's behavior does not identify the type in any of the 10 preceding rounds, we keep the type assigned in the previous round.

Figure 5: Number of users of the mechanism

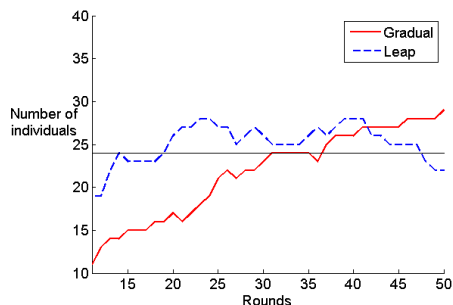
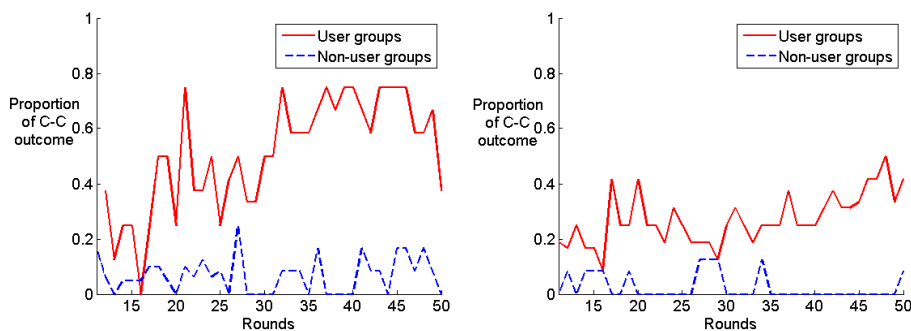


Figure 6: Proportion of $C - C$ outcomes in the matching groups in the Gradual (left panel) and in the Leap treatments (right panel)



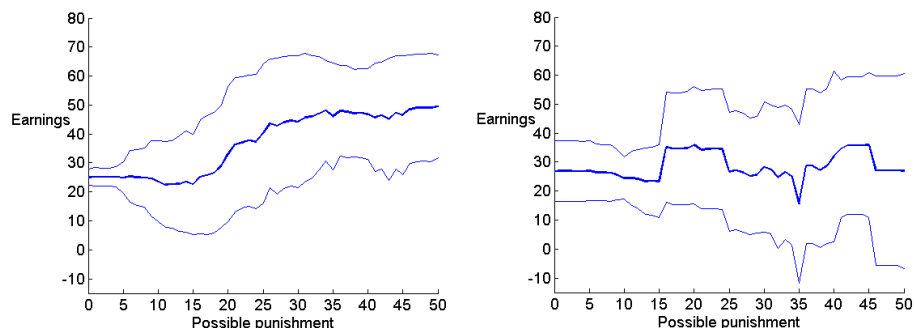
groups in Leap. Figure 6 shows the evolution of the percentage of cooperative $C - C$ outcomes over time for each treatment, separately for user and non-user groups. Unsurprisingly, user matching groups end up more often in the $C - C$ outcome than non-user groups in either treatment. Furthermore, there is no big difference between the efficiency of non-user groups across treatments. Remarkable is the substantial difference between user groups across treatments. In the second half of the experiment, the frequency of $C - C$ outcomes per user group is approximately twice as high in Gradual.

A plausible explanation is suggested by the observation that users choose higher possible punishment levels in Gradual than in Leap.²⁰ As we saw from the regressions of cooperation decisions (cf. Table 5), higher possible punishment levels above 15 significantly increase the likelihood of cooperation. These two observations together explain our finding that the Gradual mechanism is more efficient in facilitating coordination on the cooperative outcome.

Finally we shed some more light on the question of why subjects in Gradual

²⁰It is not the case that groups are more homogeneous in Gradual. We do not observe systematic differences in the treatments in how often a homogeneous group of at least 6 users formed.

Figure 7: Average earnings for the own possible punishment levels in the Gradual (left panel) and in the Leap treatments (right panel)



Notes: The data is smoothed by using weighted moving averages where the weights are the number of observations for the given own possible punishment level.

choose higher own possible punishment levels and on the question of why not everybody in a treatment uses the mechanism. For the second half of the experiment, Figure 7 shows how subjects' earnings depend on their own possible punishment level. The figure displays per treatment the average earnings as function of the own possible punishment levels, as well as the standard deviations around these averages.

In Gradual average earnings remain approximately constant until the possible punishment level of 15 which represents the theoretical threshold for cooperation to occur. From this level, average earnings steadily increase with the own punishment level. The picture is quite different for Leap; there subjects' earnings are already maximized for intermediate own punishment levels in the range of 15-25. So given the strategic uncertainty that subjects face, they are well advised to choose lower own possible punishment levels in Leap than in Gradual, just like they do.

In both treatments, the standard deviations around the averages increase with the possible punishment levels. So subjects face a clear trade-off; if they want to maximize expected payoff, they have to use the mechanism and choose very high own possible punishment levels in Gradual or intermediate own punishment levels in Leap. With this behavior they face higher risk than when they completely refrain from using the trust-building mechanism and set the own possible punishment levels equal to 0. This observation sheds light on the fact that not all subjects in a treatment switch to using the mechanism. Possibly some risk averse subjects settle for lower expected earnings with lower risk by refraining from using the mechanism.

6 Conclusion

A common finding in the psychological literature is that strangers often build relationships by the reciprocal disclosure of secrets. This process has consistently been shown to enhance feelings of intimacy. In this paper, we investigate the possibility that cooperation is fostered through a process of enhancing one's own vulnerability. By doing so, a player signals to his partner that he aims for cooperation, and provided that the partner is willing to punish an uncooperative act, the signal becomes credible. We consider the case that people can build trust conditionally in small incremental steps as well as the case in which this has to be done in a single unconditional leap of faith. Theoretically, cooperation can be supported in both cases. At the same time, a simple argument based on the iterated elimination of weakly dominated strategies suggests that higher levels of own vulnerability will be observed in the case where players can gradually build trust.

In the experiment, we observe that after some initial aversion players start using the possibility of building trust through the enhancement of own vulnerability. In the second half of the experiment, they do this to a similar extent in either variant of the trust-building process. Subjects coordinate substantially more often on the cooperative outcome with a trust-building mechanism than in the control where this is not allowed. So the trust-building mechanism works, albeit not perfectly. In this sense, its performance is comparable to other mechanisms such as Varian's (1994) compensation scheme.

In agreement with theory, we observe higher vulnerability levels when players can build trust gradually than when they have to do it in one leap. In the experiment, subjects' willingness to cooperate increases steadily with the own punishment levels after the threshold level needed for cooperation is reached. Thus, when subjects make use of the mechanism, higher mutual cooperation levels are achieved when trust is built gradually. At the same time, the gradual variant helps reducing the number of miscoordination outcomes where one partner cooperates and the other defects. These findings shed new light on why a slow, gradual process of building trust may be preferable.

Appendix A

A.1 Proofs

Proof of lemma 1:

(i) If $x_i < \underline{x}$ and $x_j > 0$, we have that both $c < t - x_i$ and $s - p < d - \beta(x_i + p)$, from which it follows that player i is better off by defecting, no matter what player j does.

(ii) If $\underline{x} \leq x_i \leq \bar{x}$ and $x_j > 0$, it holds that $c \geq t - x_i$ and $s - p \leq d - \beta(x_i + p)$. Player i then prefers to choose the same action as the other player does.

Case (iii) and the proof for $x_j = 0$ go analogously. \square

Proof of proposition 1:

We consider the two versions of the mechanism separately.

(a) *Leap mechanism.* We first show that there is no equilibrium in which $0 < x_i < \underline{x}$ or $x_i > \bar{x}$. If $0 < x_i < \underline{x}$, choosing D is a dominant strategy for player i . Realizing this, the other player j will choose D when $x_j \leq \bar{x}$ and C when $x_j > \bar{x}$. Player i is then punished for his defection with at least probability $\beta > 0$. He is therefore better off if he chooses $x_i = 0$ (and subsequently defect); this would not alter the other player's choice between C and D (which is still governed by $x_j \gtrless \bar{x}$), but avoids costly punishment. In case $x_i > \bar{x}$ player i will always cooperate, for all values of $x_j \geq 0$ (cf. Lemma 1). The best reply of the other player j is then to choose $x_j = 0$ and defect. But against this strategy of player j it is not a best response for player i to choose $x_i > \bar{x}$ and C ; he can earn more by choosing $x_i = 0$ and defect as well.

Two possibilities therefore remain: $x_i = 0$ and $\underline{x} \leq x_i \leq \bar{x}$. Suppose first that $x_i = 0$. Player i then defects for sure in the second stage. (cf. Lemma 1). Player j therefore obtains s by choosing C and $d - \beta x_j$ by choosing D . From $d > s$ she can get the most if she chooses $x_j = 0$ and D as well. Therefore, $x_1 = x_2 = 0$ followed by $D - D$ constitutes an equilibrium. (Off the equilibrium path behavior is governed by Lemma 1. A unilateral deviation of player i does not affect j 's behavior. With j being a defector, i 's response is fully determined by the lemma.)

Next suppose that $\underline{x} \leq x_i \leq \bar{x}$. From the above this necessarily requires that $\underline{x} \leq x_j \leq \bar{x}$ as well (note that against $x_j = 0$, choosing $\underline{x} \leq x_i \leq \bar{x}$ is not a best response). Both players are thus conditional cooperators. The (collapsed) second stage subgame then allows three equilibria: $C - C$, $D - D$, and a mixed equilibrium. The cooperative outcome $C - C$ strictly Pareto dominates the other two. Suppose that players coordinate on the mixed equilibrium. Then by deviating to $x_i > \bar{x}$ and making herself an unconditional cooperator (while player j remains a conditional cooperator), player i can secure outcome $C - C$. Therefore, the mixed equilibrium of the second stage subgame cannot occur on the equilibrium path of the entire game. For the same reason this also applies to outcome $D - D$. (Note that, besides $x_i > \bar{x}$, then also $x_i = 0$ would be a profitable deviation.) We thus obtain an equilibrium in which $\underline{x} \leq x_1^*, x_2^* \leq \bar{x}$ followed by $C - C$. Off the path behavior is such that after any $(x_1, x_2) \neq (x_1^*, x_2^*)$ with $\underline{x} \leq x_1, x_2 \leq \bar{x}$, both players choose D . This supports that players choose exactly (x_1^*, x_2^*) in the conditional cooperator range. Off path behavior outside the conditional cooperator range is again governed by Lemma 1.

(b) *Gradual mechanism.* Also in this case there is no equilibrium in which $0 < x_i < \underline{x}$. To see this, from Lemma 1 it follows that player i defects after choosing such punishment level. The other player j can then only achieve $\max\{s - p, d - \beta(x_j + p)\}$. For $x_j < \bar{x}$ it holds that $\max\{s - p, d - \beta(x_j + p)\} = d - \beta(x_j + p)$, which is decreasing in x_j . Player j 's best response is thus to drop out immediately after player i (such that $x_j = x_i$) and defect as well. Choosing $0 < x_i < \underline{x}$ thus yields player i a payoff equal to $d - \beta(x_i + p) < d$, strictly less than choosing zero would do.

Two possibilities therefore remain: $x_i = 0$ and $x_i \geq \underline{x}$. If $x_i = 0$ the same

reasoning as for the Leap mechanism applies; choosing $x_i = 0$ and then defecting is a best response against itself. In case $x_i \geq \underline{x}$ for $i = 1, 2$, each player is either a conditional or an unconditional cooperator. With two conditional cooperators three equilibria exist in the second stage subgame. But the same reasoning as for the Leap mechanism implies that only $C - C$ can occur on the equilibrium path of the entire game. If at least one of the players is an unconditional cooperator, this is immediate, as then the second stage subgame has only $C - C$ as equilibrium outcome. Hence $x_i^* \geq \underline{x}$ and $\delta_i(x_i^*, x_j^*) = 0$ (for $i = 1, 2$) are also equilibrium outcomes. The proposition divides this range into equilibria where both players are conditional cooperators (part (ii)), one is a conditional cooperator and the other one an unconditional cooperator (part (iii)), and equilibria in which both players are unconditional cooperators (part (iv)). \square

Proof of proposition 2:

We proceed in three steps:

(i) *Any strategy leading to dropping out first (strictly) between 0 and \underline{x} is (iteratively) weakly dominated.*

To show this, suppose player i drops out first somewhere between 0 and \underline{x} . In this case playing C in the ensuing second stage game is strictly dominated for player i . Realizing that player i will defect for sure, player j earns the most by dropping out immediately after player i and defect as well. All other responses of player j are therefore (iteratively) weakly dominated. Dropping out first between 0 and \underline{x} thus leads to outcome $D - D$ and punishment with probability $\beta > 0$. Dropping out immediately at zero and playing D then yields player i always weakly more.

(ii) *Any strategy leading to dropping out between \underline{x} and \bar{x} is (iteratively) weakly dominated.*

First, suppose that player j drops out below \underline{x} . Then player i cannot do better than dropping out immediately as well (see (i) above). Next suppose that player j drops out first somewhere between \underline{x} and \bar{x} . Then he turns himself into a conditional cooperator. If player i also drops out before \bar{x} , she becomes a conditional cooperator as well. The second stage subgame then allows three equilibria ($C - C$, $D - D$ and a mixed one) and the maximum payoff that player i can attain equals c . If instead of dropping out in the conditional cooperator range, player i chooses a possible punishment level above \bar{x} , she turns herself into an unconditional cooperator, as choosing D in the second stage game is then strictly dominated for her. Realizing this, player j will choose C in response (because choosing D is strictly dominated for player j). Staying in the mechanism until at least \bar{x} thus yields player i a payoff of c for sure and thus weakly more than dropping out in the conditional cooperator range. A similar argument applies regarding dropping out first between \underline{x} and \bar{x} ; staying in until at least \bar{x} is weakly better as this leads to outcome $C - C$ (and payoff c) for sure.

(iii) *Any strategy leading to dropping out first at zero is (iteratively) weakly dominated.*

If player i drops out first at zero, she has a dominant strategy to play D in

the second stage game. Realizing this, player j earns the most by dropping out immediately after player i and defect as well (i.e. all other responses of player j are (iteratively) weakly dominated). Dropping out first at zero thus yields a payoff equal to d . Next suppose player i does not drop out first at zero. If the other player does so, player i 's best response is to do so as well and outcome $D - D$ results. This gives player i the same payoffs as dropping out first. If the other player does not drop out at zero, from (i) and (ii) above it is then weakly dominated for both players to drop out before \bar{x} . Therefore, both will choose a possible punishment level above \bar{x} and outcome $C - C$ results. This gives both players a payoff equal to $c > d$.

From (i) through (iii) it follows that dropping out first below \bar{x} is weakly dominated. \square

For Online Publication: A.2 The case $t - c > d - s$

For expositional reasons we focused in the main text on the case where the incentive to defect is weaker if the other cooperates than if the other defects, i.e. where $t - c \leq d - s$. In this appendix we show that qualitatively the same predictions are obtained for the opposite case where $t - c > d - s$. Also then the gradual mechanism is predicted to perform better in fostering cooperation.

We first present the equivalent of Lemma 1 for the general case in which no assumptions are made on how $t - c$ compares to $d - s$.

Lemma 2 *If the other player j has chosen $x_j > 0$ in the first stage, the equilibrium strategy of player i (with $i \neq j$) in the second stage is as follows:*

(i) *If $x_i < \min \left\{ t - c, \frac{d-s+(1-\beta)p}{\beta} \right\} \equiv \underline{x}'$, then player i will be a defector, that is, $\delta_i^*(x_i, x_j) = 1$ for all $x_j > 0$;*

(ii) *If $\underline{x}' \leq x_i \leq \max \left\{ t - c, \frac{d-s+(1-\beta)p}{\beta} \right\} \equiv \bar{x}'$, then:*

(a) *if $t - c \leq \frac{d-s+(1-\beta)p}{\beta}$, player i will be a conditional cooperator; that is, player i prefers to cooperate if the other player cooperates and prefers to defect if the other player defects;*

(b) *if $t - c > \frac{d-s+(1-\beta)p}{\beta}$, player i will be a reverse cooperator; that is, player i prefers to cooperate if the other player defects and prefers to defect if the other player cooperates;*

(iii) *If $x_i > \bar{x}'$, then player i will be an unconditional cooperator, that is, $\delta_i^*(x_i, x_j) = 0$ for all $x_j > 0$.*

If the other player j has chosen $x_j = 0$ in the first stage, then for player i the same strategies hold with $\underline{x}' = \min \left\{ t - c, \frac{d-s}{\beta} \right\}$ and $\bar{x}' = \max \left\{ t - c, \frac{d-s}{\beta} \right\}$.

Proof of Lemma 2:

Cases (i), (iia) and (iii) are analogous to cases (i), (ii) and (iii) in Lemma 1. For the remaining case (iib), note that $x_i \leq t - c$ implies that player i prefers to choose D if the other chooses C while $\frac{d-s+(1-\beta)p}{\beta} \leq x_i$ implies that i prefers C if the other chooses D . \square

A first observation is that the analysis in the main text remains valid as long as part (iib) of Lemma 2 does not apply. The new part (iib) opens up the possibility that a player turns himself into a “reverse cooperator” if he chooses a punishment level in the intermediate range. A reverse cooperator prefers to do the exact opposite of what the other player does.

In the remainder we focus on the case not covered by the main text where $t - c > \frac{d-s+(1-\beta)p}{\beta}$. This corresponds to $\beta > \frac{d-s+p}{t-c+p}$.²¹ Players may then turn themselves into reverse cooperators. It can be easily seen that under the Leap mechanism they will never do so. This follows because the best response to the other player being a reverse cooperator is to choose $x_i = 0$ oneself (yielding the maximum payoff of t). This induces the other player to choose $x_j = 0$ as well. Hence no equilibria with reverse cooperators exist.

Things are different under the Gradual mechanism. Here equilibria in which both players are reverse cooperators do exist. The intuition runs as follows. If both players are reverse cooperators, the (collapsed) second stage subgame has three equilibria: $C - D$, $D - C$ and a mixed strategy equilibrium. Coordinating on one of the pure equilibria $C - D$ or $D - C$ cannot occur on the equilibrium path of the entire game, as the player supposed to cooperate earns only $s - p$ and would have been better off dropping out immediately at zero and defect, yielding d . Yet coordinating on the mixed equilibrium can occur. This holds because, unlike for the case in which the other player is a conditional cooperator (case (iia)), if the other is a reverse cooperator there is no incentive to become an unconditional cooperator oneself. This would namely lead to the disadvantageous $C - D$ outcome. Therefore, mixing is possible on the equilibrium path of the entire game (only) if expected payoffs exceed d , the minimum payoff a player can secure himself by dropping out immediately at zero. Earning at least d requires that in the mixed equilibrium, the other player j defects with sufficiently low probability. As $\frac{\partial \delta_j}{\partial x_i} < 0$, this in turn requires that player i himself chooses a high enough punishment level $x_i \geq \hat{x}$. Threshold \hat{x} is determined such that player i at least earns d in equilibrium. The following proposition makes the above intuitive predictions precise.

Proposition 3 *For all $\beta > \beta^* \equiv \frac{d-s+p}{t-c+p}$ there is a unique subgame-perfect equilibrium outcome in the Leap mechanism:*

(i) $x_i^* = 0$, and $\delta_i^*(0, 0) = 1$ for $i = 1, 2$;

In the Gradual mechanism two additional sets of equilibrium outcomes exist:

²¹For the case $t - c > d - s$ considered in this appendix, there always exist values of β such that $\frac{d-s+p}{t-c+p} < \beta \leq 1$. For $\beta \leq \frac{d-s+p}{t-c+p}$ the propositions presented in the main text apply.

- (ii) $\hat{x} \equiv \frac{(t-d)(d-s+p)-\beta p(c-d)}{(d-s+p)+\beta(c-d)} \leq x_i^* \leq t-c = \bar{x}'$,
and $\delta_i^*(x_i^*, x_j^*) = \frac{(t-x_j^*)-c}{(t-x_j^*)-c+(s-p)-[d-\beta(p+x_j^*)]}$ for $i = 1, 2$ (and $i \neq j$);
- (iii) $x_i^* > \bar{x}'$, and $\delta_i^*(x_i^*, x_j^*) = 0$ for $i = 1, 2$ (and $i \neq j$).

Proof of Proposition 3:

(a) *Leap mechanism.* From the proof of Proposition 1 it follows that there are no equilibria in which $0 < x_i < \underline{x}'$ or $x_i > \bar{x}'$ (just replace \underline{x} by \underline{x}' and \bar{x} by \bar{x}' in the proof of Proposition 1) and that $x_1 = x_2 = 0$ followed by mutual defection can be supported as equilibrium. It remains to be shown that if $\beta > \beta^*$, no equilibria with $\underline{x}' \leq x_i \leq \bar{x}'$ exist. For $\beta > \beta^*$ this range equals $\frac{d-s+(1-\beta)p}{\beta} \leq x_i \leq t-c$. In that case player i is a reverse cooperator. The best response of the other player j is then $x_j = 0$, yielding her the maximum possible payoff of t . But against $x_j = 0$, choosing $\frac{d-s+(1-\beta)p}{\beta} \leq x_i \leq t-c$ is not a best response, as choosing $x_i = 0$ would yield player i more.

(b) *Gradual mechanism.* Again it follows from the proof of Proposition 1 that there are no equilibria in which $0 < x_i < \underline{x}'$ and that $x_1 = x_2 = 0$ followed by mutual defection can be supported as equilibrium. So consider the case in which $x_i \geq \underline{x}'$ for $i = 1, 2$. Each player is then either a reverse or an unconditional cooperator. With two reverse cooperators three equilibria exist in the second stage subgame: $C-D$, $D-C$ and a mixed one. On the equilibrium path of the entire game $C-D$ and $D-C$ cannot occur; the player supposed to cooperate would then earn $s-p$, less than the amount d the player could at least get by dropping out immediately at $x_i = 0$. The mixed equilibrium may potentially occur though, because it can yield both players more than d .

In the mixed equilibrium the mixing probability δ_i of player i should make the other player j indifferent between C and D , i.e.: $(1-\delta_i) \cdot c + \delta_i \cdot (s-p) = (1-\delta_i) \cdot (t-x_j) + \delta_i \cdot (d-\beta p - \beta x_j)$. Rewriting this yields the expression for $\delta_i^*(x_i^*, x_j^*)$ in part (ii) of the proposition. Differentiating we obtain that $\frac{\partial \delta_i^*}{\partial x_j} < 0$ for $\beta > \beta^*$. A necessary condition for mixing to occur on the equilibrium path is that players at least earn d . The expected payoff of player j can be calculated from $(1-\delta_i^*) \cdot c + \delta_i^* \cdot (s-p)$. For this to exceed d it must hold that $\delta_i^* \leq \frac{c-d}{c-s+p}$. Plugging in the expression for δ_i^* and rewriting yields that $x_j^* \geq \frac{(t-d)(d-s+p)-\beta p(c-d)}{(d-s+p)+\beta(c-d)} \equiv \hat{x}$ is required (it holds that $\underline{x}' \leq \hat{x} \leq \bar{x}'$). To support $\hat{x} \leq x_i^*, x_j^* \leq \bar{x}'$ and subsequent mixing, off path behavior is governed by Lemma 2 for deviations outside $[\underline{x}', \bar{x}']$. After deviations $x_i \neq x_i^*$ within $[\underline{x}', \bar{x}']$, the for player i disadvantageous pure equilibrium follows (i.e. player i choosing C , the other player D).

One reverse cooperator together with an unconditional cooperator would lead to the $D-C$ outcome, which cannot happen on the equilibrium path (as the cooperating player gets less than d). Two unconditional cooperators can be supported on the equilibrium path by assuming that in the reverse cooperator range $[\underline{x}', \bar{x}']$ players coordinate on the mixed equilibrium of the second stage

subgame. Expected payoffs in this mixed equilibrium are below c , so neither player has an incentive to deviate in that range. \square

The equilibria of parts (i) and (iii) of Proposition 3 are as before, with only defectors and unconditional cooperators respectively. In the equilibria of part (ii) both players have turned themselves into reverse cooperators in the first stage. They subsequently use mixed strategies in the second stage subgame, with the probability of coordinating on $C-C$ (equal to $(1-\delta_1^*)(1-\delta_2^*)$) increasing in the punishment levels chosen in the first stage.

Overall we conclude that when the incentive to defect is stronger if the other player cooperates than when the other defects, coordination on a cooperative outcome can only be sustained in the Leap mechanism if β is sufficiently low (i.e. $\beta \leq \beta^*$ is needed). In contrast, under the Gradual mechanism cooperation continues to be an equilibrium outcome for all values of $\beta \in (0, 1]$. The Gradual mechanism is therefore predicted to perform better also when $t-c > d-s$. Note though that the case for cooperation under the Gradual mechanism then also becomes less focal, as an equivalent of Proposition 2 does not hold for $\beta > \beta^*$.

For Online Publication: Appendix B - Instructions

In this appendix, we present the instructions for the Gradual case, subjects read on the screen at the beginning of the experiment. The instructions for the other two treatments are available upon request from the authors.

INSTRUCTIONS PAGE 1

Welcome to this experiment on decision-making. Please read the following instructions carefully. When everyone has finished reading the instructions and before the experiment starts, you will receive a handout with a summary of the instructions. During the experiment you will be asked to make a number of decisions. Your decisions and the decisions of other participants will determine how much money you earn. At the start of the experiment you will receive a starting capital of 500 points. In addition you will earn money with your decisions. The experiment consists of 50 rounds. In each round, your earnings will be denoted in points. Your earnings in the experiment will be equal to the sum of the starting capital and your earnings in the 50 rounds. At the end of the experiment, your earnings in points will be transferred into money. For each 100 points you earn, you will receive 1 euro. Your earnings will be privately paid to you in cash.

In each of the 50 rounds of the experiment all participants are coupled in pairs. In each round you are randomly assigned to a new partner.

In each round, you and your partner will play a game in which you can choose between two options (**Cooperate** and **Defect**). Your earnings for this game will be determined by your and your partner's choice. After observing the choices in

the game, the players choose whether or not to punish their partner. However, before the game each player decides how much punishment he or she can get at maximum. This number is referred to as ‘the possible punishment level’. Both players are informed of the possible punishment levels in the own pair before they play the game.

INSTRUCTIONS PAGE 2

SEQUENCE OF EVENTS IN A ROUND

In the first phase of each round, you can determine how high punishment you might get later. You choose a possible punishment level that is at least 0 and at most 50. In the experiment, a computerized clock raises the punishment level step-by-step starting from zero. If the clock shows the punishment level you would like to choose, you have to press the SUBMIT button, and this will be your possible punishment level later on in this round. At the same time your partner has the same procedure to decide about his or her possible punishment level. If either player presses the SUBMIT button, the other player is automatically and immediately informed about it. However, this phase does not end here, the other player has to press the SUBMIT button too. If a player does not press the SUBMIT button before the clock reaches the maximum level of 50, then the player’s possible punishment level will be 50 in this round.

Notice that at the start of the second phase you and your partner are informed about each other’s possible punishment levels. Then you will play a game where you and your partner have to choose between Cooperation and Defection. When you make your decision, you will not know your partner’s decision (and neither does your partner know your decision). The earnings from this game will be determined as follows (where the first number in a cell determines your payoff, and the second your partner’s payoff):

		Decision partner	
		C	D
Own decision	C	55,55	5,70
	D	70,5	25,25

In words, this means the following:

If both of you choose C, you and your partner will each earn 55 points.

If you choose C and your partner chooses D, you will earn 5 points and your partner earns 70 points.

If you choose D and your partner chooses C, you will earn 70 points and your partner earns 5 points.

If both of you choose D, you and your partner will each earn 25 points.

After both of you have decided, your choices will be revealed. In the third phase, you can decide whether you want to punish your partner. The punishment costs you 4 points. If you decide to punish, your partner gets a point deduction of the amount he or she determined him or herself in the first phase. If you decide

not to punish, your partner does not get a point deduction, and you do not have to bear the punishment costs. At the same time, your partner also decides whether he or she wants to punish you. If he or she decides to punish you, he or she also bears a punishment cost of 4 points, and you get a point deduction of the amount you determined in the first phase. If he or she decides not to punish, you do not get a point deduction, and your partner does not have to bear the punishment costs.

After that the round is finished.

INSTRUCTIONS PAGE 3

ROUND EARNINGS

In each round, you can gain or lose points. Your earnings will consist of your payoff from the game minus the punishment you get from the other player minus the punishment costs if you punish your partner.

Note that your decision in the first phase (when you determine your punishment level) does not have a direct and immediate effect on your payoff. You get a point deduction only if your partner decides to punish you after observing your choice.

MATCHING PROCEDURE

In each round, you will be randomly matched to another participant. You will never learn with whom you are matched. The random matching scheme is chosen such that you will never be coupled to the same partner in two subsequent rounds.

INFORMATION

When you decide whether to cooperate or defect in the second phase, you know the possible punishment level your partner has chosen in the first phase. At the end of each round you will learn whether you received punishment or not. You will also be informed about the number of points you have earned in that round.

HISTORY OVERVIEW

The lower part of the screen provides an overview of the results of rounds already completed. If less than 10 rounds have been completed, this history overview contains results of all completed rounds. In case more than 10 rounds have already been completed, the history overview is restricted to the 10 most recent rounds.

The history overview contains the results of your group. The history overview contains your choices together with the choices of 7 other participants with whom you interact in the experiment. Sometimes your choices are listed as player 1's choices, sometimes as player 2's choices. Below you see an example of the history overview. The first column contains the possible punishment level of player 1 of a pair, whereas the second column shows the possible punishment level of player 2 of a pair. These punishment levels were determined in the first phase. The third column shows whether player 1 chose to cooperate or not. Column 4 shows whether player 2 chose to cooperate or not. Column 5

indicates whether player 1 received a punishment or not (yes: punishment, no: no punishment). Column 6 indicates whether player 2 received a punishment or not (yes: punishment, no: no punishment).

Player 1's punishment level	Player 2's punishment level	Player 1's decision	Player 2's decision	Was player 1 punished	Was player 2 punished
0	0	D	D	no	no
5	0	D	D	no	no
5	10	C	C	no	no
10	1	C	D	no	no
18	21	D	D	no	yes
20	10	C	D	no	yes
20	15	C	D	no	yes
25	30	C	C	no	no

The past observations in the history screen have been ordered first by the punishment level of player 1. If these numbers are the same across pairs, these observations have been ordered by player 2's punishment level. If these numbers are also the same, then they are sorted on cooperation.

On the next screen you will be requested to answer some control questions. Please answer these questions now.

References

- Altman, I. (1973). Reciprocity of Interpersonal Exchange. *Journal for the Theory of Social Behaviour*, 3(2):249–261.
- Andreoni, J. and Samuelson, L. (2006). Building rational cooperation. *Journal of Economic Theory*, 127(1):117–154.
- Andreoni, J. and Varian, H. (1999). Preplay contracting in the Prisoners’ Dilemma. *Proceedings of the National Academy of Sciences*, 96(19):10933–10938.
- Ben-Porath, E. and Dekel, E. (1992). Signaling Future Actions and the Potential for Sacrifice. *Journal of Economic Theory*, 57(1):36–51.
- Bolton, G. and Ockenfels, A. (2000). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review*, 90(1):166–193.
- Charness, G., Fréchet, G., and Qin, C.-Z. (2007). Endogenous transfers in the Prisoner’s Dilemma game: An experimental test of cooperation and coordination. *Games and Economic Behavior*, 60(2):287–306.
- Chen, Y. and Gazzale, R. (2004). When Does Learning in Games Generate Convergence to Nash Equilibria? The Role of Supermodularity in an Experimental Setting. *American Economic Review*, 94(5):1505–1535.
- Coase, R. (1960). The Problem of Social Cost. *Journal of Law and Economics*, 3:1–44.
- Derlega, V., Wilson, M., and Chaikin, A. (1976). Friendship and Disclosure Reciprocity. *Journal of Personality and Social Psychology*, 34(4):578–582.
- Dindia, K. and Allen, M. (1992). Sex differences in self-disclosure: A meta-analysis. *Psychological Bulletin*, 112(1):106–124.
- Falkinger, J. (1996). Efficient private provision of public goods by rewarding deviations from average. *Journal of Public Economics*, 62(3):413–422.
- Falkinger, J., Fehr, E., Gächter, S., and Winter-Ebmer, R. (2000). A Simple Mechanism for the Efficient Provision of Public Goods: Experimental Evidence. *American Economic Review*, 90(1):247–264.
- Fehr, E. and Gächter, S. (2000). Cooperation and Punishment in Public Goods Experiments. *American Economic Review*, 90(4):980–994.
- Fehr, E. and Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415(6868):137–140.
- Fehr, E. and Schmidt, K. (1999). A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics*, 114(3):817–868.

- Gambetta, D. (2009). *Codes of the Underworld: How Criminals Communicate*. Princeton University Press.
- Gürek, Ö., Irlenbusch, B., and Rockenbach, B. (2006). The Competitive Advantage of Sanctioning Institutions. *Science*, 312(5770):108–111.
- Gürek, Ö., Irlenbusch, B., and Rockenbach, B. (2009). Voting with Feet: Community Choice in Social Dilemmas. *IZA Discussion Paper*, No. 4643.
- Hamaguchi, Y., Mitani, S., and Saijo, T. (2003). Does the Varian Mechanism Work? - Emissions Trading as an Example. *International Journal of Business and Economics*, 2(2):85–96.
- Herrmann, I. and Palmieri, D. (2005). A haunting figure: The hostage through the ages. *International Review of the Red Cross*, 87(857):135–145.
- Iossa, E. and Spagnolo, G. (2011). Contracts as Threats: on a Rationale For Rewarding A while Hoping For B. *CEPR Discussion Papers*, DP8195.
- Kosfeld, M., Okada, A., and Riedl, A. (2009). Institution Formation in Public Goods Games. *American Economic Review*, 99(4):1335–1355.
- Laurenceau, J.-P., Barrett, L., and Pietromonaco, P. (1998). Intimacy as an Interpersonal Process: The Importance of Self-Disclosure, Partner Disclosure, and Perceived Partner Responsiveness in Interpersonal Exchanges. *Journal of Personality and Social Psychology*, 74(5):1238–1251.
- Lee, A. (1991). The Role of Hostages in Roman Diplomacy with Sasanian Persia. *Historia: Zeitschrift fr Alte Geschichte*, 40(3):366–374.
- Levine, D. (1998). Modeling Altruism and Spitefulness in Experiments. *Review of Economic Dynamics*, 1(3):593–622.
- Ostrom, E., Walker, J., and Gardner, R. (1992). Covenants With and Without a Sword: Self-Governance is Possible. *American Political Science Review*, 86(2):404–417.
- Plutarch (1992). *Essays*. London: Penguin.
- Rabin, M. (1993). Incorporating Fairness into Game Theory and Economics. *American Economic Review*, 83(5):1281–1302.
- Raub, W. (2009). Commitments by Hostage Posting. In Baurmann, M. and Lahno, B., editors, *Perspectives in Moral Science*, pages 207–225. Frankfurt am Main: Frankfurt School Verlag.
- Raub, W. and Keren, G. (1993). Hostages as a commitment device: A game-theoretic model and an empirical test of some scenarios. *Journal of Economic Behavior and Organization*, 21(1):43–67.

- Rotenberg, K. (1986). Same-Sex Patterns and Sex Differences in The Trust-Value Basis of Children's Friendship. *Sex Roles*, 15(11):613–626.
- Schelling, T. (1960). *The strategy of conflict*. Harvard University Press, Cambridge, Massachusetts.
- Sefton, M., Shupp, R., and Walker, J. (2007). The effect of rewards and sanctions in provision of public goods. *Economic Inquiry*, 45(4):671–690.
- Snijders, C. and Buskens, V. (2001). How to convince someone that you can be trusted? The role of 'hostages'. *Journal of Mathematical Sociology*, 25(4):355–383.
- Sommerfeld, R., Krambeck, H.-J., Semmann, D., and Milinski, M. (2007). Gossip as an alternative for direct observation in games of indirect reciprocity. *Proceedings of the National Academy of Sciences*, 104(44):17435–17440.
- Sutter, M., Haigner, S., and Kocher, M. (2010). Choosing the Carrot or the Stick? Endogenous Institutional Choice in Social Dilemma Situations. *Review of Economic Studies*, 77(4):1540–1566.
- Varian, H. (1994). A Solution to the Problem of Externalities When Agents Are Well-Informed. *American Economic Review*, 84(5):1278–1293.
- Weber, R. (2006). Managing Growth to Achieve Efficient Coordination in Large Groups. *American Economic Review*, 96(1):114–126.
- Williamson, O. (1983). Credible Commitments: Using Hostages to Support Exchange. *American Economic Review*, 73(4):519–540.
- Yamagishi, T. (1986). The provision of a sanctioning system as a public good. *Journal of Personality and Social Psychology*, 51(1):110–116.