Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

FAKULTÄT
FÜR WIRTSCHAFTS- UND
SOZIALWISSENSCHAFTEN

Adversity is a school of wisdom: Experimental evidence on cooperative protection against stochastic losses

Sonja Köke
Andreas Lange
Andreas Nicklisch

WiSo-HH Working Paper Series
Working Paper No. 22
February 2015

WiSo-HH Working Paper Series
Working Paper No. 22
February 2015

Adversity is a school of wisdom: Experimental evidence on cooperative protection against stochastic losses

Sonja Köke, Universität Hamburg
Andreas Lange, Universität Hamburg
Andreas Nicklisch, Universität Hamburg

ISSN 2196-8128

Font used: „TheSans UHH" / LucasFonts

Kontakt:

WiSo-Forschungslabor
Von-Melle-Park 5
20146 Hamburg
E-Mail: experiments@wiso.uni-hamburg.de
Web: http://www.wiso.uni-hamburg.de/forschung/forschungslabor/home/

Adversity is a school of wisdom: Experimental evidence on cooperative protection against stochastic losses*

Sonja Köke[§] Andreas Lange[‡] Andreas Nicklisch[¶]

February 18, 2015

Abstract

We investigate the dynamics of voluntary cooperation to either reduce the size or the probability of a stochastic shock. For variants of a repeated four-person prisoner's dilemma game, we show that cooperation is larger and more stable when it affects the probability rather than the size of damages. We provide crucial insights on behavioral adaptation following adverse events: defecting players are more likely to switch to cooperation after experiencing an adverse event, while existing cooperation is reinforced when the damage does not occur. This behavior is consistent with simple learning dynamics based on ex post evaluations of the chosen strategy.

Keywords: ex post rationality, experiment, cooperation, repeated prisoner's dilemma, regret learning, stochastic damages
JEL codes: C92, H41, Q54

*The authors thank participants of several conferences and seminars at several universities for useful comments.

[§]Department of Economics, University of Hamburg, e-mail: sonja.koeke@wiso.uni-hamburg.de

[‡]*Corresponding author:* Department of Economics, University of Hamburg, von-Melle-Park 5, 20146 Hamburg, Germany, e-mail: andreas.lange@wiso.uni-hamburg.de

[¶]Department of Economics, University of Hamburg, e-mail: andreas.nicklisch@wiso.uni-hamburg.de

1 Introduction

Protection against common stochastic losses is an ubiquitous challenge for modern societies. The fight against extreme events triggered by climate change, hurricane prevention, protecting public security against terror, international health cooperation against pandemic diseases, or simply forest fire prevention are important contemporary examples. Given the stochastic nature of damage events, several scholars have studied the behavioral responses of individuals and organizations to damage occurrences. Birkland (2006) investigates determinants of policy changes in diverse policy fields from aviation security and terrorism to preparedness to earthquakes or hurricanes. He interprets accidents, naturally occurring disasters or deliberately caused catastrophes as "focussing events" (see Kingdon, 1995) which may induce increased attention to a policy problem and thereby possibly trigger policy changes.

When considering stochastic damages and behavioral reactions that they trigger, two qualitatively different channels exist through which actions may impact future damage events: they may impact the *size* of damages while leaving the probability of an adverse event unaffected (e.g., preparing for earthquakes, adaptation for climate change) or they may change the *probability* that such an adverse event occurs and thereby may fully prevent the damage event from happening (e.g., aviation security, terrorism, mitigation of climate change). In this paper we investigate how the individual willingness to cooperate depends on how preventive actions affect stochastic losses. We concentrate on a voluntary cooperation setting as protective actions against probabilistic losses often require the cooperation of members of a community. Similar incentive structures can be found in many different settings, from team production settings, where possible payoffs depend on the group performance, to problems of non-point source pollution, where fines can only be triggered based on ambient, rather than individual pollution levels. We are particularly interested in the evolution of behavior over time, that is, how experiencing adverse events affects subsequent decisions. Although problems of repeated cooperation to reduce probabilistic losses are common place, there is surprisingly little known about how people actually behave when facing this type of challenges.

For this purpose, we provide experimental evidence within variants of a repeated n-person prisoner dilemma game with stochastic payoffs: people may (indefinitely) repeatedly choose to invest in protective actions which benefit the entire group. In the short run (one-shot), players have incentives to free-ride on the investments of others, while the (indefinitely) repeated interaction will allow for positive cooperation levels sustained in subgame-perfect equilibria. Specifically, we compare a setting where individual cooperation reduces the damage size in cases of certain damages (*CertDam*) with stochastic settings in which cooperation affects expected damages through reducing either the damage size (*DamRed*) or the probability of damages of fixed size (*ProbRed*). Expected payoffs conditional on the number of cooperators in the group are held constant across treatments. We compare individual behavior between treatments and the determinants of its evolution over time.

Our experimental results show significant differences between cooperation rates in *CertDam* and *DamRed* versus *ProbRed*: subjects are more likely to cooperate to reduce the probability of the all-or-nothing damage, rather than to marginally reduce the size of a certain or stochastic damage. These differences between treatments get more pronounced over time. When cooperation reduces the probability of an adverse event, cooperation remains rather stable over the series of interactions. In sharp contrast, cooperation rates decline over time when cooperation reduces the size of a certain damage or a stochastic damage which occurs with fixed probability.

In line with our motivating examples and the (German) proverb which inspired the title of our paper, we demonstrate that experiencing adverse events in treatments with stochastic damages is of particular importance for the dynamics of individual behavior: (i) non-cooperating players are more likely to switch to cooperation following a damage event. This tendency is particularly strong in *ProbRed* and is consistent with ex post regret if the player believes to have been pivotal in triggering the damage. (ii) The occurrence of damages makes it less likely for cooperating players to continue cooperation. In other words, the absence of the damage reinforces existing individual cooperation. Players therefore appear to assess their actions from an ex post perspective when deciding about future actions. As such, we demonstrate that our findings on cooperation rates and their dynamics deviate from conventional game theoretic equilibrium concepts. Rather, the treatment differences and the dynamics of decisions are largely consistent with combinations of behavioral motives of anticipated regret (e.g., Loomes and Sugden, 1988; Zehlenberg, 1999; Filiz-Ozbay and Ozbay, 2007) and simple learning dynamics which also link back to notions of ex post regret (e.g., Selten and Chmura, 2008; Chmura et al., 2012).

The remainder of the paper is structured as follows: section 2 gives an overview of the related literature. Section 3 describes the experimental setting: after describing the game in section 3.1, we derive predictions in section 3.2, before detailing the experimental design in section 3.3. Our experimental results are presented and discussed in section 4. Section 5 concludes.

2 Related Literature

Our experiment relates to several different strands of theoretical and experimental literature. First, it is related to the extensive literature on "self-insurance" and "self-protection." Following the seminal article by Ehrlich and Becker (1972), the function of protective and preventive actions as complements or substitutes for market insurance are analyzed at the individual level for purely private goods (Dionne and Eeckehoudt, 1985; Jullien et al., 1999; Briys and Schlesinger, 1990)[1] or related to some forms of externalities (Muermann and Kunreuther, 2008). Lohse et al. (2012) extend this literature by including a public good structure of the risky event: while they separately consider actions that reduce the size or the probability of a loss for all members in the group, they do not provide predictions on how behavior in the two cases compares.[2] Focussing completely on loss prevention, Keser and Montmarquette (2008) analyze individual contributions that reduce the risk of correlated public losses. They compare risky decisions with decisions under ambiguity varying initial probabilities and initial endowments. They show that contributions decrease in initial loss probability and with ambiguity (in comparison to risk), while they increase with endowment. Likewise, Dickinson (1998) compares public good games with probabilistic and certain gains from contributions and finds that risk decreases contributions. None of these paper provides a comparison of protective and preventive behavior in group settings.

Our study also relates to studies on the influence of group liability on individual behavior as we analyze probabilistic losses. For example, the environmental economics literature studies policy instruments for dealing with non-point source pollution (e.g., Segerson, 1988; Miceli and Segerson, 2007; Barrett, 2011), i.e. where fines can only be put on ambient pollution levels. Similarly, incentives for cooperative behavior in groups have been discussed in the context of indus-

[1]For a setting of a single decision maker, Friesen (2012) shows by building on Becker's (1968) theory of crime that risk averse participants are deterred more by an increase in fine than by an increase in the probability of being caught which leads to an identical expected fine. When translating the model to our setting, one would expect that cooperation is highest in the damage size reduction setting and lower in the probability reduction, exactly the opposite of our findings.

[2]While most of the papers use independent risks (uncorrelated realization of the loss), which makes sense when assuming an insurance market in the private good case and represents examples like public security and cancer research in the public good setting, Muermann and Kunreuther (2008) have started to analyze partly correlated risks. In our setting, we are interested in fully correlated risks, which do rather capture the case of natural catastrophes, epidemics and wars. The impact of risk aversion on the efficient level as compared to the individual choice of self-insurance, self-protection and market insurance has been analyzed by Dionne and Eeckehoudt (1985), Jullien et al. (1999) and Briys and Schlesinger (1990). The choice of insurance types when both self-insurance and self-protection are available but no market insurance have been analyzed by Ihori and McGuire (2010).

trial organization and team production (e.g., Holmstrom, 1982; Rasmussen, 1987; Varian, 1990). This mostly theoretical literature considers typically the threat of group penalties to prevent shirking of group members in one-shot rather than repeated settings, whereas participants in our setting may choose to cooperate to avoid being potentially penalized by increased free-riding of other group members in the consecutive periods. Also inspired by environmental problems is recent experimental research on threshold public good games (e.g., Milinski et al., 2008; Tavoni et al., 2011). These papers are related to our investigation as in our *ProbRed* treatment, damages are avoided if an ex ante unknown threshold of cooperating players is reached. Dannenberg et al. (2014) consider settings with commonly known horizons but unknown thresholds which differ from our study as we consider indefinitely repeated games in which cooperation could be sustained as an equilibrium.

With its repeated game structure with random termination rule, our experiment also is related to Bò and Fréchette (2011) who study long sequences of prisoners' dilemma games with probabilistic termination as well. In a setting with positive deterministic payoffs, they provide evidence suggesting that the existence of a cooperative equilibrium may be a necessary (but not sufficient) condition for persistent cooperation or even cooperation levels which increase with experience. Rather, their results demonstrate that individual learning paths matter crucially for the occurrence of joint cooperation. For our analysis of individual decision dynamics, we apply learning dynamics following ex post rationality (e.g., Selten and Chmura, 2008; Chmura et al., 2012; Selten and Stoecker, 1986; Roth and Erev, 1995; Camerer and Ho, 1999; Beggs, 2005). While all these models can explain that cooperation is declining for *CertDam* and *DamRed* while being more stable in *ProbRed*, decision dynamics at the individual level can be best explained by models where the ex post optimal individual action receives a positive reinforcement (e.g., Selten and Chmura, 2008; Chmura et al., 2012).[3] As such, our article also complements Bereby-Meyer and Roth (2006) who relate individual learning patterns for stochastic or certain payoffs to the frequency of payoffs: their results suggest that subjects learn faster to defect in a two-player prisoners' dilemma game when payoffs are certain rather than probabilistic.[4] Based on their finding that the frequency of the payoff is more important than the payoff size for learning to defect, we can expect cooperation to deteriorate faster under certain losses than when losses are stochastic.

[3] An earlier experimental study by Aoyagi and Fréchette (2009) finds for sequences of prisoners' dilemma games with probabilistic termination and positive deterministic payoffs that the best fitting model analyzing individual choices simply relies on most recent experience.

[4] Unlike our setting, Bereby-Meyer and Roth consider in their stochastic treatment only cooperation influencing the probability but not the payoffs: cooperation impacts the probability of gaining 1 token instead of 0 tokens in the range of 0.5% to 17.5%.

3 Game and Predictions

3.1 Experimental Treatments

The starting point of our setting is a repeatedly played simultaneous move four-person prisoners' dilemma ($n = 4$). At the beginning of each period, each player is endowed with E tokens. At the end of each period, a damage of D tokens occurs with probability p and reduces the endowment of each player. Damages are fully correlated across the four players; that is, either all players or no player within a group incur the damage in a given period and damages are independent over time[5]. With their decisions, players may reduce either the size or the probability of the damage, depending on the treatment.

For this purpose, each player is asked before the damage realizes, whether she wants to cooperate or defect.[6] The action of individual i in period t is, therefore, the binary contribution choice $q_i^t \in \{0, 1\}$ with $q_i^t = 1$ being cooperation and $q_i^t = 0$ defection. Cooperation costs the individual player c tokens. The sum of cooperators in a group and period is denoted by $Q^t = \sum_{j=1}^{n} q_j^t = q_i^t + Q_{-i}^t$. The potential damage, $D^{Treat}(Q^t)$, and the probability of its occurrence, $p^{Treat}(Q^t)$, depend on the total cooperation level and differ between treatments ($Treat$). With this, the general payoff structure of individual i in period t for a certain treatment condition is given by

$$\pi_{i,t}(q_i^t, Q_{-i}^t, s^t) = E - cq_i^t - s^t D^{Treat}(Q^t) \tag{1}$$

where $s^t \in \{0, 1\}$ reflects the state of nature where the damage has ($s^t = 1$) or has not ($s^t = 0$) occurred.

In the experiment, we differentiate between three treatments which are calibrated to guarantee equivalence in expected damages, that is, $p^{Treat}(Q^t)D^{Treat}(Q^t)$ is equivalent for all treatments. In the first treatment, denoted as *CertDam*, expected damages occur with certainty: $D^{CertDam}(Q^t) = p_0 D_0 - p_0 dQ^t$ and $p^{CertDam}(Q^t) \equiv 1$. In the second treatment, hereafter denoted as *DamRed*, each player's cooperation leads to a reduction of the damage by d, while the initial probability is kept constant, that is, we have $D^{DamRed}(Q^t) = D_0 - dQ^t$ and $p^{DamRed}(Q^t) \equiv p_0$. In the third treatment, hereafter denoted as *ProbRed*, cooperation leads to a reduction of the probability of the damage ($p^{ProbRed}(Q^t) = p_0 - xQ^t$) while its level is fixed at $D^{ProbRed}(Q^t) \equiv D_0$. Equivalence of the expected payoffs is guaranteed by setting $dp_0 = xD_0$ which leads to expected damages in all treatments being given by $p_0(D_0 - dQ) = (p_0 - xQ)D_0$.

[5]For simplicity reasons in the experiment, we do not introduce the structure of a stock pollutant in this paper.

[6]In the experiment, we use neutral wording; the exact wording is "take/not take an action".

In order to guarantee the prisoners' dilemma structure, we assume $np_0d > c > p_0d$ and $nxD_0 > c > xD_0$. In other words, cooperation is socially beneficial in terms of expected payoffs, but does not pay off individually. Further, we assume that even full cooperation ($Q^t = n$) does not reduce the damage nor its probability to zero ($p_0 - nx > 0$, $D_0 - nd > 0$).

In our experiment, players get information about their own cooperation decision q_i^t, the resulting cost they incurred, and the total level of cooperation Q^t after each period. They also get to know whether the damage event occurred or not and are informed about their individual payoff. With this information, players in *CertDam* and *DamRed* can calculate the payoff that they would have received if they had changed their own decision. This is different in *ProbRed*: for example after observing a damage event, a defecting player cannot know if the damage also would have occurred if she individually had cooperated. Conversely when no damage occurred, a cooperating player does not know if she was pivotal in preventing the damage event. In order to control for the impact of players' being informed about their marginal impact on the payoff, we introduce a fourth treatment condition *ProbRed*$^+$ which is identical with *ProbRed* in the mapping of cooperation into probability and damage, but gives players additional feedback after each period: players are informed whether the damage would have occurred if zero, one, two, three, or four players had cooperated. Therefore, *ProbRed*$^+$ increases the subjects' awareness about their decision's marginal impact on the payoff.

Table 1 gives an overview of the damage and probability functions as well as the resulting expected damages for all treatments.

Treatment	$D^{Treat}(Q^t)$	$p^{Treat}(Q^t)$	$D^{Treat}(Q^t)\ p^{Treat}(Q^t)$
CertDam	$p_0D_0 - p_0dQ^t$	1	$p_0(D_0 - dQ^t)$
DamRed	$D_0 - dQ^t$	p_0	$p_0(D_0 - dQ^t)$
ProbRed	D_0	$p_0 - xQ^t$	$(p_0 - xQ^t)D_0$
ProbRed$^+$	D_0	$p_0 - xQ^t$	$(p_0 - xQ^t)D_0$

Table 1: *Summary of damage size $D^{Treat}(Q^t)$ and damage probability $p^{Treat}(Q^t)$ for the respective treatments ProbRed, ProbRed$^+$, DamRed, and CertDam.*

In all treatment conditions our setting mimics infinite play. For this purpose, we apply the random stopping rule for supergames (e.g., Bò and Fréchette, 2011). In our experiment, the number of supergames is not known to the players. At the beginning of each supergame, players are randomly re-matched into new groups. Each supergame consists of several periods of the game described above. A supergame has a publicly known termination probability δ after each period. That is, after each period, the supergame terminates with probability δ, and a new supergame starts in new randomly re-matched groups, whereas with probability $1-\delta$ the supergame continues in the same group constellation. Therefore, the setting allows players to learn in changing group compositions across supergames. At the same time, they cannot predict the termination of the specific supergame, while they have to discount future payoffs from cooperation with a common "discount factor".

We set the parameter values as follows: termination probability $\delta = 0.2$, initial damage probability $p_0 = 0.5$, probability reduction $x = 0.1$, initial damage size $D_0 = 20$, damage reduction $d = 4$, initial endowment $E = 25$ and cost $c = 5$.

3.2 Predictions

We are interested in the level of cooperation as well as in dynamics of behavior across periods. To derive predictions about cooperation behavior, we use two conceptually different tools. In a first step, we use the standard equilibrium concept of subgame perfection, which is derived under the assumption of perfectly rational and forward looking individuals who define strategies that condition the action in each period on group members' behavior in the previous period. We derive the number of cooperators that can be supported in equilibrium for risk-neutral and risk-averse individuals as well as for individuals who show ex-ante regret aversion (e.g., Loomes and Sugden, 1982). In the second step, we consider theoretical approaches of ex post rationality, in particular impulse balance learning (e.g., Selten and Chmura, 2008; Chmura et al., 2012).[7] These approaches assume that players ex post assess the success of their previously chosen action and adapt their strategy accordingly. While departing from standard equilibrium concepts (which may lead to predictions on the extent of cooperation), these approaches allow to make predictions on the dynamics of individual decisions.

3.2.1 Subgame-Perfect Nash Equilibria (expected utility maximizers)

Since all treatments are identical in the mapping of cooperation decisions into expected payoffs, the equilibria for risk neutral players do not differ between treatments. It is obvious that the game has a subgame perfect equilibrium in which *all* players *always* defect: as in the one-shot prisoner's dilemma game, no player individually has an incentive to cooperate. The repeated nature of the game allows, however, for additional subgame perfect equilibria so that there is no unique equilibrium in our game. To demonstrate the possible extent of cooperation in a subgame perfect equilibrium, we assume that a set of $Q \leq n$

[7]We also relate to concepts of reinforcement learning (e.g., Roth and Erev, 1995; Erev and Roth, 1998) and experience-weighted attraction learning (e.g., Camerer et al., 1999; Ho et al., 2008).

players follow a modified grim trigger strategy: they cooperate as long as at least $Q-1$ other players cooperate, otherwise they defect in all subsequent periods. The remaining $n-Q$ players always defect. This strategy involves the maximally possible punishment for a once-deviating player and can sustain cooperation as a subgame perfect equilibrium if

$$\begin{aligned}
& \sum_{t=0}^{\infty}(1-\delta)^t\left[E-c-p_0(D_0-dQ)\right] \\
& \geq \left[E-p_0(D_0-d(Q-1))\right]+\sum_{t=1}^{\infty}(1-\delta)^t(E-p_0D_0) \\
\Leftrightarrow \quad & \frac{1}{\delta}\left[E-c-p_0(D_0-dQ)\right] \\
& \geq \left[E-p_0(D_0-d(Q-1))\right]+\frac{1-\delta}{\delta}(E-p_0D_0) \qquad (2)
\end{aligned}$$

Here, the left-hand side states the expected payoff of any cooperating player i if all Q players continue to cooperate forever. The first expression of the right hand side states the payoff of a deviator i in the period in which he deviates, while the second term states the expected continuation payoff if all players play defect, starting in the next period. Note that given the defection of other players, the deviating player does not have an incentive to return to cooperation. Therefore, if condition (2) is satisfied, Q players playing the modified grim trigger strategy and $n-Q$ players always defecting establishes a subgame perfect equilibrium.

Condition (2) can be rewritten as

$$Q \geq Q^{min} = \frac{c}{p_0d(1-\delta)}+\frac{\delta}{1-\delta} \qquad (3)$$

which – given our parameter choice – leads to equilibria of $Q \geq 3$ cooperating players which are supported by the modified grim trigger strategy ($Q^{min} = 2.875$).[8]

While this analysis suggests that zero or at least three (risk-neutral) players cooperate, it does not generate any predictions on treatment differences as those are identical in expected payoffs. Differences may occur if subjects are risk-averse (or risk-loving). Intuitively, one may expect levels of cooperation to be higher

[8]Naturally, the multiplicity of equilibria may motivate further discussions on equilibrium selection. While not being the focus of the paper, we note that the equilibrium which supports $Q = 4$ is not "renegotiation proof" as – following the defection of one player – the remaining three players collectively would not have an incentive to follow through with the punishment as it lowers their payoffs, while the cooperation of these three players can still be supported by the modified grim trigger strategies.

in *DamRed* than in *ProbRed* for risk-averse subjects: while the expected utility of a player for $Q = 0$ is identical in the two treatments ($p = p_0$, $D = D_0$), it is larger in *DamRed* than in *ProbRed* and *ProbRed*$^+$ if $Q > 0$, suggesting that the willingness to cooperate is higher in *DamRed* than in *ProbRed* and *ProbRed*$^+$.[9] However, this does not take into account the repeated nature of the game. We assume that risk preferences apply to the total payoff of the individual in the experiment ($\mathbb{E}\big(u(\sum_t \pi_i^t)\big)$). We concentrate on constant absolute risk aversion (CARA, $u(x) = -\frac{1}{\sigma}\exp(-\sigma x)$) which has two properties that allow for an analytic solution of the equilibrium conditions: (i) optimal strategies for players do not depend on the wealth already accumulated in the experiment; (ii) constant absolute risk aversion preferences lead to behavior which is not affected by (independent) background risk (Gollier and Pratt, 1996) such that equilibrium conditions do not depend on the number of subsequently played supergames.

In order to study the stability of cooperation under risk aversion, we again concentrate on modified grim trigger strategies that have been introduced above. Under the assumption of CARA preferences, closed form solution can be derived for value functions which capture the expected utility to the player from the continuation of the game within the same supergame (see Appendix A). For the parameters used in our experiment, Figure 1 depicts the minimal cooperation level Q^{min} needed in the respective treatments to make cooperation attractive for a subject of a given level of risk aversion σ. We see that all of the curves collapse for risk-neutral players ($\sigma = 0$) for which we again obtain the level given in equation (3) of $Q^{min} = 2.875$. For all treatments, the threshold Q^{min} is increasing in σ such that more risk-averse subjects are less likely to cooperate, controlling for the cooperation level of other subject in the same group.

Already for small levels of risk-aversion ($\sigma > 0.02$)[10] cooperation cannot be stabilized in any of the treatments. For levels $0 < \sigma < 0.02$, we see that in order to make an individual cooperate, more cooperating players are needed in *CertDam* than in *ProbRed* and *ProbRed*$^+$ than in *DamRed*. For risk-loving players ($\sigma < 0$) the relationship is reversed. Inverting the depicted functions, we see that cooperation is harder to sustain in equilibrium if agents are more risk-averse. That is, the more risk-averse a player is, the less likely she belongs to a potentially existing subset of cooperating players. This decreasing effect of risk-aversion on the propensity to cooperate is intuitive as more risk-averse subjects implicitly put more weight on situations where small payoffs occur. As this is the case when the game stops, they put less weight on the continuation of the supergame such that defection becomes more attractive as this realizes maximal payoffs in the current

[9] Note that $(p_0 - xQ)u(E - D_0 - cq_i) + (1 - p_0 + xQ)u(E - cq_i) \geq p_0 u(E - D_0 + dQ - cq_i) + (1 - p_0)u(E - cq_i)$ due to the concavity of $u(\cdot)$ such that (collective) cooperation is more beneficial. This argument follows an analysis of individual decision making by Friesen (2012).

[10] Note that this level of risk-aversion for experimental Taler translates into a CARA-level of 2 in terms of Euro.

period.

Prediction 1. ***(SPNE)*** *(a) Risk-averse players are predicted to be less likely to cooperate in all treatments. (b) Sustained cooperation is most likely in DamRed, less so in ProbRed and ProbRed$^+$, and least in CertDam for risk aversion, while the opposite order holds for risk-loving players.*

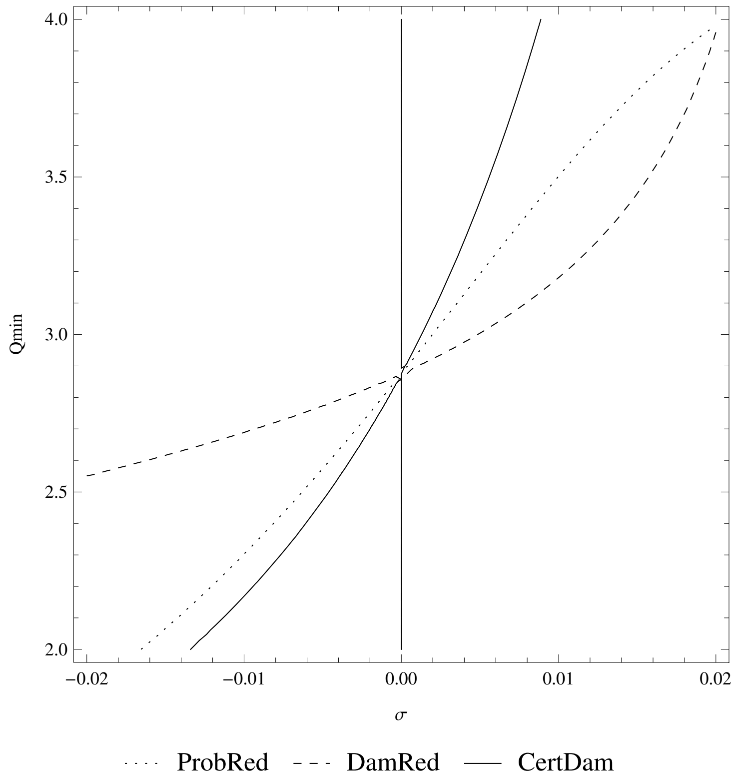

Figure 1: *Minimal cooperation level Q^{min} required to stabilize cooperation as a function of risk aversion σ. The parameter values are those given at the end of section 3.1.*

3.2.2 Subgame-Perfect Nash Equilibria (anticipated regret)

We now complement the derivations of subgame-perfect Nash equilibria for standard expected utility maximizers by considering players who suffer from anticipated regret (e.g., Loomes and Sugden, 1982). After having chosen q_i^t and after the realization of payoff $\pi_{i,t}(q_i^t, Q^t, s^t)$ where $s^t = 1$ if damages have and $s^t = 0$ if they have not occurred, the player may experience some regret. The ex post utility is assumed to be given by

$$\begin{aligned} u_i(q_i^t, Q_{-i}^t, s^t) &= \pi_{i,t}(q_i^t, Q_{-i}^t, s^t) \\ &\quad -R_i(\max\{0, \pi_{i,t}(1-q_i^t, Q_{-i}^t, \hat{s}^t) - \pi_{i,t}(q_i^t, Q_{-i}^t, s^t)\}) \end{aligned} \tag{4}$$

where $R_i(\cdot)$ is increasing and convex to capture regret aversion ($R_i(0) = 0$) and $\hat{s}$ reflects the damage event under the alternative action $(1-q_i^t)$. That is, the player may experience regret if the alternative action would have led to a higher payoff. In *CertDam* and *DamRed*, note that $\hat{s}^t = s^t$ and that defection is a dominant strategy: the payoff difference to cooperation is 3 in *CertDam*, and 1 or 5 in *DamRed*, depending on the damage event. In *ProbRed* and *ProbRed*$^+$, however, player i may be pivotal in triggering the damage event such that $\hat{s}^t$ can differ from s^t. This happens with 10% probability[11] and would lead to cooperation being the superior action (payoff difference $D_0 - c = 20 - 5 = 15$). With 90% probability, the player cannot affect the damage event in which case defection ex post would have been the better choice (payoff difference $c = 5$). Again assuming modified grim trigger strategies, equation (2) can therefore be modified to

$$\begin{aligned} &\frac{1}{\delta}\left[E - c - p_0(D_0 - dQ) - R_i(3)\right] \\ &\geq \left[E - p_0(D_0 - d(Q-1))\right] + \frac{1-\delta}{\delta}(E - p_0 D_0) \end{aligned} \tag{5}$$

in *CertDam*,

$$\begin{aligned} &\frac{1}{\delta}\left[E - c - p_0(D_0 - dQ) - 0.5(R_i(1) + R_i(5))\right] \\ &\geq \left[E - p_0(D_0 - d(Q-1))\right] + \frac{1-\delta}{\delta}(E - p_0 D_0) \end{aligned} \tag{6}$$

in *DamRed*

$$\begin{aligned} &\frac{1}{\delta}\left[E - c - p_0(D_0 - dQ) - 0.9R_i(5)\right] \\ \geq\ &\left[E - p_0(D_0 - d(Q-1)) - 0.1R_i(15)\right] \\ &+\frac{1-\delta}{\delta}(E - p_0 D_0 - 0.1R_i(15)) \end{aligned} \tag{7}$$

[11] Imagine the damage occurs when a random draw between 0 and 1 is smaller than $p(Q)$. The impact of one more individual cooperating on $p(Q)$ is -0.1. The random draw determining the damage occurrence lies in this impact range with probability of 0.1. So the probability of the individual being pivotal in preventing the damage is 10%.

in *ProbRed* and *ProbRed*$^+$. Due to the convexity of $R_i(\cdot)$, we would therefore expect players to be less likely to choose cooperation in *DamRed* than in *CertDam* ($R_i(3) \leq 0.5(R_i(1) + R_i(5))$). The relationship to *ProbRed* and *ProbRed*$^+$ cannot be signed unambiguously: if $R_i(15)$ is sufficiently large ($R_i(\cdot)$ sufficiently convex), *ProbRed* and *ProbRed*$^+$ can be expected to lead to the largest individual propensity to cooperate (if $0.9R_i(5) - 0.1R_i(15) \leq R_i(3)$). That is, if players anticipate sufficient regret from triggering the damage event in *ProbRed* and *ProbRed*$^+$, we could expect equilibria with positive cooperation levels to be more likely to occur in *ProbRed* and *ProbRed*$^+$ than in *CertDam* than in *DamRed*.

3.2.3 Learning Dynamics

The analysis of equilibria in the two previous sections relies on strategies of players which are chosen exclusively by considering *future* (expected) payoffs. That is, the past is only relevant to the extent that it gives players information about the strategies of other players such that they can condition their own actions, for example, on the number of cooperating players. Thereby, no effect of damage events on future actions could be explained, contrary to our introductory examples. In this section, we assume that players apply *ex post rationality* (cf. Selten and Stoecker, 1986): players assess the success of their previously chosen action ex post and adapt the strategy accordingly. There are several theoretical approaches following the general idea of the *Law of Effect* stating that actions that (would) have been successful in the past will be reinforced and dissatisfying actions will be weakened (see Herrnstein, 1970): reinforcement learning (e.g., Roth and Erev, 1995; Erev and Roth, 1998), experience-weighted attraction learning (e.g., Camerer et al., 1999; Ho et al., 2008), and impulse balance learning (e.g., Selten and Chmura, 2008; Chmura et al., 2012). In the following, we will concentrate on the latter as it is closest to the idea of regret used in the previous section. We illustrate how cooperation may evolve over time in the respective treatments.

Formally, there is an initial attraction $A_{i,0}(q)$ of player i to play action $q \in \{0, 1\}$. Selten and Chmura (2008) assume that the attraction of action q evolves according to

$$A_{i,t+1}(q) = A_{i,t}(q) + \max\{0, \pi_{i,t}(q, Q^t_{-i}, s^t(q)) - \pi_{i,t}(1-q, Q^t_{-i}, s^t(1-q)), \quad (8)$$

where $s^t(q)$ ($s^t(1-q)$) denotes the state of the damage event if action q $(1-q)$was chosen. That is, an action is reinforced if it would have been the better strategy. The probability of action q being played in period $t+1$ is simply its attraction relative to the sum of the attractions of both actions available to individual i:

$$P_{i,t+1}(q) = \frac{A_{i,t}(q)}{A_{i,t}(0) + A_{i,t}(1)}. \quad (9)$$

Note that the extent of reinforcement in (9) equals the obtained payoff difference. With the same arguments as in section 3.2.2, only defection is reinforced in *CertDam* (with 3 in each period t) and in *DamRed* (with 1 or 5, both with 50% probability in each period t) such that only $A_{i,t}(0)$ grows over time (in expectation by 3 per period) and we expect cooperation to be phased out in the long run:[12]

$$\mathbb{E}[P_{i,t+1}(1)] = \frac{A_{i,0}(1)}{A_{i,0}(0) + 3t + A_{i,0}(1)} \rightarrow_{t\rightarrow\infty} 0. \tag{10}$$

In *ProbRed*$^+$ and in *ProbRed* (if players behave according to the correct probability of having been pivotal), however, cooperation is reinforced in 10% of the periods (with a payoff difference of 15) while in the remaining 90% of the cases defection is reinforced (by 5). Therefore, in expectation $A_{i,t}(0)$ grows by 4 per period and $A_{i,t}(1)$ by 1.5, such that the expected probability of cooperation after t periods is

$$\mathbb{E}[P_{i,t+1}(1)] = \frac{A_{i,0}(1) + 1.5t}{A_{i,0}(0) + A_{i,0}(1) + 6t} \rightarrow_{t\rightarrow\infty} 0.25. \tag{11}$$

Therefore, we would not expect cooperation to be phased out in the long run. Instead the likelihood of cooperation converges towards 25%. Note that when a damage event occurs, cooperating players have (ex post) obviously chosen the wrong action such that defection will be reinforced. Conversely, their choice may have been correct if no damage occurred. Differently, defecting players may believe (in *ProbRed*) or know (in *ProbRed*$^+$) that they have been pivotal and therefore positively reinforce cooperation for the next period, while definitely the absence of a damage ex post confirms their choice. We therefore can formulate the following prediction.

Prediction 2. ***(Learning dynamics)*** *(a) For CertDam and DamRed we expect cooperation to be phased out in the long run, while cooperation can be sustained in ProbRed and ProbRed*$^+$*. The rate of decline of cooperation rates can be expected to be similar in CertDam and in DamRed, while being less pronounced in ProbRed and ProbRed*$^+$*. (b) In ProbRed and ProbRed*$^+$*, a damage event can be expected to make (i) defecting players more likely to switch to cooperation and (ii) cooperating players more likely to switch to defection.*

[12]The alternative learning dynamics (e.g., Erev and Roth, 1995) lead to similar insights for *CertDam* and *DamRed* as Beggs (2005) shows that weakly dominated strategies (as cooperation in our case) are phased out over time. In *ProbRed* and *ProbRed*$^+$, however, the likelihood of cooperation would also be predicted to converge to zero as defection is still dominant in expected payoff terms.

3.3 Experimental Design

In total, we ran 12 experimental sessions between January and March 2014 at the Experimental Laboratory of the School of Business, Economics and Social Sciences at the University of Hamburg. Three sessions were conducted for each of the treatment conditions that we described in section 3.1. A total of 280 students from the University of Hamburg participated in the experiment, with a maximum of 24 and a minimum of 16 subjects per session. Median age was 24 years, 53% were female participants.

We applied the same sequence of periods and supergames across all sessions and treatments which we randomly determined prior to the first experimental session. Overall, all participants played seven supergames (participants did not know the total number of supergames beforehand), the supergames consisted 5, 3, 7, 4, 7, 3 and 5 periods, respectively. We organized the rematching at the end of each supergame such that two new groups were randomly formed from a matching unit of 8 participants which remained constant for the entire duration of the session. This gave us 9 independent observations in *ProbRed*, *DamRed*, and *CertDam*, as well as 8 independent observations in *ProbRed*$^+$. Subjects earned an average of 10.50 Euro in the repeated prisoners' dilemma part, with a maximum of 12.70 Euro and a minimum of 8.25 Euro. After the main experiment, we assessed participants' risk preferences following Eckel and Grossman (2008) and Dave et al. (2010) with an average payoff of 38 Cent (minimum 2 Cent, maximum 70 Cent), before adding a survey which included questions on socio-demographic variables. During the experiment, participants played for Taler, at the end of the experiment, the sum of the payoffs in all rounds were converted into Euros at an exchange rate of 1 Taler for 1 Euro-Cent and paid out privately. Each session lasted for about 60 minutes. The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007), recruitment took place with hroot (Bock et al., 2014). The instructions (translated from German to English) can be found in the Appendix B.

4 Results

We structure our discussion of the results by first considering the average treatment differences, before explicitly exploring the individual learning dynamics.

Figure 2 shows the mean cooperation rates per period and treatment. Table 2 summarizes the average cooperation rates across all periods as well as for the first and last periods of the supergames. It is immediately seen that cooperation rates in *ProbRed* and *ProbRed*$^+$ are substantially higher than in *DamRed* and *CertDam*. Overall, cooperation rates across all periods are 59% in *ProbRed*, 54%

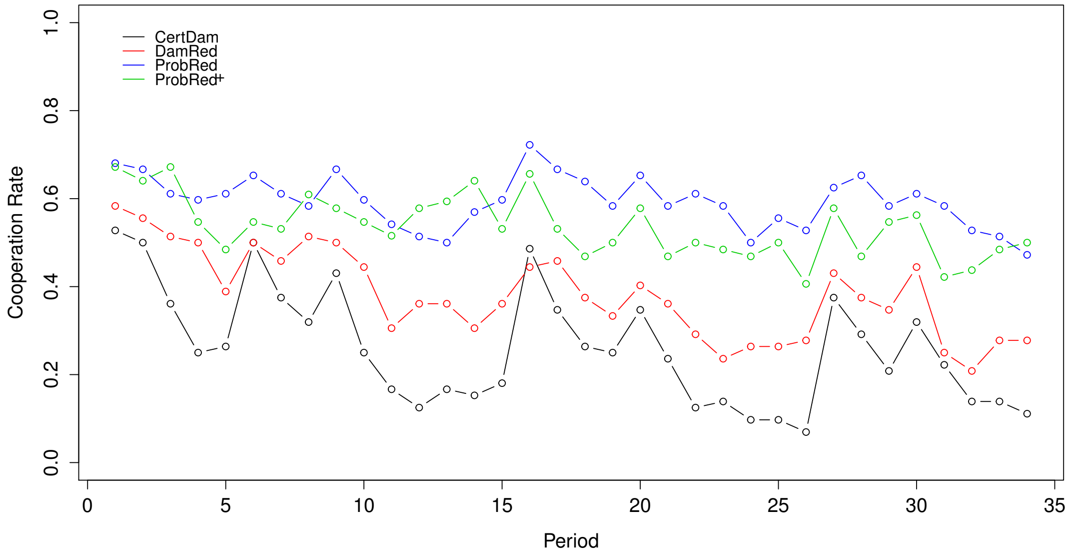

Figure 2: *Mean cooperation frequency per period by treatment.*

	all periods	first periods	last periods
(1) *CertDam*	.26	.43	.20
(2) *DamRed*	.38	.47	.36
(3) *ProbRed*	.59	.66	.57
(4) *ProbRed*$^+$	.54	.60	.51
tests	(3),(4)$>^{**}$(2) (3),(4)$>^{***}$(1)	(3)$>^{**}$(2) (4)$>^{*}$(2) (3),(4)$>^{***}$(1)	(3)$>^{**}$(2) (4)$>^{*}$(2) (3),(4)$>^{***}$(1)

Table 2: *Average cooperation rates by treatments over the entire experiment (left panel), over the first periods of all supergames (middle panel), and over the last periods of all supergames (right panel), tests refer to two-sided Wilcoxon Mann-Whitney rank sum tests,* *** *indicates significance at a* $p < 0.01$ *level,* ** *at a* $p < 0.05$ *level and* * *at a* $p < 0.1$ *level.*

in *ProbRed*$^+$, 38% in *DamRed*, and 26% in *CertDam*. More specifically, cooperation rates in *ProbRed* and *ProbRed*$^+$ are significantly larger than in *CertDam* ($p < 0.01$)[13] and *DamRed* ($p < 0.05$). No significant difference exists between *ProbRed* and *ProbRed*$^+$. These results are largely robust to concentrating on the first or the last periods of supergames as is displayed in Table 2. We therefore formulate our first result:

Result 1. *Cooperation rates are larger when cooperation affects the probability of a damage event (ProbRed and ProbRed*$^+$*) rather than affecting the size of a stochastic damage (DamRed) or when it leads to a certain damage reduction (CertDam).*

Result 1 is consistent with our predictions based on anticipated regret (section 3.2.2) and the learning dynamics in section 3.2.3. They could not be explained based on SPNE predictions as derived for expected utility maximizers in section 3.2.1. In fact, we find no significant impact of risk aversion on cooperation decisions in any of the treatments: Table 3 reports results from both a regression analyzing decisions in the first period of the first supergame (left panel) as well as a random effect regression (errors are clustered at the matching group level) analyzing decisions in all periods and all supergames. While confirming Result 1, the behavior does not appear to be driven by the individual's risk-aversion.[14]

The treatment differences reported in Result 1 qualitatively occur already in the very first period of the experiment: while 68% cooperate in *ProbRed*, 67% in *ProbRed*$^+$, only 58% cooperate in *DamRed* and 53% in *CertDam*. At the individual level (since each subject provides an independent observation in the first period of the first supergame), the differences between *CertDam* and *ProbRed* ($p = 0.06$) and *ProbRed*$^+$ ($p = 0.09$) are weakly significant based on two-sided Wilcoxon Mann-Whitney rank sum tests.[15]

The treatment differences are further strengthened over time as can be seen in Figure 2 as well as in Figure 3 which shows cooperation rates in the first period of the respective supergames. We find a negative trend of cooperation rates in the first periods of supergames in *DamRed* and *CertDam* (both $p = 0.05$, based on Cuzick's non-parametric test for trends), while the negative trend is not significant for the probability reduction treatments ($p = 0.19$ and $p = 0.13$, respectively).

Table 4 reports further evidence for the cooperation trends both within and across supergames based on a random-effects regression of the individual cooper-

[13]Throughout the paper and unless specified otherwise, statistical significance is assessed by two-sided Wilcoxon Mann-Whitney rank sum tests relying on matching unit averages.

[14]Risk attitudes are measured by the lottery choice in the second part of the experiment (the variable *risk*, ranges from one to six, such that the lottery choice with larger numbers indicates more risk tolerance).

[15]No significant differences occur when controlling for risk-aversion (see Table 3).

	dependent variable: q_i^t	
	only first period	*all periods*
DamRed	.251 (.206)	.11 (.148)
ProbRed	.256 (.19)	.278 (.169)
ProbRed$^+$	.272 (.189)	.191 (.137)
risk	.04 (.033)	−.022 (.017)
risk × *DamRed*	−.053 (.048)	.006 (.03)
risk × *ProbRed*	−.029 (.048)	.016 (.039)
risk × *ProbRed*$^+$	−.036 (.047)	.024 (.029)
constant	.386*** (.13)	.339*** (.092)
obs	280	9520
n	280	280
F-test/Wald-Chi^2-test	.92	68***

Table 3: *Left panel: linear regression of cooperation behavior in the first period, right panel: random effects regression of cooperation behavior in all periods of the experiment; coefficients are reported along standard errors in parenthesis (errors are clustered at the matching group level);* *** *indicates significance at a* $p < 0.01$ *level,* ** *at a* $p < 0.05$ *level and* * *at a* $p < 0.1$ *level. obs reports the number of observations while* n *reports the number of subjects; models' fitness are assessed by F-test and Wald-Chi*^{2}*-tests.*

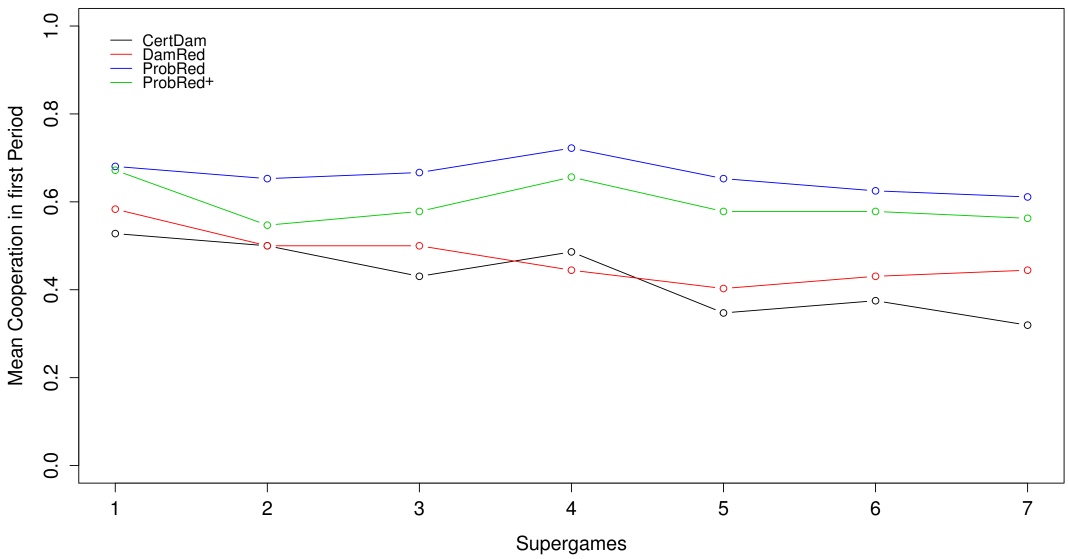

Figure 3: *Mean cooperation in the first period of all supergames across treatment conditions*

ation decision on the supergame (*supergame*, ranging form 1 to 7) and the period within a supergame (*period in supergame*, ranging from 1 to 7) as well as on dummies for the treatments and the corresponding interaction terms. We find negative time trends across supergames in *DamRed* and *CertDam*, and a significantly less negative trend in *ProbRed*$^+$, while there is no significant trend in *ProbRed*.[16] Negative time trends also occur within supergames in all treatments.[17] The downward trend within supergames is largest in *CertDam*, significantly smaller in both *DamRed* and *ProbRed* and weakest in *ProbRed*$^+$.

	dependent variable: q_i^t
DamRed	.06 (.071)
ProbRed	.147* (.078)
ProbRed$^+$	.099 (.08)
supergame	$-$.031*** (.006)
supergame × *DamRed*	$-$.004 (.01)
supergame × *ProbRed*	.02** (.009)
supergame × *ProbRed*$^+$	.01 (.008)
period in supergame	$-$.055*** (.007)
period in supergame × *DamRed*	.025*** (.009)
period in supergame × *ProbRed*	.033*** (.01)
period in supergame × *ProbRed*$^+$	.043*** (.008)
constant	.558*** (.057)
obs	9520
n	280
Wald-Chi^2-test	173***

Table 4: *Random-effects linear regression of time trends for individual cooperation decision q_i^t; coefficients are reported along standard errors in parenthesis (errors are clustered at the matching group level); *** indicates significance at a $p < 0.01$ level, ** at a $p < 0.05$ level and * at a $p < 0.1$ level. obs reports the number of observation while n report the number of subjects; model's fitness is assessed by a Wald-Chi^2-test.*

Result 2. *Cooperation rates follow different time trends: the downward trend is strongest in CertDam, less strong in DamRed and least in ProbRed$^+$ and ProbRed.*

The different time trends, in particular across supergames, are consistent with the predicted treatment differences due to different learning dynamics as discussed

[16] According to F-Tests, testing that *superg* × *treatment* + *supergame* is statistically different from zero for all treatments at $p < 0.05$ except *ProbRed* ($p = 0.146$).

[17] According to F-Tests, testing that *period in superg* × *treatment* + *period in supergame* is statistically different from zero for all treatments at $p < 0.05$.

in section 3.2.3. They support the idea that probability reduction in *ProbRed* and *ProbRed*$^+$ results in a slower average learning of defection than in *DamRed* and *CertDam*, both within supergames and across supergames.

However, the discussion so far has concentrated on average time trends while the predicted learning dynamics were based on individual experiences. To gain further insights into the different time trends, we therefore now investigate determinants of behavioral adjustments at the individual level. Given the prediction in section 3.2.3, we particularly expect the occurrence of a damage event to be important in determining future choices in *ProbRed* and *ProbRed*$^+$ when the player was or believes to have been pivotal. For both *DamRed* and *CertDam*, the learning rules would always predict a move towards defection which in *CertDam* is less strong following a damage event.

In a first step, we consider the conditional frequencies of $q_i^{t+1} = 1$ given q_i^t and the occurrence of the damage s^t. Table 5 summarizes the frequencies by treatment conditions as well as the significant differences based on nonparametric Mann-Whitney tests.

	damage in t: $s^t = 1$		no damage in t: : $s^t = 0$	
	$q_i^t = 0$	$q_i^t = 1$	$q_i^t = 0$	$q_i^t = 1$
(1) *CertDam*	.13	.60	-	-
(2) *DamRed*	.16	.65	.17	.76
(3) *ProbRed*	.22	.79	.13	.90
(4) *ProbRed*$^+$	.23	.68	.31	.74
tests	(3),(4)$>^{**}$(1) (4)$>^{**}$(2)	(3)$>^{***}$(1),(2)	(4)$>^{**}$(2) (4)$>^{***}$(3)	(3)$>^{***}$(2),(3)

Table 5: *Mean q_i^{t+1} given q_i^t and the occurrence of the damage s^t; tests refer to two-sided Wilcoxon Mann-Whitney rank sum tests, *** indicates significance at a $p < 0.01$ level, ** at a $p < 0.05$ level and * at a $p < 0.1$ level.*

Overall, it seems that the effect of probability reduction on cooperation is twofold: it leads to more stable cooperation of those players who already cooperate (their frequency to choose $q_i^{t+1} = 1$ is about 15% higher in *ProbRed*), and it induces non-cooperating players to cooperate after a damage event occurred (the frequency to choose $q_i^{t+1} = 1$ is 6-10% higher in *ProbRed* and *ProbRed*$^+$). That is to say, the "all-or-nothing" damage of *ProbRed* and *ProbRed*$^+$ prevents players from choosing defection and additionally leads more defecting players to switch to cooperation.[18]

[18]Notice that there is also a surprising effect in *ProbRed*$^+$ for non-cooperators if the damage did not occur: here, the frequency of cooperation in $t + 1$ is 14-18% higher than in the other treatment conditions. While we are lacking a clear explanation, this finding may be driven by

For a detailed analysis of individual learning in our game, we estimate a series of Arellano-Bond panel regressions, for each treatment condition separately.[19] This allows us to analyze endogenous regressors (see Arellano and Bond, 1991): the dependent variable is q_i^{t+1} (i.e., the decision whether to cooperate or defect in the consecutive period), dependent variables are q_i^t (i.e., the decision whether to cooperate or defect in the current period), Q_{-i}^t (i.e., the number of cooperators except i in the current period), the occurrence of the damage in t (i.e., we compute a dummy variable s^t which is one if the damage occurred in t and zero otherwise; omitted in *CertDam*), and interaction terms $q_i^t \times s^t$, as well as $Q_{-i}^t \times s^t$ to see whether actually realized damage reduction (in *DamRed*), or unsuccessful cooperation (in *ProbRed* and *ProbRed*$^+$) leads to different learning than when a damage does not occur.

To access the additional information provided in *ProbRed*$^+$, we additionally introduce a variable measuring the number of cooperators exceeding the necessary number to avoid the realization of the damage. That is, the variable $\Delta cooperator$ computes the difference between the actual players cooperating and the cooperators required by nature for the absence of the damage. $\Delta cooperator$ is zero if the number of cooperators just coincides with the number required to avoid the damage, it is negative if too few players cooperate to prevent the damage and is positive if even a smaller number of cooperators were necessary to prevent the damage. Hence, we test whether players coordinate their cooperation onto the sufficient number of cooperators in the previous period. Estimations for coefficients along standard errors in parenthesis are reported in Table 6.

the additional information that these players receive relative to *ProbRed*. We control for this effect in our following analysis.

[19] Arellano-Bond is typically applied to continuous rather than discrete dependent variables. However, we are not aware of a fully consistent method which can both incorporate the lagged contribution variable as well as control for the interdependencies at the individual and matching unit level. Our results are, however, robust to alternative specifications like random effects probit model, or OLS regressions with individually clustered errors.

	dependent variable: q_i^{t+1}				
	CertDam	*DamRed*	*ProbRed*	*ProbRed*$^+$	*ProbRed*$^+$
q_i^t	$.215^{***}$ (.022)	$.368^{***}$ (.03)	$.29^{***}$ (.027)	$.06^{**}$ (.028)	$.087^{***}$ (.029)
Q_{-i}^t	$.077^{***}$ (.012)	$.043^{***}$ (.016)	$.029^{**}$ (.013)	$.06^{***}$ (.016)	$.089^{***}$ (.018)
s^t		$.106^{***}$ (.03)	$.114^{***}$ (.036)	.045 (.043)	$-.086$ (.055)
$q_i^t \times s^t$		$-.395^{***}$ (.038)	$-.26^{***}$ (.032)	$-.179^{***}$ (.042)	$-.186^{***}$ (.042)
$Q_{-i}^t \times s^t$		.017 (.02)	.017 (.018)	.01 (.023)	.003 (.023)
$\Delta cooperator$					$-.04^{***}$ (.011)
const	$.130^{***}$ (.012)	$.194^{***}$ (.023)	$.360^{***}$ (.029)	$.407^{***}$ (.033)	$.413^{***}$ (.033)
obs	2232	2232	2232	1984	1984
n	72	72	72	64	64
Wald-Chi^2-test	188^{***}	183^{***}	147^{***}	42^{***}	55^{***}

Table 6: *Estimation results for an Arellano-Bond panel regressions with dependent variable q_i^{t+1}; coefficients are reported along standard errors in parenthesis;* *** *indicates significance at a $p < 0.01$ level,* ** *at a $p < 0.05$ level and* * *at a $p < 0.1$ level. Standard errors are clustered at the matching group level. obs reports the number of observation while n reports the number of subjects; models' fitness are assessed by Wald-Chi^2-tests.*

The estimation results in Table 6 confirm our previous findings in Table 5. They indicate that cooperation is highly path dependent in all treatment conditions: if a player cooperates in period t, it is very likely that she cooperates in period $t+1$ as well (significant positive marginal effect of q_i^t). For all treatments, we also find evidence for conditional cooperation (significant positive coefficients for Q_{-i}^t). Experiencing a damage event also triggers behavioral changes: non-cooperators are more likely to switch to cooperation following a damage event in both *ProbRed* and *DamRed* (significant positive coefficients for s^t) which is consistent with the learning dynamics detailed in section 3.2.3.[20] We further find significant negative coefficients for the interaction $q_i^t \times damage^t$. Hence, a damage event typically reduces the likelihood for cooperators to continue cooperation (or at least does lead to significantly smaller increases than found for defectors). Again, this is consistent with section 3.2.3 for *ProbRed* and *ProbRed*$^+$ where cooperators may regret their action as it did not prevent the damages. The negative coefficient for *DamRed* cannot be explained by the discussed learning dynamics.[21] In addition, we can support our expectation concerning the coordination of behavior in *ProbRed*$^+$: the significant negative coefficient of $\Delta cooperator$ suggests that players condition their cooperativeness on the number of cooperators needed to prevent the damages in the previous period: if there are more (less) players than needed to avoid the damage, the likelihood to cooperate decreases (increases).

Result 3. *In all treatment conditions with stochastic payoffs, the prevention of the damage reinforces existing cooperation while the occurrence of a damage stimulates a strategy switch of players from defection to cooperation.*

5 Conclusion

This paper investigates determinants of cooperation in repeated social dilemmas with stochastic damages. Inspired by environmental problems like climate policy, or hurricane prevention, but also other challenges like public security protection against terror, or international health cooperation against pandemic diseases, we study the evolution of cooperation when the entire group benefits from individual cooperation while individual players have incentives to free-ride and may cooperate only due to (indefinitely) repeated interactions. With stochastic damages, players may take actions which either reduce the size of damages or reduce the probability that such adverse events occur.

[20] This effect seems to be dominated by the coordination of cooperation in *ProbRed*$^+$.

[21] One may speculate that cooperating players following a damage event in *DamRed* realize that even here defection was a dominant strategy and therefore better learn the rules of the game.

Our results show that cooperation on probability reduction leads to significantly higher cooperation rates than cooperation on damage reduction. For probability reduction, the cooperation rates are sustained over time, whereas they decline over time for damage reduction as well as in a setting where damages are certain. This result contradicts predictions based on subgame perfect Nash equilibria for expected utility maximizers. However, we show a combination of ex ante regret aversion and a learning dynamics which reinforces the ex post optimal action can explain the difference between the two settings. The experimental data suggest that the absence of a damage event tends to reinforce individual cooperation, while non-cooperating players are more likely to switch behavior following an adverse event.

Overall, our results may provide some optimistic view on the prospects of voluntary cooperation in dilemma situations: differently from situations where cooperation leads to (continuous) changes in the size of damages (or payoffs), more sustained cooperation can be expected if it may lead to a discrete payoff change as an adverse event may be prevented with some probability. Cautiously interpreting the results from our lab experiment in terms of current debates on climate policy, our findings suggest that shifting the public attention from activities which are likely to reduce the occurrence of extreme negative events (mitigation activities) to measures which reduce their impact (e.g., adaptation) may lead to a declining individual willingness to cooperate. Our experiment also indicates that in line with our introductory discussion of natural disasters or accidents serving as "focussing events" (Kingdon, 1995; Birkland, 2006), experiencing adverse events indeed may lead to behavioral changes, even if cooperation is voluntary and not enforced by modified policies.

References

[1] Aoyagi, M. and G. Fréchette (2009). Collusion as public monitoring becomes noisy: Experimental evidence. *Journal of Economic Theory* 144, 1135-1165.

[2] Arellano, M., and S. Bond (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *The Review of Economic Studies* 58, 277-297.

[3] Barrett, S. (2011). Avoiding disastrous climate change is possible but not inevitable, *PNAS* 108, 11733-11734.

[4] Becker, G.S. (1968). Crime and punishment, an economic approach. *Journal of Political Economy* 76, 169-217.

[5] Beggs, A.W. (2005). On the convergence of reinforcement learning. *Journal of Economic Theory* 122, 1-36.

[6] Bereby-Meyer Y., and A.E. Roth (2006). The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation. *The American Economic Review* 96, 1029-1042.

[7] Birkland, T.A. (2006). Lessons of disaster: poliocy change after catastrophic events. Georgetown University Press, Washington DC.

[8] Bò, P.D. and G. Fréchette (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *The American Economic Review* 101, 411-429.

[9] Bock, O., A. Nicklisch and I. Baetge (2014). hroot: Hamburg registration and organization online tool. *European Economic Review* 71, 117-120.

[10] Briys, E. and H. Schlesinger (1990). Risk aversion and the propensities for self-insurance and self-protection. *Southern Economic Journal* 57, 458-467.

[11] Camerer, C.F. and T.H. Ho (1999). Experience-weighted attraction learning in normal form games. *Econometrica* 67, 837-874.

[12] Chmura, T., S. Goerg and R. Selten (2012). Learning in experimental 2x2 games. *Games and Economic Behavior* 76, 44-73.

[13] Dannenberg, A., A. Löschel, G. Paolacci, C. Reif, and A. Tavoni (2014). On the provision of public goods with probabilistic and ambiguous thresholds. *Environmental and Resource Economics*.

[14] Dave, C., C.C. Eckel, C.A. Johnson and C. Rojas (2010). Eliciting risk preferences: When is simple better? *Journal of Risk and Uncertainty* 41, 219-243.

[15] Dickinson, D.L. (1998). The voluntary contribution mechanism with uncertainty group payoffs. *Journal of Economic Behavior and Organization* 35, 517-533.

[16] Dionne, G. and L. Eeckhoudt (1985). Self-insurance, self-protection and increased risk aversion. *Economics Letters* 17, 39-42.

[17] Eckel, C.C. and P.J. Grossman (2008). Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior and Organization* 68, 1-17.

[18] Ehrlich, I. and G.S. Becker (1972). Market insurance, self-insurance, and self-protection. *Journal of Political Economy* 80. 623-648.

[19] Erev, I. and A. Roth (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review* 88, 848-881.

[20] Filiz-Ozbay, E., and Ozbay, E. Y. (2007). Auctions with anticipated regret: Theory and experiment. *The American Economic Review*, 1407-1418.

[21] Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics* 10, 171-178.

[22] Friesen, L. (2012). Certainty of punishment versus severity of punishment: An experimental investigation. *Southern Economic Journal* 79, 399-421.

[23] Gollier, C. and J.W. Pratt (1996). Risk vulnerability and the tempering effect of background risk. *Econometrica* 64, 1109-1123.

[24] Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior* 13, 243-266.

[25] Ho, T.H., Camerer, C.F. and J.-K Chong (2008). Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory* 133, 177-198.

[26] Holmström, B. (1982). Moral hazards in teams. *The Bell Journal of Economics* 13, 324-340.

[27] Ihori, T. and M.C. McGuire (2010). National self-insurance and self-protection against adversity: Bureaucratic management of security and moral hazard. *Economics of Governance* 11, 103-122.

[28] Jullien, B., B. Salanie and F. Salanié (1999). Should more risk-averse agents exert more effort? *The Geneva Papers on Risk and Insurance Theory* 24, 19-28.

[29] Lohse, T., J.R. Robledo and U. Schmidt (2012). Self-insurance and self-protection as public goods. *Journal of Risk and Insurance* 79, 57-76.

[30] Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 805-824.

[31] Keser, C. and C. Montmarquette (2008). Voluntary contributions to reduce expected public losses. *Journal of Economic Behavior and Organization* 66, 477-491.

[32] Kingdon, J. W. (1995). Agendas, alternatives, and public policities. 2nd edition. Harper Collings, New York.

[33] Miceli, T. J. and K. Segerson (2007). Punishing the innocent along with the guilty: The economics of individual versus group punishment. *Journal of Legal Studies* 36, 81-106.

[34] Milinski, M., R.D. Sommerfeld, H.-J. Krambeck, F.A. Reed and J. Marotzke (2008). The collective-risk social dilemma and the prevention of simulated dangerous climate change. *PNAS* 105, 2291-2294.

[35] Muermann, A. and H. Kunreuther (2008). Self-protection and insurance with interdependencies. *Journal of Risk and Uncertainty* 36, 103-123.

[36] Rasmussen, E. (1987). Moral hazard in risk averse teams. *Rand Journal of Economics* 18, 428-435.

[37] Roth, A. and I. Erev (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8, 164-212.

[38] Segerson, K. (1988). Uncertainty and incentives for non-point pollution control. *Journal of Environmental Economics and Management* 15, 87-98.

[39] Selten, R. and T. Chmura (2008). Stationary concepts for experimental 2x2-games. *The American Economic Review* 98, 938-966.

[40] Selten, R. and R. Stöcker (1986). End behavior in sequences of finite prisoner's dilemma supergames. *Journal of Economic Behavior and Organization* 7, 47-70.

[41] Tavoni, A., A. Dannenberg, G. Kallis and A. Löschel (2011). Inequality, communication and the avoidance of disastrous climate change in a public goods game. *Proceedings of the National Academy of Sciences* 108, 11825-11829.

[42] Varian, H.R. (1990). Monitoring agents with other agents. *Journal of Institutional and Theoretical Economics* 46, 153-174.

[43] Zeelenberg, M. (1999). Anticipated regret, expected feedback and behavioral decision making. *Journal of behavioral decision making*, 12(2), 93-106.

Appendix A: Derivation of equilibrium conditions for risk-averse players

We define value functions as capturing the expected utility to the player from the continuation of the game (within the same supergame): $V^{Treat}(0,0)$ if cooperation already broke down, $V^{Treat}(0,Q-1)$ if the player defects in a period while $Q-1$ players still cooperate, anticipating that cooperation fully breaks down in subsequent periods and $V^{Treat}(1,Q)$ if the player cooperates and a total of Q players cooperate.

Given our notation, $V^{Treat}(0,0)$ is implicitly defined by:

$$\begin{aligned} V^{Treat}(0,0) &= \left[p^{Treat}(0)\exp(-\sigma(E-D^{Treat}(0))) + (1-p^{Treat}(0))\exp(-\sigma E)\right] \\ &\quad \times \left[(1-\delta)V^{Treat}(0,0) + \delta u(0)\right] \end{aligned} \tag{12}$$

Here the first bracket reflects the period payoff, while the second bracket the expected continuation utility.[22] Denoting

$$\begin{aligned} A^{Treat}(q,Q) &= p^{Treat}(Q)\exp(-\sigma(E-cq-D^{Treat}(Q))) \\ &\quad +(1-p^{Treat}(Q))\exp(-\sigma(E-cq)), \end{aligned} \tag{13}$$

equation (12) can be rewritten as

$$V^{Treat}(0,0) = \frac{\delta u(0)A^{Treat}(0,0)}{1-(1-\delta)A^{Treat}(0,0)}. \tag{14}$$

$V^{Treat}(0,Q-1)$ then is given by (taking into account the modified grim trigger strategies):

$$\begin{aligned} V^{Treat}(0,Q-1) &= A^{Treat}(0,Q-1)\left[(1-\delta)V^{Treat}(0,0)+\delta u(0)\right] \\ &= \frac{\delta u(0)A^{Treat}(0,Q-1)}{1-(1-\delta)A^{Treat}(0,0)}. \end{aligned} \tag{15}$$

[22]Note that the utility function allows for the following decomposition of the expected utility in a term containing x_0 and another one containing $\sum_{t=1}^{\infty} x_t$: $u(\sum_{t=0}^{\infty} x_t) = \exp(-\sigma x_0)u(\sum_{t=1}^{\infty} x_t)$.

Finally, anticipating stable cooperation, we obtain

$$\begin{aligned} V^{Treat}(1,Q) &= A^{Treat}(1,Q)\left[(1-\delta)V^{Treat}(1,Q)+\delta u(0)\right] \\ &= \frac{\delta u(0)A^{Treat}(1,Q)}{1-(1-\delta)A^{Treat}(1,Q)} \end{aligned} \tag{16}$$

With this notation, under constant absolute risk aversion preferences, cooperation is indeed stable if

$$\begin{aligned} V^{Treat}(1,Q) &\geq V^{Treat}(0,Q-1) \\ \Leftrightarrow \frac{\delta u(0)A^{Treat}(1,Q)}{1-(1-\delta)A^{Treat}(1,Q)} &\geq \frac{\delta u(0)A^{Treat}(0,Q-1)}{1-(1-\delta)A^{Treat}(0,0)}. \end{aligned} \tag{17}$$

This expression gives the condition under which the incentive to cooperate is higher than to defect and is used as the basis for the simulations reported in Figure 1.

Appendix B: Experimental Instructions for the DamRed Treatment (English translation)

In the following, we report an English translations of the experimental instructions for the *DamRed* treatment.

General instructions for the participants

You are now taking part in an economic science experiment. If you carefully read the following instructions, you can - depending on your decisions - earn a not inconsiderable amount of money. Therefore, it is very important that you carefully read the following instructions.

The instructions that we gave you are solely meant for your private information.**During the experiment, communication is completely prohibited.** If you have any questions, please raise your hand out of the cabin. Someone will then come to you to answer your question. Violation of this rule leads to exclusion from the experiment and from all payments.

During the experiment we do not have Euro but Taler. Your total income will first be computed in Taler. The total amount of Taler that you earned during the experiment will be converted into Euro in the end, such that

100 Taler = 1 Euro.

At the end of the experiment you will be paid in cash the total amount of Taler that you earned (converted into Euro) plus 5 Euro for participation. We will conduct the payment such that no other participant will see your payment.

The experiment is divided into two parts. Here, we give the instructions for the 1st part. You will get the instructions for the 2nd part on your computer screen after the 1st part is finished. The two parts are not related with respect to their content.

Explanations for the 1st part of the experiment

The 1st part of the experiment is divided into **phases**. You do not know, however, how many phases there are in total. Each phase is divided into **rounds**. The number of rounds in a phase is random. After each round, the phase ends with a probability of 20%.

More concretely, this means that: after the first round there is a second round with a probability of 80% (which is on average in four cases out of five). So, with a probability of 20% (which is on average in one case out of five) the phase ends after the first round. After the second round (if there is one) there is a third one with a probability of 80%. So, with a probability of 20%, the phase ends after the second round and so on...

At the beginning of each phase, participants are randomly assigned into groups of four. Thus, your group has three other members in addition to you. **During one phase, the constellation of the group remains unchanged. It only gets randomly rematched at the beginning of a new phase.**

Information on the structure of a round

All rounds in all phases are always structured in the exact same way. In the following we describe the structure of one round.

At the beginning of each round, every participant gets an **income** of 25 Taler.

At the end of each round, a **damage** might occur, which reduces the income by 20 Taler.

The damage occurs with a probability of 50% (which is on average in one of two rounds). For this, in each round, the computer randomly determines whether the damage occurs. The occurrence of the damage is only valid in the respective round and does not influence the probability of the next rounds. The occurrence of the damage is determined jointly for the whole group, such that either all or no group members suffer the damage.

All group members are able to reduce the potential damage through their decisions. **For this, at the beginning of each round, i.e. before the damage occurs, each group member has to decide whether it does or does not carry out a damage-reducing action** (see figure 1 at the end of the instruction).

Each damage-reducing action costs the group member taking the action 5 Taler (independent of whether the damage occurs or not). Each damage-reducing action reduces the personal damage of **each** group member (not only of the group member taking the action) by 4 Taler. For you, personally, this means that each damage-reducing action that has been carried out in your group reduces your damage (if it occurs) by 4 Taler, independent of whether you have taken such an action yourself. **A damage-reducing action which you carry out costs you 5 Taler for sure. In return, you reduce your damage and the damage of each other group member by 4 Taler, if the damage occurs.**

The personal damage, if it occurs, amounts to **20 Taler** if no one in your group carried out an action, **16 Taler** if one person carried out an action, **12 Taler** if two persons took the action, **8 Taler** if three persons took the action and **4 Taler** if all group members took the action.

Your round income (in Taler) is calculated as follows

- If the damage **does not** occur and you **did not** take the damage-reducing action:

  25

- If the damage **does not** occur and you did take the damage-reducing action:

  $25 - 5 = 20$

- If the damage occurs and you **did not** take the damage-reducing action:

  $25 - 20 + 4*$ [sum of all damage-reducing actions in the group],

- If the damage occurs and and you did take the damage-reducing action:

  $25 - 5 - 20 + 4*$ [sum of all damage-reducing actions in the group],

<u>4 examples:</u>

The damage probability always is 50%.

- **You** and **one other** group member take a damage-reducing action in your group, the damage **does not** occur. Your round income is $25 - 5 =$ **20** Taler.

- **Only you** take a damage-reducing action in your group, the damage occurs. Your round income is $25 - 5 - 20 + 4 * 1 =$ **4** Taler.

- **You** and **two other** group members take a damage-reducing action, the damage occurs. Your round income is $25 - 5 - 20 + 4 * 3 =$ **12** Taler.

- **Two other** group members take a damage-reducing action, but you do not, the damage occurs. Your round income is $25 - 20 + 4 * 2 =$ **13** Taler.

At the end of a round, each participant receives information on whether he/she took an action him- or herself, how many other group members took an action, if the damage occurred and what the round income is. Then, a new round starts in the same group constellation or in a new group constellation if a new phase begins.

The **sum of all your round incomes** will be paid out to you in private at the end of the experiment.

Before the experiment starts, we would like to ask you to answer some control questions on the computer to make sure you understand the rules.

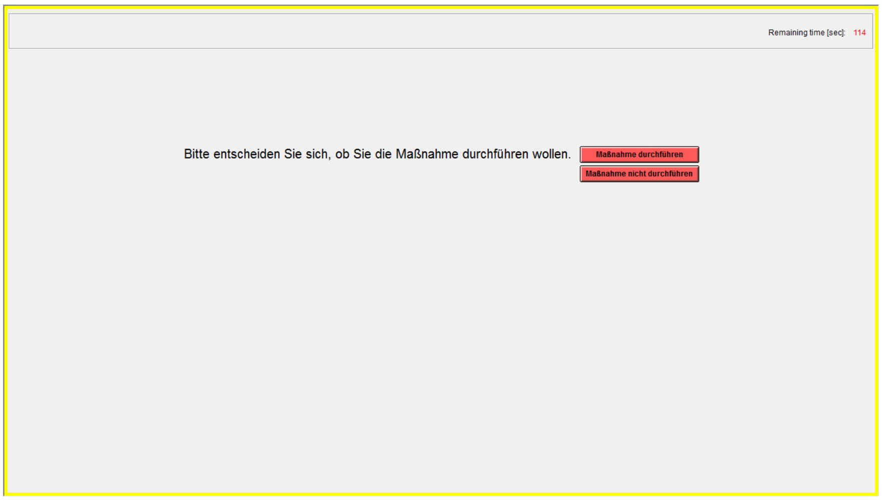

Figure 4: Decision screen for taking the action in a round in Part 1 of the experiment.

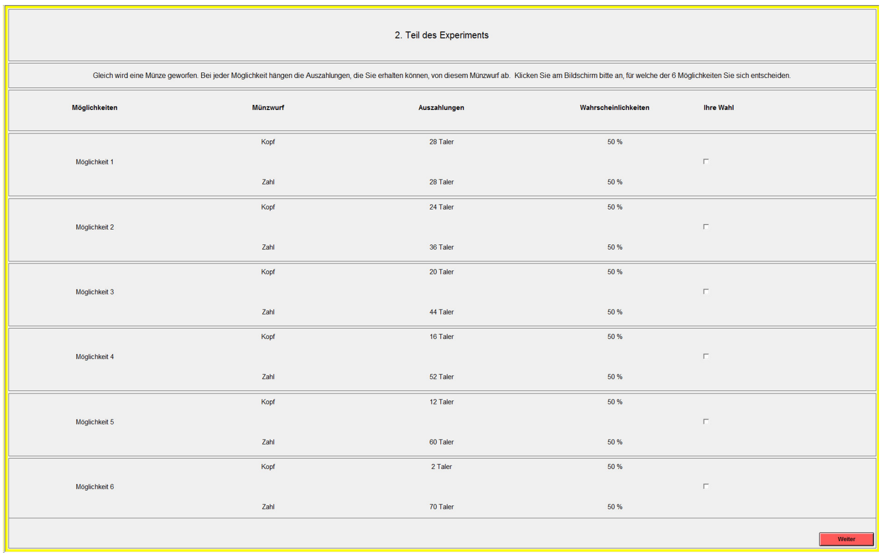

Figure 5: Decision screen for the risk-assessment task in Part 2 of the experiment.