Payoffs, Beliefs, and Cooperation in Infinitely Repeated Games

Preregistration

Maximilian Andres* Lisa Bruttel† Juri Nithammer‡

January 11, 2024

Abstract

This is the preregistration of a laboratory experiment designed as an empirical test of our paper "Payoffs, Beliefs, and Cooperation in Infinitely Repeated Games," available at SSRN (see https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4491762). This document contains a description of the experimental design and procedures, the hypotheses and planned statistical analysis, the results of the power analysis, as well as the instructions from the experiment.

*University of Potsdam, Germany. maximilian.andres@uni-potsdam.de
†University of Potsdam, Germany. lisa.bruttel@uni-potsdam.de
‡University of Potsdam, Germany. juri.nithammer@uni-potsdam.de

1 Research Question and General Setup

In our model (see Andres et al., 2023), we show that a player's belief about the probability of cooperation by their opponent moderates the effect of changes in the payoff parameters on cooperation in a prisoner's dilemma. If beliefs are optimistic, increasing the gain from unilateral defection has a large negative effect on cooperation, while increasing the loss from unilateral cooperation has a negligible effect. However, if beliefs are pessimistic, increasing the gain has only a negligible effect, while increasing the loss has a large negative effect on cooperation.

Testing the hypothesis that participants with different beliefs react differently to changes in the payoffs requires an exogenous variation of beliefs. In our experimental setup, we induce beliefs by informing participants truthfully that we will match them to a partner only after they have chosen their strategy in the infinitely repeated game and by telling them ex ante the probability that we will match them to a partner who plays Grim or AlwaysDefect, respectively. The probability that defines this matching is equivalent to the belief p as laid out in Andres et al. (2023).

Inducing beliefs in this way has two advantages in comparison to using elicited own beliefs of the participants. First, we avoid that beliefs vary with the payoffs (Aoyagi et al., 2023; Andres, 2023), which would affect how changes in the payoffs affect cooperation, see Proposition 2 in Andres et al. (2023). Second, imposed beliefs are exogeneous variables in the econometric sense, i.e. they are neither correlated with the cooperation decision (the dependent variable) nor with individual characteristics that may be unobserved but also correlated with the cooperation decision. See Costa-Gomes et al. (2014) for a discussion of the endogeneity of beliefs.

2 Experimental Design and Procedures

Basic Setup Participants in the experiment face different versions of the infinitely repeated prisoner's dilemma game. The left part of Table 1 shows the game in its general form, where parameters are such that $T > R > P > S$ and $2 \cdot R > T + S$. Following Stahl II (1991), we can normalize the payoffs to reduce the prisoner's dilemma to a function $\Gamma(g, l)$ of the gain from unilateral defection g and the loss from unilateral cooperation $-l$, with $g, l > 0$. The normalization subtracts the punishment payoff P from the original payoff R, S, T or P and then divides by $R - P$. The right part of Table 1 presents the resulting payoffs after the normalization: The gain from unilateral defection is $\frac{T-P}{R-P} = 1+g$, and the loss from unilateral cooperation is $\frac{S-P}{R-P} = -l$, while the reward and punishment payoffs are $\frac{R-P}{R-P} = 1$ and $\frac{P-P}{R-P} = 0$, respectively. In the following, we refer to this normalized form of the game.

(a) Original

	C	D
C	R, R	S, T
D	T, S	P, P

(b) Normalized

	C	D
C	$1, 1$	$-l, 1+g$
D	$1+g, -l$	$0, 0$

Table 1: Stage-Game Payoffs in the Prisoner's Dilemma Γ

Experimental Design All participants choose strategies in three different prisoner's dilemma games that vary the stage-game payoff-parameters (see Table 2). The baseline game, named Base, has parameters $g = l = 0.2$. The HighGain game has parameters $g = 1,\ l = 0.2$. The HighLoss game has parameters $g = 0.2,\ l = 1$. In the experiment, the stage-game payoff tables are shown to the participants in the standard notation, using T, S, P, R, indicating the amount of points they receive. The order in which participants face these three games is randomized between sessions, but participants always start with Base.

Table 2: Row Player's Stage-Game Payoffs in the Three Games

Base

	C	D
C	75	45
D	80	50

	C	D
C	1	-0.2
D	$1+0.2$	0

HighGain

	C	D
C	75	45
D	**100**	50

	C	D
C	1	-0.2
D	$1+\mathbf{1}$	0

HighLoss

	C	D
C	75	**25**
D	80	50

	C	D
C	1	$-\mathbf{1}$
D	$1+0.2$	0

In the experiment, participants play infinitely repeated versions of all three games. For each of the stage-games, participants choose among the two strategies Grim and AlwaysDefect. They make this choice for nine possible discount factors $\delta \in \{0.1, ..., 0.9\}$. Thus, for each participant, we can derive three critical discount factors from their choices, one for each stage-game parametrization.

Treatments We compare three treatments in which beliefs vary exogenously between participants. At the beginning of a session, each participant is informed about the probability with which they will be matched to a player that has played Grim/AlwaysDefect in a given decision. This probability stays constant and participants are reminded of it on each decision screen. The probability is randomly assigned, drawn individually from a discrete uniform distribution over $\{0.1, 0.5, 0.9\}$, resulting in the three treatments LowBelief, MedBelief and HighBelief.

Filtering Rule If participants' preferences satisfy the independence axiom, it has to hold that if a participant prefers Grim over AlwaysDefect for a specific discount factor $\bar{\delta}$, then they must prefer Grim for all $\delta \geq \bar{\delta}$ and AlwaysDefect for all $\delta < \bar{\delta}$. Only then a unique critical discount factor can be computed for this participant. Participants' actual decisions in the experiment could violate the independence axiom in two ways. First, they might switch more than once between strategies. Second, they may switch into the "wrong" direction, i.e., choose Grim for low discount factors and AlwaysDefect for high ones. We deal with this issue in the following way.

We enforce that participants switch at most once between the strategies because with more than one switch we would obtain more than one critical discount factor for the same decision by the same individual, which we could not interpret in a meaningful way. However, we allow that participants switch into the "wrong" direction because such switches provide a clear-cut filtering rule as it is impossible to calculate a critical discount factor for such a participant. Therefore, we will exclude participants that choose AlwaysDefect for low discount factors and switch to Grim for higher ones in at least one of the three games from all analyses, since our outcome variable is undefined for these participants. The sample size N we specify below refers to valid observations in the sense that critical discount factors can be computed for these participants in all three games.

Procedures Experiments will be conducted at the Potsdam Laboratory for Economic Experiments (PLEx). The subject pool consists mainly of students enrolled at the University of Potsdam or other universities in Potsdam. We expect that sessions last about 75 minutes including welcoming and payoff procedures. Participants will be paid in cash at an exchange rate of 1:18, plus a showup fee of eight euros. We expect that participants will earn about 19 euros on average. Payments are administered privately. Participants will be asked to sign an informed consent sheet, detailing data protection issues and including a very broad description of the experiment.

If our subject pool does not provide enough participants, we will collect the remaining observations at the laboratory of the Technical University of Berlin, making sure that treatment assignment is balanced across the two locations.

Instructions The experimental instructions describe the underlying game and the concept of strategies to the participants. Actions are denoted as "X" (C) and "Y" (D), and strategies as "A" (Grim) and "B" (AlwaysDefect). Participants are told that they decide between strategies, denoted as "plans" and we explain to them what these strategies mean and what each of the four possible joint strategies entails. We also explain how stage-game payoffs decrease over time, depending on the discount factor. On each decision screen, participants have access to a payoff-calculator for the infinitely repeated game in addition to the stage-game payoff-table. By request of the participant, this payoff-calculator displays the discounted payoffs for the four possible joint strategies and each possible discount factor.

In the beginning of the experiment, participants go through a guided tour of the experiment on their computer screen. During this tour, they have to try out the payoff-calculator in order to answer some comprehension questions. Furthermore, they learn how to enter decisions in the price-list format in which the strategies are presented.

In the instructions, we also explain the matching procedure in detail. We tell participants that only one of their 27 decisions (three games × nine discount factors) will be payoff-relevant, and that they will be matched to another participant only after the decision and based on the other participant's decision. With a given probability, they will be matched to a participant that has chosen plan A or plan B in this specific supergame for this specific discount factor. If no such participant can be found, then the computer steps in. We also remind them that this procedure is the same for all participants in the lab, and that therefore, it is very well possible that the other participant will not be matched to them in return. See Appendix A.1 for an English translation of the instructions.

Pretests We conducted three pretest sessions with a total of 51 participants in Potsdam in December 2023. These pretests led to changes in the instructions and display on screen to improve comprehension. The data from the pretests will not be included in the analysis.

3 Hypotheses and Data Analysis

For each participant, we elicit their critical discount factor for each of the three possible stage-game parameter combinations. We define the critical discount factor as the discount factor for which a given participant first plays Grim. If a participant chooses Grim for 0.6 and everything above, their critical discount factor is 0.6. If they always choose Grim, their critical discount factor is 0.1. If they always choose AlwaysDefect, their critical discount factor is 1.

Based on our model in Andres et al. (2023), we can calculate point predictions for the critical discount factors in each of the three treatments. The first three rows in Table 3 present these predictions for the three different stage-game parameter combinations all participants face.

The fourth and fifth row in Table 3 show the differences in the critical discount factors between games. These differences are the point predictions for our two main outcome variables, i.e. the differences in the critical discount factors between the Base game and the HighGain game, and between the Base game and the HighLoss game, respectively:

$$\Delta\delta^*(p)|_{\Delta g} = \delta^*(p, g = 1, l = 0.2) - \delta^*(p, g = 0.2, l = 0.2) \quad (1)$$

$$\Delta\delta^*(p)|_{\Delta l} = \delta^*(p, g = 0.2, l = 1) - \delta^*(p, g = 0.2, l = 0.2) \quad (2)$$

To test whether the belief moderates the effect of the **gain** on the critical discount factor, we take the differences between the critical discount factors in the Base game and in the HighGain game and compare these differences across the LowBelief and the MedBelief treatment (see H_{Gain}). Similarly, for a test of the moderating effect of the **loss** on the critical discount factor, we take the differences between the critical discount factors in the Base game and in the HighLoss game and compare these differences across the HighBelief and the MedBelief treatment (see H_{Loss}).

Table 3: Point Predictions and Hypotheses

	LowBelief $p = 0.1$	MedBelief $p = 0.5$	HighBelief $p = 0.9$
(1) Base: $\delta^*(p, g = 0.2, l = 0.2)$	0.67	0.29	0.18
(2) HighGain: $\delta^*(p, g = 1, l = 0.2)$	0.74	0.55	0.51
(3) HighLoss: $\delta^*(p, g = 0.2, l = 1)$	0.90	0.55	0.24
(2) – (1) Effect of the gain: $\Delta\delta^*(p)\|_{\Delta g}$	**0.07**	**0.26**	0.32
(3) – (1) Effect of the loss: $\Delta\delta^*(p)\|_{\Delta l}$	0.24	**0.26**	**0.06**

H_{Gain}: The positive effect of increases in the gain on the critical discount factor increases from LowBelief to MedBelief.
H_{Loss}: The positive effect of increases in the loss on the critical discount factor decreases from MedBelief to HighBelief.

The experiment is calibrated such that we can use data from the Base game for the test of both hypotheses. Further, we are comparing symmetric changes in the stage-game parameters, which makes a comparison across the parameters more practical. Finally, we focus our analysis on the two comparisons where the theory predicts the largest differences, but we will report, analyze and discuss the remaining data, too. More specifically, we will also test whether there is a difference in $\Delta\delta^*(p)|_{\Delta g}$ between MedBelief and HighBelief, and whether there is a difference in $\Delta\delta^*(p)|_{\Delta l}$ between LowBelief and MedBelief.

Data Analysis We will employ one-sided t-tests for both of our main hypotheses, comparing the average differences in critical discount factors across treatments. Note that taking the differences in critical discount factors at the individual level as our outcome variables removes variation that is constant for a given participant. Note further that taking differences in the same way in all three treatments controls for any order effects that might arise because all participants face the Base game first.

For the t-tests, we will collapse the data to one observation per participant, resulting in a data set of length N. The hypotheses are directional and the respective data sets only partially overlap. Therefore, we do not correct for multiple hypothesis testing.

	H_0	H_1
H_{Gain}:	$\Delta\delta^*(p=0.1)\|_{\Delta g} \geq \Delta\delta^*(p=0.5)\|_{\Delta g}$	$\Delta\delta^*(p=0.1)\|_{\Delta g} < \Delta\delta^*(p=0.5)\|_{\Delta g}$
H_{Loss}:	$\Delta\delta^*(p=0.5)\|_{\Delta l} \leq \Delta\delta^*(p=0.9)\|_{\Delta l}$	$\Delta\delta^*(p=0.5)\|_{\Delta l} > \Delta\delta^*(p=0.9)\|_{\Delta l}$

For further analysis, we will employ linear regressions, estimated by OLS, with the two outcome variables, $\Delta\delta^*(p)|_{\Delta g}$ and $\Delta\delta^*(p)|_{\Delta l}$, as the dependent variables in two separate models. These models will include the belief in linear and squared form, thus explicitly allowing the effect of the belief on the effect of the loss to be non-monotonic, as predicted in Proposition 1b in Andres et al. (2023): If a player gets more optimistic that their opponent will cooperate, they will place less importance on the loss. However, if they are extremely pessimistic, they will not cooperate anyway, and thus also not place that much importance on the loss.

In these models, we plan to include data from the questionnaire. We will ask participants which potential strategy choice of the other person was more influential in their decision making process. As a manipulation check, we can test whether the treatments affect this importance variable. Further, as a test of the mechanism, we can use the importance variable to estimate the same model as above. We will also ask participants for a written explanation of their behavior, which we may use in an exploratory analysis using machine-learning methods to quantify the content of these statements. We are especially interested in whether players mention thoughts about their beliefs, and what conclusions they draw from having a high vs. a low belief.

4 Power Analysis

We base the power analysis on a Monte Carlo simulation with our theoretical model as the underlying data generating process. By sampling repeatedly, varying the effect sizes and the number of observations in the samples, we can compute the power as the share of samples that yield statistically significant results.

The hypothetical participants in the simulation are randomly assigned to treatments and then report their critical discount factors $\delta^*(p, g, l)$ based on the theory in Andres et al. (2023). As in the experiment, simulated participants can only report critical discount factors between 0.1 and 1, and changes in the critical discount factors between the three games are computed as described above. For each decision, it is randomly determined whether the true critical discount factor is reported or a random critical discount factor is taken instead. We take the probability of this as our main measurement of effect-size: the share of randomly reported critical discount factors. However, this can also be transformed approximately into Cohens' d (see below). In pseudo code, the simulation works like this:

1. For each combination of the error-rate and the number of subjects, repeat 10,000 times {

   i. Generate a random sample of size N and assign each individual to one of the three treatments LowBelief, MedBelief or HighBelief. For each individual in that sample, calculate their critical discount factor in each of the three games. This is a function of g, l and p.

   ii. For each decision, determine whether the true critical discount factor is reported or some random discount factor is reported. This is done with a given probability, which is the share of randomly reported critical discount factors. We call this the error rate.

   iii. Determine whether both t-tests reject the respective H_0 of H_{Gain} and H_{Loss} at 5% and store this in an indicator variable.

   }

2. Calculate the mean of the indicator variable for each combination of the error-rate and the number of subjects. This is the respective level of power.

See Figure 1 for the results of this simulation. Power decreases in the error-rate and increases in the number of observations. In Figure 1a, we display the results of the simulation described above, indicating the power for both t-tests of rejecting the null simultaneously. This means a random draw only counts as a success if both differences, for the gain and the loss, are statistically significant at 5%. Figure 1a also includes the approximate average Cohen's d on the second x-axis. In each sample of the simulation, we take the average over the effect-sizes for both comparisons and then take the average over all 10,000 samples. In Figure 1b, we can also see that the power for each individual hypothesis is higher, and converges to 5% as the error-rate approaches 1.

Figure 1: Power for H_{Gain} and H_{Loss}

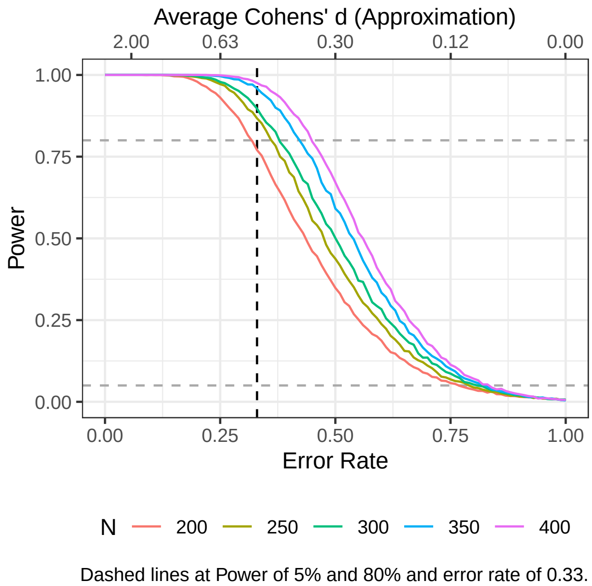

(a) Both t-tests simultaneously

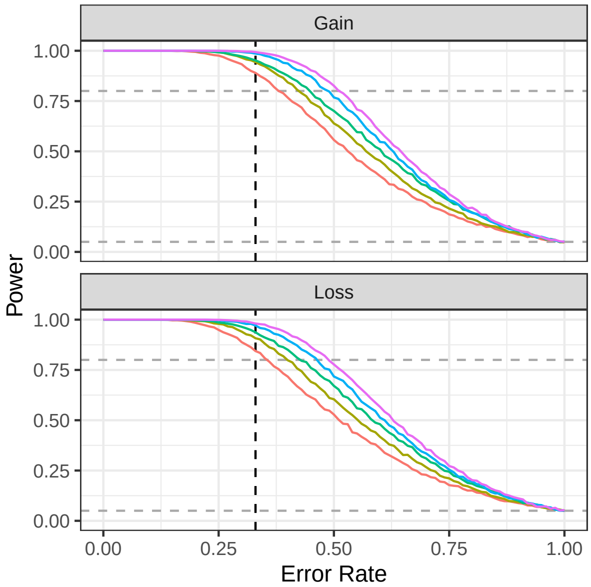

(b) Both t-tests separately

For an error rate of 0.33, we can see that $N = 250$ yields a power of well above 80%. An error rate of 0.33 corresponds to an approximate average effect size of $d \approx 0.5$ of both comparisons, or about $d \approx 0.51$ and $d \approx 0.48$ for the gain and the loss, respectively.

A Appendix

A.1 Instructions

Instructions

These instructions are identical for all participants.

Welcome to this experiment!

You will receive a monetary compensation for participating in this experiment. The amount you receive depends on your decisions and the decisions of other participants. It is therefore important that you read the instructions on the following pages carefully. Please take enough time to do this – the decision-making environment in this experiment requires a detailed explanation, which we will guide you through step by step.

For the entire duration of the experiment, you are not allowed to communicate with other participants. We therefore ask you not to talk to each other. Please also refrain from using your cell phones. Violation of these rules will result in exclusion from the experiment and payment.

If you have a question, please give us a hand signal. We will then come to you and answer your question personally.

During the experiment, we do not talk of euros, but of points. Your total income will initially be calculated in points. Your score will be converted into euros at the end of the experiment, using the following conversion rate:

18 points = 1 euro

At the end of today's experiment, you will receive the points you have scored from the experiment converted into euros in cash. In addition, you will receive 8 euros today for being on time for the experiment. The payment procedure is organized in such a way that the other participants will not see the amount you receive.

The experiment consists of two parts:

- In the first part, you can familiarize yourself with the experiment. This part has no influence on your payout. We call this part the test phase.
- In the second part, the actual experiment, you will make various decisions. These decisions will determine which payout you receive.

We will then ask you to complete a short questionnaire. You will then receive your payout in cash.

The decision situation

The decisions you will make in the experiment concern different versions of a specific decision situation. We would like to present these to you first.

You and another person make a decision between actions X and Y at the same time. The other person will later be assigned to you by the computer program. Neither you nor the other person will know anything about the identity of the other person.

Your payout depends on what you choose and what the other person chooses. The payouts are shown in the following table:

Your action	**Action of the other person**	**Your payout**	**Payout of the other person**
X	X	**75**	75
Y	X	**80**	45
X	Y	**45**	80
Y	Y	**50**	50

The payouts in the decision situations you face in the experiment will partially differ from those in this table. However, you and the other person will always receive the same payout if you both choose the same action, and different payouts if you choose different actions.

Many repetitions

When you and the other person make the decision described above in the experiment, you don't just do it once, but many times in succession. After each decision, you find out what the other person has done. The payouts that we have described above become smaller and smaller from round to round. This happens very evenly: After each round, all four possible payouts are multiplied by a number that we call the residual factor.

The residual factor is a number between 10% and 90%. The higher the residual factor, the more is left over after each round. The smaller the residual factor, the faster the payouts shrink. With a residual factor of 90%, the four payouts from the table above shrink in each round to 90% of the value in the previous round. With a residual factor of 50%, they are only half as large in each round as in the previous round.

On the next page we show you an overview of how this shrinkage looks for the value 50, depending on the residual factor. This is shown for the first 30 rounds. In the table you can see that the payouts become very small at some point.

	Runde (1-15)														
Rest-faktor	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**
0,9	50	45	40,5	36,45	32,81	29,52	26,57	23,91	21,52	19,37	17,43	15,69	14,12	12,71	11,44
0,8	50	40	32	25,6	20,48	16,38	13,11	10,49	8,39	6,71	5,37	4,29	3,44	2,75	2,2
0,7	50	35	24,5	17,15	12,01	8,4	5,88	4,12	2,88	2,02	1,41	0,99	0,69	0,48	0,34
0,6	50	30	18	10,8	6,48	3,89	2,33	1,4	0,84	0,5	0,3	0,18	0,11	0,07	0,04
0,5	50	25	12,5	6,25	3,13	1,56	0,78	0,39	0,2	0,1	0,05	0,02	0,01	0,01	0
0,4	50	20	8	3,2	1,28	0,51	0,2	0,08	0,03	0,01	0,01	0	0	0	0
0,3	50	15	4,5	1,35	0,41	0,12	0,04	0,01	0	0	0	0	0	0	0
0,2	50	10	2	0,4	0,08	0,02	0	0	0	0	0	0	0	0	0
0,1	50	5	0,5	0,05	0,01	0	0	0	0	0	0	0	0	0	0

	Runde (16-…)															
Rest-faktor	**16**	**17**	**18**	**19**	**20**	**21**	**22**	**23**	**24**	**25**	**26**	**27**	**28**	**29**	**30**	**…**
0,9	20,59	18,53	16,68	15,01	13,51	12,16	10,94	9,85	8,86	7,98	7,18	6,46	5,81	5,23	4,71	…
0,8	3,52	2,81	2,25	1,8	1,44	1,15	0,92	0,74	0,59	0,47	0,38	0,3	0,24	0,19	0,15	…
0,7	0,47	0,33	0,23	0,16	0,11	0,08	0,06	0,04	0,03	0,02	0,01	0,01	0,01	0	0	…
0,6	0,05	0,03	0,02	0,01	0,01	0	0	0	0	0	0	0	0	0	0	…
0,5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	…
0,4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	…
0,3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	…
0,2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	…
0,1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	…

[In the instructions, this table filled an entire page in landscape.]

Deciding on plans, not actions

The decisions between action X and action Y in the many repetitions are not all made individually in the experiment. Instead, at the beginning of a series of repetitions, you choose a plan once, as you would like to decide in the many rounds. You have two plans to choose from: Plan A and Plan B.

Plan A stipulates that you choose action X in the first round. If the other person also chooses action X in the first round, you stick with action X in the subsequent rounds. If the other person chooses action Y in the first round, you choose action Y from the second round onwards.

Plan B is that you always choose action Y, regardless of what the other person does. You start with action Y and choose action Y in every round.

As you and the other person both choose between plan A and plan B, there are only four ways in which the decisions can be made in the many repetitions:

Your plan	**Plan of the other person**	**Progression of the actions in the many repetitions**
A	A	In the first round and in all subsequent rounds, you both select action X.
B	A	In the first round, you choose action Y and the other person chooses action X. From the second round onwards, you both choose action Y for all subsequent rounds.
A	B	In the first round, you choose action X and the other person chooses action Y. From the second round onwards, you both choose action Y for all subsequent rounds.
A	B	In the first round and in all subsequent rounds, you both select action Y.

In the experiment, you never decide between actions X and Y in individual rounds, but only ever decide between plan A and plan B for the entire duration of all repetitions. The computer program uses this to recognize how you and the other person's decisions will play out and calculates your payout directly from this.

To do this, the computer program assumes mathematically that there are an infinite number of repetitions of the decision situation. As you can see in the table, however, the payouts per round shrink to very small values after relatively few rounds, so that a large part of your total payout is determined in the first repetitions.

You make the decision between plan A and plan B several times in the experiment – for a total of three different decision situations, which differ in the the payouts, and in each case for the residual factors 90%, 80%, 70%, 60%, 50%, 40%, 30%, 20% and 10%.

For each of the remaining factors, you specify whether you choose plan A or plan B. You may only switch between the plans once per decision situation. This means: If, for example, you want to choose plan A for the residual factor 40% and plan B for the residual factor 30%, you must also choose plan A for the residual factors 90% to 50%; and you must also choose plan B for the residual factors 20% to 10%. You can also select the same plan for all residual factors.

Here we show you a screenshot of the first decision situation on the computer.

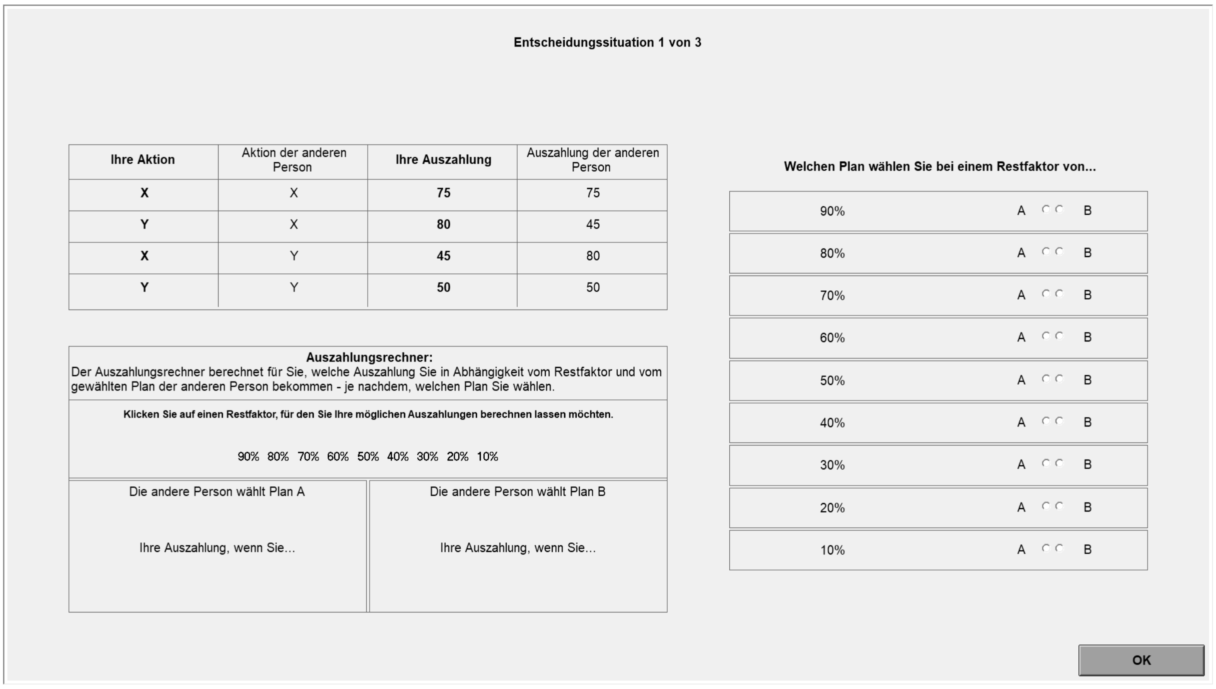

In the experiment, you will have a payout calculator at your disposal that will show you how high the total payout will be over all rounds – depending on the residual factor and which plan you and the other person choose.

Payout of a single decision

In the experiment, you make the decision between plan A and plan B for three different versions of the decision situation. For each of these three versions, you decide for each of the nine different residual factors between the two plans. This means that you make a total of $3 \cdot 9 = 27$ decisions in the experiment. The payout you receive for your participation in the experiment is determined by exactly one of these 27 decisions. Which one this is is decided by a random mechanism of the computer, which gives all 27 decisions the same probability of being drawn as a decision relevant to the payout. Your payout corresponds to your points from this decision converted into euros. You will receive this payout in addition to the 8 euros you receive for showing up on time for the experiment.

This means that each of your decisions can be the one that determines your payout! You will only be told which one at the end of the experiment.

What do you learn about the other person?

The other person, whose decision together with your own decision determines your payout, will only be assigned to you after you have made your payout-relevant decision. You initially make all your decisions between Plan A and Plan B on your own. The allocation is based on the decisions made by the other people.

At the beginning of the experiment, you are told how likely it is that you will be assigned to a person who has chosen plan A in the decision relevant to your payout, or a person who has chosen plan B. This probability is then always the same for all decisions and does not change.

You therefore know from the outset for all decisions with what probability the other person who is ultimately assigned to you for the calculation of your payout will have chosen plan A or plan B. This probability does not depend on your own decision.

The payouts for the other people are calculated in exactly the same way. This means that for all other people, too, one of the 27 decisions is initially selected by a random computer mechanism. After the decision, each of the other people – just like for you – is assigned a second person whose decision for plan A or plan B determines the first person's payout. It is therefore not unlikely that the person who is assigned to you for the calculation of your payout will not be assigned to you for the calculation of this person's payout, but to another person from today's experiment.

During this procedure, it may happen that there is no person who has opted for the plan to be assigned to you. In such a case, the computer takes over the role of the other person and selects the corresponding plan so that your payout is also calculated normally.

Test phase

At the beginning of the experiment, all participants can familiarize themselves with the display of the decision on the computer screen. In the test phase, you will be shown the screen for the first of the three versions of the decision situation.

In the test phase, quiz questions are asked on the screen. The experiment only begins when all participants have answered all the quiz questions correctly. Your answers to the quiz questions have no consequences for your payout at the end of the experiment!

If you have a question, please give us a hand signal. We will then come to you and answer your question personally.

References

Andres, M. (2023). Communication in the infinitely repeated prisoner's dilemma: Theory and experiments. *Working Paper, Available at arXiv 2304.12297*.

Andres, M., L. Bruttel, and J. Nithammer (2023). Payoffs, beliefs, and cooperation in infinitely repeated games. *Working Paper, Available at SSRN 4491762*.

Aoyagi, M., G. R. Fréchette, and S. Yuksel (2023). Beliefs in repeated games: An experiment. *Working Paper, Available at SSRN 3784841*.

Costa-Gomes, M. A., S. Huck, and G. Weizsäcker (2014). Beliefs and actions in the trust game: Creating instrumental variables to estimate the causal effect. *Games and Economic Behavior 88*, 298–309.

Stahl II, D. O. (1991). The graph of prisoners' dilemma supergame payoffs as a function of the discount factor. *Games and Economic Behavior 3*(3), 368–384.