

Sustaining Cooperation with Correlated Information: An Experimental Test*

Yongping Bao
University of Bremen
yongping.bao@uni-bremen.de

Fabian Dvorak
University of Konstanz
fabian.dvorak@uni-konstanz.de

Sebastian Fehrler[†]
University of Bremen
sebastian.fehrler@uni-bremen.de

December 2024

Abstract

In infinitely or indefinitely repeated games with noisy signals about others' actions, sustaining cooperation is difficult. Theoretical work shows that cooperation can be maintained if the signals are correlated and the degree of correlation depends on the actions. In this study, we implement such an information structure in a laboratory experiment and investigate whether subjects are able to sustain cooperation by conditioning their behavior on it. A substantial number of subjects adopt strategies accounting for the correlation, but this does not increase cooperation compared to a control treatment without correlation, as behavior with independent signals is more lenient.

Keywords: infinitely repeated games, monitoring, cooperation, strategic uncertainty, prisoners' dilemma, correlated signals.

JEL: C72, C73, C92, D83.

*The study obtained ethical approval by the GfEW (No. C6C1g6Kj). A pre-analysis plan of the experiment was registered at the AEA registry as AEARCTR-0005369. Fabian Dvorak gratefully acknowledges funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2117 - 422037984. Sebastian Fehrler gratefully acknowledges funding from the Juniorprofessorenprogramm Baden-Württemberg and from the Federal Ministry of Labour and Social Affairs (Germany) through the funding network Interdisciplinary Social Policy Research (FIS).

[†]corresponding author

1 Introduction

In a tightly controlled laboratory experiment, we study strategy choices and resulting cooperation rates in an indefinitely repeated Prisoners' Dilemma with imperfect public monitoring and correlated signals, where the degree of correlation depends on the actions of both players (as in Awaya and Krishna, 2016, 2019).

A large literature in game theory investigates cooperation in infinitely (or indefinitely) repeated games. In some cases, players monitor the actions of the opponents or partners perfectly. In many situations, however, opponents' actions are monitored only imperfectly. For example, in teamwork, players repeatedly put effort into a joint project. Their partners do not observe their actual effort, but instead the output which noisily represents the effort (Sekiguchi, 1997; Compte and Postlewaite, 2015). In the oligopoly model of Green and Porter (1984), firms observe the market price but not the output choices of their competitors. In both examples, shocks or other factors may confound the observed information, such that instead of directly observing the actual choice of the opponent, players only observe noisy signals that represent the opponents' actions.

Understanding how cooperation can be sustained in indefinitely repeated games under imperfect monitoring has become an important topic of the literature, because sustaining cooperation is difficult when information is noisy. Awaya and Krishna (2016, 2019) point out that one way to sustain cooperation is through correlated information. While the literature commonly assumes that noisy signals solely depend on players' own chosen actions, they assume that the noisy signals depend on all actions. Under their assumption, the signals are correlated in a way that they are similar when players' actions are similar, and become dissimilar, as players' actions diverge. One example is a duopoly. The two firms have private information regarding their own sales, but their sales are correlated as they are sensitive to the prices on the market, and the degree of correlation depends on whether they chose similar actions.

The fundamental insight in Awaya and Krishna (2016, 2019) is that exploiting correlation in information can help to improve monitoring and sustain cooperation. In this study, we adopt and adapt their assumption of information correlation and explore the empirical implications. While their focus is on private monitoring, our design uses public monitoring.¹ They show

¹Private monitoring, as in Sekiguchi (1997); Compte and Postlewaite (2015), means that noisy signals are only privately observable. The secret price cutting of oligopolies described by Stigler (1964) is a classic example.

that communication prior to every round can sustain cooperation based on pure strategies. A necessary condition is that the private signals are more strongly correlated if players choose the same action. By publicly reporting private signals, players can draw inferences about past actions based on these reports. Under public monitoring, signals are publicly observable and there is no role for truthful communication of signals. In this sense, our implementation represents the case of fully truthful communication under private monitoring.

Our central focus is thus not on the communication strategies players devise when they know how to exploit information contained in correlation. In contrast, we focus on the question whether people make use of correlated information for sustaining cooperation given that everyone reports truthfully. In other words, our design does not directly answer the research question of Awaya and Krishna (2016, 2019), that is what effect does the exchange of private information have on cooperation when signals are correlated. Instead, we experimentally test the underpinnings of their theory by addressing the question to what extent subjects can exploit the information about others' actions contained in the correlation of signals.²

For this purpose, we implement an indefinitely repeated Prisoners' Dilemma of imperfect public monitoring, where payoffs depend on the own action and the signal regarding the action of the other player. Signals are publicly observable and noisy, that is they falsely represent the real action of the opponent with an exogenously given and fixed probability. We distinguish between two settings. In one experimental treatment we implement a correlation structure which can be exploited to support full cooperation as a subgame perfect equilibrium. Signals are systematically and perfectly correlated, that is the two public signals are the same, if both players choose the same action. If their actions differ, signals are independently drawn. We compare subjects' decisions in this treatment to a control treatment with independent signals using the same stage game parameters and continuation probability. The control treatment without correlation has a similar design as Fudenberg et al. (2012), Aoyagi et al. (2019) and Dvorak and Fehrler (2024).³

Intuitively, if publicly observable signals are perfectly correlated, that is identical when actions are the same and independently drawn otherwise, players can look for differences in signals to

²Findings from other studies of different contexts suggest that subjects may struggle to make correct inferences from correlated information (e.g., Enke and Zimmermann, 2017; Fehrler et al., 2024).

³Other recent experimental literature on repeated games includes Dal Bó and Fréchette (2019), Romero and Rosokha (2023), Aoyagi et al. (2024), Backhaus and Breitmoser (2024), and Bland (2024). For a survey of earlier contributions see Dal Bó and Fréchette (2018).

maintain cooperation. In such cases, simple cooperative strategies can be used that condition on the correlation of the signals. For instance, cooperation can be enforced by a grim-trigger strategy which triggers if signals diverge. If both players start with cooperation, both will receive the same signal, and will continue cooperating as long as the two signals match. Once the signals become different, they infer from the dissimilarity of the signals that different actions have been chosen in this period. The cooperating player therefore detects the defection of the other player with certainty and defects in all future periods. As this strategy is similar to the grim-trigger strategy under perfect monitoring except that it takes the correlation of information into account, we refer to this strategy as the correlated (signals) grim-trigger strategy (CGRIM). We also consider other correlation-based strategies such as correlated tit-for-tat (CTFT), correlated win-stay-lose-shift and a correlated trigger strategy with two periods of punishment (CT2). In the experiment, we use a combination of stage-game parameters and continuation probability for which the perfect correlation structure described above allows full cooperation based on CGRIM, while at the same time no subgame perfect cooperative equilibria exist in the treatment without signal correlation.

We implement pre-play communication before each interaction in a form of free chat. Communication facilitates cooperation (e.g., Cooper et al. 1992; Rabin 1994; Ellingsen and Östling 2010). Pre-play communication, in particular, helps players to coordinate on cooperative equilibria, and raises cooperation rates by reducing strategic uncertainty (Kartal and Müller, 2024). In a previous study of a noisy, indefinitely repeated Prisoners' Dilemmas with uncorrelated signals, Dvorak and Fehrler (2024) observe high cooperation rates in the first round with pre-play communication followed by a steady and substantial decline in the subsequent rounds. Communication is not a treatment variable in our design, but it gives us a higher chance of observing cooperative strategies, and allows for the comparison of the decline of cooperation rates between treatments. We further classify the communication content and analyze its impact on the evolution of strategy choices.

Our main finding is that more than half of all subjects understand the value of the correlation structure for maintaining cooperation and use the correlation-based strategies when signals are perfectly correlated. The majority of these subjects use the CGRIM strategy, while a few also use the CTFT strategy. However, we do not find that correlation structure of signals promotes cooperation, because (i) CGRIM is neither lenient nor forgiving, which triggers many

punishments, and (ii) subjects' play is in general very lenient, and especially in the absence of correlation, where defection cannot be detected with certainty.

The rest of the paper is organized as follows. In the next section, we introduce the stage-game, and derive theoretical predictions for both treatments. In Section 3, we describe the experimental design and clarify the reasons for our choice of stage-game parameters. Section 4 summarizes research questions and methodology. We present the empirical results in Section 5. The key findings are summarized and discussed in Section 6.

2 Repeated Prisoners' Dilemma with Imperfect Public Monitoring

Consider an indefinitely repeated Prisoners' Dilemma with two players who repeatedly play a 2×2 stage game. The game terminates with a constant exogenously given probability of $0 < 1 - \delta < 1$ after every round. In each round, the two players $i = \{1, 2\}$ choose from an action set $A_i = \{C, D\}$ and their action is denoted by a_i .

There is ex-post uncertainty regarding the actual choice of an opponent. Under public monitoring (Green and Porter, 1984), instead of observing the actual choice, each player's action is translated into a noisy signal and both signals are publicly announced to players. More technically, let $\Omega_i \in \{c, d\}$ be a set of signals of player i 's action, where signal c represents action C , and d represents D . We denote the realized signal from this set by ω_i . For each action profile $a = (a_i, a_{-i})$, a conditional probability distribution $\pi(\omega|a)$ is assigned over signals. In our experiment, the two treatments differ with respect to π , and therefore in the quality of signals.

Imperfect public monitoring with uncorrelated signals (*NoCor*) For the treatment without correlation, signals are conditionally independent of $a = (a_1, a_2)$. That is, signals are drawn independently for each of the chosen actions and signal ω_i depends only on a_i but not on a_{-i} . Signals are noisy and indicate the opposite of the chosen action with probability $\epsilon = 0.2$, that is: when the player chooses C (D), the signal indicates C (D) with probability $1 - \epsilon = 0.8$, and indicates D (C) with probability 0.2. The conditional probability distribution of ω_i is thus $\pi(\omega_i = c|a_{-i} = C) = \pi(\omega_i = d|a_{-i} = D) = 1 - \epsilon$ and $\pi(\omega_i = c|a_{-i} = D) = \pi(\omega_i = d|a_{-i} = C) = \epsilon$.

Imperfect public monitoring with correlated signals (*Cor*) In this setup, both players receive identical signals if they choose the same action. Yet the signals still indicate the opposite of the chosen action with probability 0.2. The conditional probability distribution of ω_i is given by $\pi(\omega_i = c|a_i = a_{-i} = C) = \pi(\omega_i = d|a_i = a_{-i} = D) = 1 - \epsilon$ and $\pi(\omega_i = c|a_i = a_{-i} = D) = \pi(\omega_i = d|a_i = a_{-i} = C) = \epsilon$ when both players choose the same action. If their actions differ, signals are drawn independently, as in *NoCor*. Therefore, if the two signals differ, it is certain that different actions were chosen. If the signals are the same, players do not know for sure the opponent's action.

Player i 's stage game payoff g_i is defined by i 's action a_i and the noisy signal ω_{-i} reflecting the other player's action, that is $g_i : A_i \times \Omega_{-i} \rightarrow \mathbb{R}$. Hence player i cannot infer the action of the other player from her payoff. The expected stage game payoff of player i is given by $u_i : A \rightarrow \mathbb{R}$, $u_i(a) = \sum_{\omega_{-i} \in \Omega_{-i}} g_i(a_i, \omega_{-i}) \pi(\omega_{-i}|a)$.

The expected stage-game payoff profile (u_i, u_{-i}) which we consider has the form of a Prisoners' Dilemma:

Table 1: The prisoner's dilemma stage game

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

Expected payoffs in Table 1 are normalized. Parameters g and l are both positive and $g < 1+l$.

Denote by $\Omega^t = \{\Omega_i, \Omega_{-i}\}^t$ the set of public histories up to round t . A public strategy for player i is a mapping $\sigma_i : \bigcup_{t \geq 0} \Omega^t \rightarrow \Delta A_i$. A strategy profile $\sigma = (\sigma_i, \sigma_{-i})$ is a perfect public equilibrium (PPE) if σ_i is a public strategy and for any history up to round t , σ is a Nash equilibrium in all rounds following t , in other words PPE is a subgame perfect Nash Equilibrium (SPE) that depends only on public signals.

The equilibrium solution for *NoCor* Mutual cooperation can be sustained by non-lenient and non-forgiving grim-trigger strategies. Players start with C but deviate to D for all subsequent rounds if the private history $(a_i, \omega_i, \omega_{-i})^t \neq (C, c, c)^t$. We refer to mutual cooperation as the reward state, and to mutual defection as the punishment state. This strategy is a PPE if the long-run incentive of cooperation $\frac{1}{1-\delta(1-\epsilon)^2}$ is as least as big as the gain from defection $\frac{1+g}{1-\delta\epsilon(1-\epsilon)}$, that is δ should be sufficiently large:

$$\delta \geq \delta_{NoCor}^{PPE} = \frac{g}{(1+g)(1-\epsilon)^2 + \epsilon^2 - \epsilon} \quad (1)$$

The equilibrium solution for *Cor* Mutual cooperation can still be enforced by grim-trigger strategies. The class of cooperative SPE strategies defined in the above paragraph still exists. In addition, correlated grim-trigger strategies exist which can be used to achieve mutual cooperation in all rounds. Specifically, we define the correlated grim-trigger strategy (CGRIM) as follows: the player starts with C and will continue choosing C as long as the two public signals match, that is $\omega_i = \omega_{-i}$. Otherwise the player chooses D in all subsequent rounds. This strategy can be used to construct a symmetric subgame perfect equilibrium if the long-term gains of cooperation are at least as big as the short-term incentive to defect. The threshold δ^{PPE} is given by

$$\delta \geq \delta_{Cor}^{PPE} = \frac{g}{1 - 2\epsilon + 2\epsilon^2 + g} \quad (2)$$

See Appendix A for a proof.

Some remarks are in order. It is easier to sustain cooperation under *Cor* than *NoCor* in terms of grim-trigger strategies. For $\epsilon \in (0, 1)$ and g fixed, whenever (1) holds, condition (2) holds as well. The correlation structure of signals changes the equilibrium condition in two aspects. First, it increases the value in the continuation path of the reward state. Even when both players cooperate, bad signals occur with positive probability in *NoCor* which triggers the punishment state. Under *NoCor*, the continuation value on the cooperative path is the expected utility of reward state and punishment state, which will be reached with positive probability. With correlation, however, signals always match if both cooperate, which means punishment state will never be reached, and therefore the continuation value is larger.

Second, note that *Cor* makes unilateral defection look more attractive as the probability of staying in reward state after unilateral defection is higher than that in *NoCor*.⁴ When signals differ, a cooperating player can know for sure that the opponent has defected. However, when signals match, it is impossible to make such inference. Even when players choose different actions, signals still match with a probability of $2\epsilon(1 - \epsilon)$. This is the probability of remaining in the reward state after unilateral defection, which corresponds to the probability of the signal profile being either $(\omega_i, \omega_{-i}) = (c, c)$ or (d, d) , that is, $2\epsilon(1 - \epsilon)$. The *NoCor* treatment renders this probability to be $\epsilon(1 - \epsilon)$, that is, cooperation in the next round is only possible when (c, c) is observed. Despite that *NoCor* triggers punishment with a higher probability when one player defects unilaterally, it does not compensate for the reduced expected utility in the reward state due to punishments on the equilibrium path.

3 Experimental Design

We implement a noisy Prisoners' Dilemma game with public monitoring in a laboratory experiment.⁵ The experiment has two between-subject treatments. The two treatments vary in the conditional distribution of signals:

NoCor Signals are public and independent.

Cor Signals are public and perfectly correlated if both actions are the same, and independent otherwise.

In every round, two players choose their actions $a_i \in \{C, D\}$ simultaneously. Payoffs depend on the player's own action a_i and the received signal about the other player's action $\omega_{-i} \in \{c, d\}$. Signals are noisy and indicate the wrong action with probability $\epsilon = 0.2$, which do not vary between treatments.

The continuation probability δ of the repeated game is 0.8. The stage-game payoff matrix is given in the upper panel of Figure 1. The payoffs are in experimental currency units. The lower

⁴This only applies to the absolute probabilities to stay in the reward state. The difference in the probability to stay in reward state when choosing defection instead of cooperation is larger in *Cor* ($1 - 0.32 = 0.68$) than in *NoCor* ($0.64 - 0.16 = 0.48$). Therefore, unilateral defection is more risky in *Cor* compared to *NoCor*. Since the reward state is also more valuable in *Cor*, unilateral defection is less reasonable in *Cor*.

⁵The treatments have been pre-registered on the AEA RCT Registry (Bao et al., 2020, AEARCTR-0005369). We also present our pre-analysis plan in Appendix F.

panel of Figure 1 shows the expected stage-game payoffs for action profiles. The normalized expected stage-game parameters are $g = l = 0.8$.

Figure 1: Stage-Game Parameters and Predictors of Cooperation

	c	d
C	32	2
D	40	10

	C	D
C	26, 26	8, 34
D	34, 8	16, 16

Notes: Payoffs are in experimental currency units with an exchange rate of 50 ECU = 1 EUR. Both matrices are shown to subjects during their decision-making.

Under such parameterization, the existence thresholds of PPE are $\delta_{NoCor}^{PPE} = 0.81$ and $\delta_{Cor}^{PPE} = 0.54$. These parameters make sure that the condition outlined in Equation (2) holds but the one defined in Equation (1) does not, that is there exists an equilibrium in which both players play correlated grim-trigger strategy when signals are correlated, and no cooperative PPE exists if there is no correlation.

In all treatments, subjects engage in a pre-play communication-stage before the first round of every supergame. A supergame is one indefinitely repeated interaction. Communication is possible via a chat-box interface for 120 seconds.

In every session of both treatments, subjects are randomly assigned to 3 matching groups with 8 participants in each group. Subjects only interact with the members in the same matching group throughout the entire session. Thus, one matching group can be treated as one independent observation. Subjects play 7 supergames of pre-determined lengths. At the beginning of every supergame, subjects are matched with a new partner from their matching group using perfect stranger matching such that they do not play with the same partner from the matching group for a second time. To keep the length of supergames constant across treatments, we generate 3 sequences of random numbers beforehand, and use them to determine the length L_i of each

supergame.⁶ To increase the number of observations per supergame, we adapt the block-random-termination method (Fréchette and Yuksel, 2017). Subjects play a block of five rounds at the beginning of every supergame. If the true length L_i is smaller or equal than 5, the supergame ends at the end of round 5 and only the first L_i rounds are payoff relevant. If L_i is larger than 5, the supergame continues until round L_i has been reached and all rounds are payoff relevant. Before the end of round 5, subjects are not informed about whether the supergame ends or not. The modified block-random-termination method allows us to collect data of at least five rounds in every supergame.

At the end of each of the first four rounds within every supergame, subjects receive feedback on $\{a_i, \omega_i, \omega_{-i}\}$ plus the stage profit g_i . Starting from round 5, in addition to $\{a_i, \omega_i, \omega_{-i}\}$ and g_i , they are informed of the pre-generated random number of that round. The supergame does not end until that number is smaller or equal than $1 - \delta$, which is 0.2 under our parameterization.

Before the game begins, subjects are provided with a detailed explanation of the experimental procedure and must answer control questions to self-check their understanding of the process (see Appendix E). At the end of the experiment, subjects answer a short survey that elicits basic socio-economic characteristics, such as age and gender.

We conducted six sessions at Lakelab of University Konstanz, with three sessions for *Cor* and *NoCor* respectively. A total of 144 subjects participated in our experiment.⁷ Subjects were students of the University of Konstanz.⁸ Each experimental session lasted approximately 2 hours, and the average earning was 18.67 Euros.

⁶We use Stata to generate 3 sequences of uniformly distributed random numbers between 0 and 1 with seeds 3, 4, and 5 (Seeds 1 and 2 have been used in: Dvorak and Fehrler, 2024). Denote the 3 sequences as $\{r_n\}_i = \{r_1, r_2, \dots, r_x\}_i$, where $i = 3, 4, 5$ indicates the seed underlying the sequence and $n \in \mathbb{N}$. The first supergame has x_1 rounds if $r_{x_1} \leq 0.2$ and for all $n < x_1$, $r_n > 0.2$. The second supergame has $x_2 - x_1$ rounds if $r_{x_2} \leq 0.2$ and for all $x_1 < n < x_2$, $r_n > 0.2$. And so forth. The resulting (lengths of the) sequences are SQ1 (2, 8, 1, 5, 7, 1, 7), SQ2 (4, 2, 2, 21, 4, 3, 5) and SQ3 (2, 3, 1, 1, 4, 6, 6).

⁷After data collection, we found out that one subject participated twice in two *NoCor* sessions. We exclude the matching group of this subject's second participation from the data analysis, using the data of the remaining 136 subjects (mean age = 22 years, 54.41% female).

⁸The experiment was programmed in z-Tree (Fischbacher, 2007) and subjects were recruited via hroot (Bock et al., 2014).

4 Research Questions & Methods

We address three research questions.

Question 1: *Do strategy choices differ when cooperation can be sustained based on correlated signals?*

We start with an analysis of the cooperation rates after memory-one histories to shed light on treatment differences in leniency and forgiveness. In our setup, a memory-one history consists of a player’s own action and the two public signals $\{a_i, \omega_{-i}, \omega_i\}$. We have a vector of nine possible memory-one states ($\emptyset, ccc, ccd, cdc, cdd, dcc, dcd, ddc, ddd$). The first element \emptyset is the initial round with no history of play. The rest of the elements are nonempty histories representing $\{a_i, \omega_{-i}, \omega_i\}$. For instance, cdc describes a state of a player who chose cooperation, receives signal d and sends out signal c to the other player. We look at the cooperation probabilities after each of the possible histories and we represent the cooperation probabilities by a vector $(\sigma_\emptyset, \sigma_{ccc}, \sigma_{ccd}, \sigma_{cdc}, \sigma_{cdd}, \sigma_{dcc}, \sigma_{dcd}, \sigma_{ddc}, \sigma_{ddd})$. We focus on nine states to allow behavior to be conditioned on the combination of action and public signals. This allows us to investigate whether the cooperation rates are in line with correlation-based strategies. Unjustifiable defection is estimated by $1 - \sigma_{ccc}, 1 - \sigma_{ccd}$ without signal correlation, and by $1 - \sigma_{ccc}, 1 - \sigma_{cdd}$ with signal correlation. $\sigma_{cdc}, \sigma_{cdd}$ are estimations for leniency when signals are uncorrelated. When signals are correlated, estimations for leniency are σ_{ccd} and σ_{cdc} . For both uncorrelated and correlated signals, $\sigma_{dcc}, \sigma_{dcd}, \sigma_{ddc}$ and σ_{ddd} are estimations for forgiveness.

To analyze strategy choices, we use the R package `stratEst` (Dvorak, 2022), which can be used to implement the SFEM of Dal Bó and Fréchette (2011) based on the EM algorithm (Dempster et al., 1977). We restrict our attention to the first five rounds of the last 3 supergames when subjects have gained experience of play. In the main text, we report models which include the six pure strategies studied by Dal Bó and Fréchette (2011) and their corresponding correlation-based variants as our candidate strategies. As this set of strategies was originally developed for the case of perfect monitoring, we adapt these pure strategies so that they depend on own actions and public signals. Descriptions of the candidate strategies are summarized in Table 2. As a robustness check, we extend the candidate set to include the 20 strategies analyzed in Fudenberg et al. (2012) for the imperfect public monitoring case. A description of each strategy and its automaton representation can be found in the tables of Appendix B.

Table 2: Overview of the Candidate Strategies

Strategy	Acronym	Description
Always Defect	ALLD	Always play D .
Always Cooperate	ALLC	Always play C .
Grim	GRIM	Play C until the signal representing the partner's action is d , then play D forever.
Tit-for-Tat	TFT	Play C unless the signal representing the partner's action is d in the last round.
Win-Stay-Lose-Shift	WSLS	Play C if own choice is the same as the signal representing the partner's choice in the last round, otherwise play D .
T2	T2	Play C until either signal is d , then play D twice and return to C (regardless of all actions and signals during the punishment rounds).
Correlated Grim	CGRIM	Play C until the public signals do not match, then play D forever.
Correlated Tit-for-Tat	CTFT	Play C unless the public signals do not match in the last round.
Correlated Win-Stay-Lose-Shift	CWSLS	Play C if the public signals match, otherwise play D .
Correlated T2	CT2	Play C until the signals do not match, then play D twice and return to C (regardless of all actions and signals during the punishment rounds).

Question 2: *Does correlation in signals increase cooperation?*

According to the theoretical prediction outlined in Section 2, there should be more cooperation in the *Cor* treatment. In line with the theory, our pre-registered main hypothesis is the following:

H1: *The average cooperation rate will be higher in Cor than in NoCor.*

H1 follows directly from the critical continuation probabilities defined in Equations (1) and (2), and our stage-game parameters ($g = l = \delta = 0.8$). We test this hypothesis by comparing the average cooperation rates in the first five rounds of the last three supergames between treatments. We further analyze the decline of cooperation rates to compare the stability of cooperation.

Question 3: *How does communication affect cooperation?*

To approach this question, we had two research assistants classify the subjects' communication content into 39 sub-categories, which we then grouped into four main categories. The tables of Appendix D illustrate the categories. We expect that subjects who coordinate in the *Cor* treatment, explicitly or implicitly talk about punishments for the case when the two public signals differ. Our major focus is on the main categorical level, which we take as explanatory variables and analyze through logistic regressions whether subjects who talk about certain topics are more likely to play cooperatively.

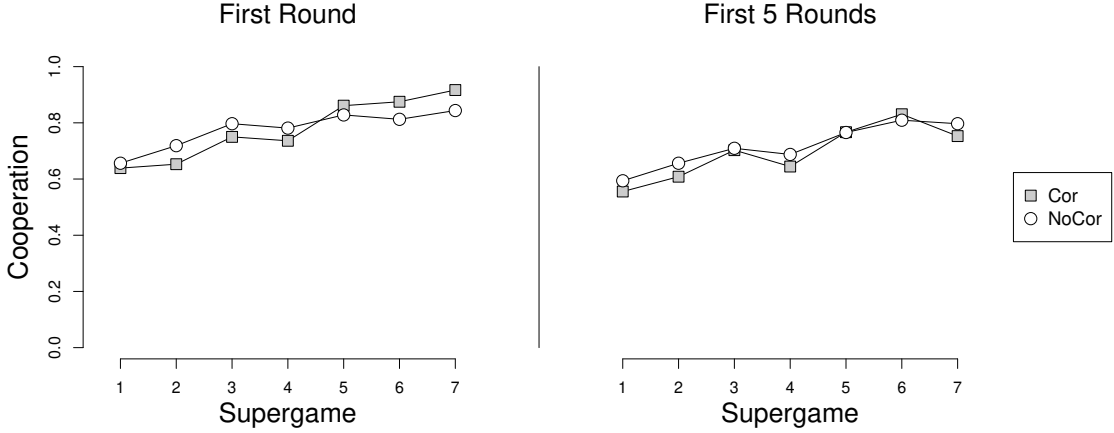
5 Experimental Results

In line with our pre-analysis plan, we focus on the last three supergames for cooperation rates and strategy frequency analysis, when subjects have gained experience with the game. Figure 2 depicts the average cooperation frequency over all seven supergames. The average cooperation rates increase in both treatments as subjects gain more experience. In addition, since each supergame consists of at least five rounds, our main focus is on the first five rounds in the supergames such that the transition from the block to random termination does not confound the result.

5.1 Strategy Choices

We estimate the probability of cooperation after different memory-one histories. For both treatments, cooperation rates condition on $\{a_i, \omega_{-i}, \omega_i\}$, which has nine possible memory-one histories. Table 3 reports the cooperation rates following each history. σ_\emptyset indicates the probability of cooperating in the first round. It is high in both treatments, meaning that subjects are generally cooperative in the initial round. The unjustified defection ($1 - \sigma_{ccc}$, $1 - \sigma_{cdd}$ in *Cor* and $1 - \sigma_{ccc}$, $1 - \sigma_{ccd}$ in *NoCor*) rates are low in both treatments. The largest treatment difference is observed within the *cdc* state. Subjects in *NoCor* are substantially more cooperative than in *Cor*. Since σ_{cdc} is an estimation of leniency in both treatments, a higher level indicates that subjects are less likely to defect when a bad signal occurs. State *ccd* and *cdd* are also estimations of leniency in *Cor* and *NoCor* respectively. Treatment differences between these

Figure 2: Evolution of Cooperation over Supergames



Notes: The lines depict the first round and average frequency of cooperation over seven supergames for both treatments.

states are small. We do not observe much difference in states *dcc* and *dcd*. The cooperation rates in these states are relatively low in both treatments. In terms of forgiveness (*dcc*, *dcd*, *ddc* and *ddd*), the willingness to return to cooperation after defection is low in both *Cor* and *NoCor*.

The probabilities of ending in these states are unequal for all strategies. We therefore cannot interpret the cooperation probabilities after memory-one histories as strategy choices and explain whether subjects play a certain strategy or not. For instance, σ_{ccd} and σ_{cdc} are very different in the correlation treatment. Subjects are willing to cooperate when they receive a cooperative signal from their partners but tend to defect when the signal regarding the partner's choice is defective. If subjects play CGRIM, σ_{ccd} and σ_{cdc} should equalize. But we still cannot conclude that subjects in *Cor* do not play CGRIM due to the above reason. To find out what strategies they play, our next step is to investigate the heterogeneity of strategy choices. We are interested in the specific strategies subjects employ, which contribute to our observation in Table 3 and explain why subjects in *NoCor* are substantially more lenient on average following *cdc*. All candidate strategies condition on $\{a_i, \omega_{-i}, \omega_i\}$.

We fit a strategy estimation model on the candidate strategy set and select the subset of strategies that explains subjects' choices best using the Bayesian Information Criterion (BIC) for model comparison. Table 4 shows the estimated maximum-likelihood shares of the selected

Table 3: Cooperation Rates After Memory-One Histories

	σ_{\emptyset}	σ_{ccc}	σ_{ccd}	σ_{cdc}	σ_{cdd}	σ_{dcc}	σ_{dcd}	σ_{ddc}	σ_{ddd}	$\ln L$
<i>Cor</i>	0.88 (0.03)	0.95 (0.01)	0.83 (0.22)	0.33 (0.06)	0.71 (0.04)	0.41 (0.08)	0.39 (0.09)	0.17 (0.21)	0.35 (0.07)	-409.28
<i>NoCor</i>	0.83 (0.04)	0.91 (0.02)	0.95 (0.03)	0.79 (0.04)	0.86 (0.07)	0.22 (0.13)	0.46 (0.10)	0.11 (0.09)	0.32 (0.09)	-386.73

Notes: This table summarizes the response probabilities of cooperation following the 9 possible histories: \emptyset , *ccc*, *ccd*, *cdc*, *cdd*, *dcc*, *dcd*, *ddc*, *ddd*. Bootstrapped standard errors are from 10000 iterations and are presented in parentheses. The log likelihood of the model is summarized in the last column.

strategies for both treatments. We test whether the strategy shares differ between treatments with a likelihood-ratio test, based on the assumption that the strategy shares are identical in both treatments. The p -value of this test is < 0.001 , which suggests that the strategy shares differ between treatments. In Table 4, ALLC attracts a substantial share in *NoCor*, whereas the non-lenient and non-forgiving correlation-based variant of the grim-trigger strategy, CGRIM attracts the highest share in *Cor*. This estimation result explains and confirms the low leniency of subjects choices in *Cor* following state *cdc*. The behavior of a substantial share of subjects in *Cor* is consistent with playing CGRIM. They defect immediately when signals mismatch and never return to cooperation.

Table 4: SFEM with Strategy Selection

	ALLD	ALLC	TFT	WSLS	CGRIM	CTFT	γ	BIC	$\ln L$
<i>Cor</i>	0.06 (0.03)	0.26 (0.08)	0.06 (0.04)	0.07 (0.05)	0.40 (0.10)	0.16 (0.08)	0.10	856.59	-415.47
<i>NoCor</i>	0.12 (0.04)	0.76 (0.06)	0.07 (0.04)	-	0.05 (0.03)	-	0.10	737.02	-360.19

Notes: The table reports the maximum-likelihood shares of the strategies of Dal Bó and Fréchet (2011) and their correlation-variants with data from the first 5 rounds of the last 3 supergames. All strategies condition on action-public signal profile $\{a_i, \omega_{-i}, \omega_i\}$. The estimation procedure assumes constant strategy use. γ is the estimated tremble probability, which avoids likelihood shares of zero when subjects deviate from a choice pattern. Strategies are selected based on Bayesian Information Criterion. Strategies attracting zero shares are omitted (-). Standard errors are reported in parentheses. Values may not add up to one because of rounding.

To check the robustness of the SFEM result, we expand the candidate strategy set to include another 14 pure strategies. We additionally include CGRIM, CTFT, CWLS and CT2 into the candidate set as alternative correlation-based strategies. Table C1 in Appendix D presents the results. In general, ALLC remains the most popular strategy in *NoCor*, while in *Cor*, although a lenient variant of the grim-trigger strategy (GRIM2) is now the most prevalent strategy, CGRIM and CTFT attract significant shares.

Result 1: *Strategy choices differ when signals are correlated. Subjects move from lenient strategies used in the absence of correlation to not-lenient and unforgiving correlation-based strategies with correlation.*


5.2 Cooperation

Figure 3 shows the average cooperation rate in both treatments. Bars indicate the means, and error bars indicate two-way clustered standard errors of the mean (Cameron et al., 2011). The two clusters we use for the standard errors are participants and matches. Figure 3 indicates that the average cooperation rate does not differ between treatments. The average cooperation rate is 0.78 in *Cor*, 0.79 in *NoCor*. The two-sided z-test comparing the frequency of cooperation in both treatments is not significant ($z = -0.15, p = 0.88$).

Figure 4 shows the average frequency of cooperation over the first five rounds. The cooperation rates start high (above 78 percent in all supergames, and above 82 percent in the last three supergames) and decline over rounds. In the last three supergames, the average cooperation rate has reduced by 19 percentage points in *Cor* and 9 percentage points in *NoCor*, which indicates more stability in *NoCor*. We regress cooperation on the round number, treatment dummies and their interaction term to test whether the difference in stability is statistically significant using two-way clustered standard errors. The result shows that the decline of cooperation is steeper in *Cor* than in *NoCor* ($p = 0.02$).

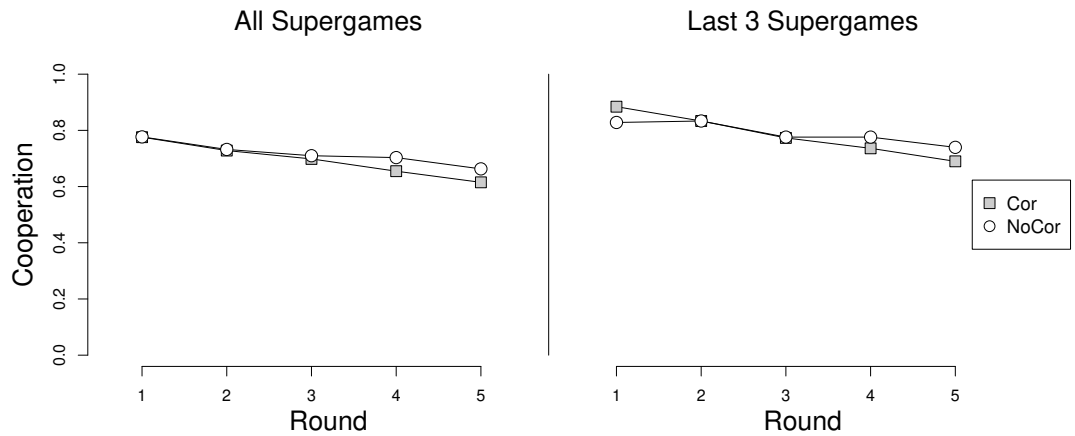
Result 2: *Signal correlation does not increase the cooperation rate. The decline in cooperation is slightly steeper in the correlation treatment than in the no correlation treatment.*

Figure 3: Average Cooperation Rate Between Treatments



Notes: Bars show the average frequency of cooperation. Error bars represent two-way clustered standard errors for mean cooperation rate on subject and match level.

Figure 4: Stability of Cooperation over Rounds



Notes: The lines show the average frequency of cooperation in the first 5 rounds of all and the last 3 supergames.

5.3 Communication

Next, we take a closer look at the pre-play communication content and explore whether and in what way the average cooperation rate and strategy choices are influenced by the opportunity to communicate. Two research assistants classify subject-round communication observations into 39 sub-categories, with 5.67 classifications on average for each subject-round observation and 2313 classifications in total for the last three supergames. We merge the 39 sub-categories into four main categories: *Coordination*, *Deliberation*, *Relationship* and *Trivia*.⁹ We determine Cohen’s κ on the main categorical level.¹⁰ The average κ is 0.35, indicating a fair level of agreement between the raters. *Coordination* is the effort to coordinate on future behavior. It includes explicit and implicit punishment threats when signals differ. According to Table 5, for both treatments, the category *Coordination* has been covered in almost all pre-play chats in the last three supergames. Observations concerning the discussion of actions and strategies are classified into the *Deliberation* category. The *Relationship* category contains all chats with promises, expressions of trust and distrust, and requests for trustworthy behavior. This category is the least frequent among the four. Another frequently observed category is *Trivia*. This category covers small talk, off-topic talks, and expressions of boredom.

Table 5: Frequency of Valid Codings per Individual-Round Observation

	<i>Cor</i>	<i>NoCor</i>
<i>Coordination</i>	0.99	0.99
<i>Deliberation</i>	0.37	0.26
<i>Relationship</i>	0.15	0.03
<i>Trivia</i>	0.89	0.96

Notes: Frequency of valid codings on main categorical level of subject-round observations. Data is from the last three supergames. A classification is considered to be valid if the classifications from the two raters agree. It is possible that an observation belongs to more than one category, resulting in both column sums to be larger than 1.

⁹Detailed description of the sub-categories and the mapping from sub-categories to main categories are in Table D1 in Appendix D. We use the same classification scheme as Dvorak and Fehrer (2024) except two differences. First, to account for signal correlation, we add sub-categories regarding matching of signals. Second, because communication in our sessions has a pre-play structure, we eliminate their category “Information” and its sub-categories which regard the share of information in repeated communication.

¹⁰When the rating is random, agreement occurs with probability $p_i^2(Yes) + p_i^2(No)$, $i = 1, 2$, where $p_i(Yes)$ is the frequency of rater i classifying objects into any categories, and $p_i(No) = 1 - p_i(Yes)$.

Looking closely at sub-categorical levels (Table D1 in Appendix D), shows that most subjects propose the mutual play of C in both treatments. A small fraction of subjects in Cor talk about non-lenient GRIM punishment, while subjects in $NoCor$ do not attempt to agree on a punishment plan. Interestingly, although subjects play CGRIM, topics like the matching of the signals, or implicit and explicit punishment when signals differ, are rarely discussed. The frequency of their occurrence is below 0.1%.

To answer the question how communication is related to cooperation, we conduct logistic regressions. The marginal effects are presented in Table 6, with the dependent variable being the first-round cooperation, and the explanatory variables being the dummies whether or not the pre-play communication falls into a main category. We control for the supergame indexes and socio-demographic and other subject-related characteristics. Data include all supergames, because the cooperation in round one is so high in the last three supergames that little variation exists. Standard errors are bootstrapped with 1000 repetitions and are two-way clustered on subjects and matches (Cameron et al., 2011). In the treatment Cor , *Deliberation* is positively correlated to the first-round cooperation. In treatment $NoCor$, a positive correlation is found for the category *Trivial*.

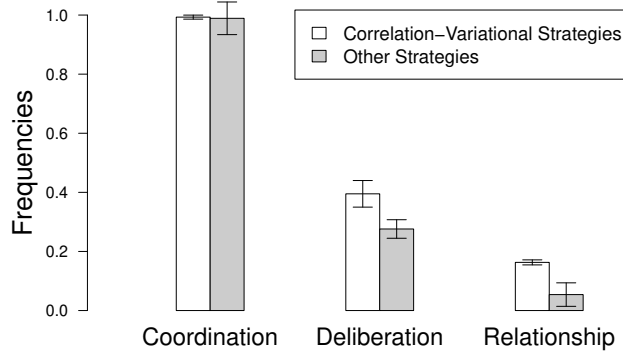
Table 6: Communication and First Round Cooperation

	<i>Cor</i>	<i>NoCor</i>
<i>Coordination</i>	0.96	1.38
<i>Deliberation</i>	0.65**	0.09
<i>Relationship</i>	0.34	-0.21
<i>Trivial</i>	0.25	0.71*
Supergame	0.30***	0.11

Notes: The table presents the result of Logistic regression: the marginal effects of being in the main categories on first period cooperation. Data include all supergames. We control for socio-demographic and subject-related characteristics. Bootstrapped standard errors are calculated with two-way clustering, with cluster dimensions being subjects and match (1000 repetitions). Significance is based on the critical values of one-sided t statistics. *** (**,*) indicates significance on the 1 (5,10)% level.

We further assign subjects to strategies on the basis of posterior probabilities, that is we map each individual to a strategy that has the maximum posterior probability based on SFEM.

Figure 5: Frequency of Communication of Correlation-Variants and Other Strategies



Notes: Bars show the frequencies of communication of subjects who use correlation-based strategies (CGRIM or CTFT) and other strategies. Individuals are assigned strategies with highest posterior probabilities. Data is from the last three supergames. Error bars represent two-way clustered standard errors at the subject and match levels.

Among the 136 subjects, 49 use correlation-based strategies (CGRIM or CTFT), 87 use non-correlation related strategies (ALLD, ALLC, TFT and WSLS). We look at the frequencies of main categorical communication regarding correlation-based strategies users and users of other strategies. Figure 5 presents the result. Subjects who use correlation-related strategies have a similar communication pattern as the other subjects. There is only a small difference in *Deliberation* and *Relationship* ($p < 0.001$ for both categories). Subjects who are assigned as users of CGRIM or CTFT deliberate more, and are involved in talks to build up relationship more often.

Result 3: *Subjects use pre-play communication to coordinate behavior. Subjects who engage in deliberation are more cooperative in the treatment with signal correlation, where deliberation is associated with the use of correlation-based strategies.*

6 Conclusion

Sustaining cooperation is difficult in indefinitely repeated games when players cannot perfectly observe their opponents' actions. One way to sustain cooperation found in the literature is

through correlated information (Awaya and Krishna, 2016, 2019). In a laboratory experiment, we investigate the effect of correlated information on strategy choice and cooperation in an indefinitely repeated Prisoner’s Dilemma with public monitoring.

In the experiment, information is correlated in a simple way that is easy for subjects to understand. When players choose the same action, they receive identical signals about the actions of both players, which are the result of the same noisy process. If their actions are different, the two signals are independently determined by the same noisy process. With correlation, there is a simple grim-trigger strategy that can be used to efficiently support cooperation in equilibrium. This grim-trigger strategy cooperates as long as the two public signals exchanged in each round of the supergame match, and fails otherwise. Without correlation, such a strategy does not exist and cooperation cannot be supported in equilibrium. The theory therefore predicts more cooperation when signals are correlated. However, other studies have shown that people struggle to make correct inferences based on correlated information (e.g., Enke and Zimmermann, 2017; Fehrler et al., 2024). Therefore, the hypothesis that correlation can promote cooperation by improving the quality of imperfect monitoring was an open empirical question.

We find no treatment difference in the average frequency of cooperation between treatments. Cooperation starts high in *Cor* and has a steeper decline. An analysis of participants’ strategies shows that a significant fraction of participants in the *Cor* treatment take the correlation of signals into account when making their decisions. However, they do so by following a non-lenient punishment scheme for defection. Since not all participants in the *Cor* treatment are cooperative, punishments occur frequently. This reveals two opposing roles of correlation. On the one hand, correlation makes it possible to detect defection with certainty, which makes efficient cooperation possible. On the other hand, defection becomes unambiguously detectable, and the response to defection is not lenient because there is no wiggle room around bad signals.

Note that the discount factor we implemented in the lab is very close to the threshold above which cooperative equilibria exist in the absence of correlation, which may explain why participants are so cooperative in the *NoCor* treatment. However, it is still unclear why participants are so lenient and forgiving in the *NoCor* treatment. Dvorak and Fehrler (2024) also find high leniency under private monitoring with pre-play communication. This relates to the broader question of why participants are lenient towards bad signals under imperfect monitoring with pre-play communication. One possible explanation is that pre-play communication allows them to connect with each other and social preferences become more important. Future research could

further investigate the effects of pre-play communication on cooperation and strategy choice in our setup.

References

- Aoyagi, M., Bhaskar, V., and Fréchette, G. R. (2019). The impact of monitoring in infinitely repeated games: Perfect, public, and private. *American Economic Journal: Microeconomics*, 11(1):1–43.
- Aoyagi, M., Fréchette, G. R., and Yuksel, S. (2024). Beliefs in repeated games: An experiment. *American Economic Review*, 114(12):3944–75.
- Awaya, Y. and Krishna, V. (2016). On communication and collusion. *American Economic Review*, 106(2):285–315.
- Awaya, Y. and Krishna, V. (2019). Communication and cooperation in repeated games. *Theoretical Economics*, 14(2):513–553.
- Backhaus, T. and Breitmoser, Y. (2024). Individual Strategy Choice in the Repeated Prisoner’s Dilemma . *Discussion Paper Series – CRC TR 224*, 341.
- Bao, Y., Dvorak, F., and Fehrler, S. (2020). Exchanging private information to sustain cooperation in noisy, indefinitely repeated interactions. *AEA RCT Registry*, 5369-1.0.
- Biernacki, C., Celeux, G., and Govaert, G. (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22:719–725.
- Bland, J. (2024). Meta-study in the shadow of the future: the effect of continuation probability on strategies. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.5042608>.
- Bock, O., Baetge, I., and Nicklisch, A. (2014). hroot: Hamburg registration and organization online tool. *European Economic Review*, 71:117 – 120.

- Breitmoser, Y. (2015). Cooperation, but no reciprocity: Individual strategies in the repeated prisoner’s dilemma. *American Economic Review*, 105(9):2882–2910.
- Cameron, a. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249.
- Compte, O. and Postlewaite, A. (2015). Plausible cooperation. *Games and Economic Behavior*, 91:45–59.
- Cooper, R. W., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in coordination games. *Quarterly Journal of Economics*, 107(2):739–771.
- Dal Bó, P. and Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1):411–429.
- Dal Bó, P. and Fréchette, G. R. (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, 56(1):60–114.
- Dal Bó, P. and Fréchette, G. R. (2019). Strategy choice in the infinitely repeated prisoner’s dilemma. *American Economic Review*, 109(11):3929–3952.
- Dempster, A., Laird, N., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B*, 39(1):1–38.
- Dvorak, F. (2022). *stratEst: Strategy Estimation*. R package version 1.1.6.
- Dvorak, F. (2023). *stratest: a software package for strategy frequency estimation*. *Journal of the Economic Science Association*, 9:337–349.
- Dvorak, F. and Fehrler, S. (2024). Negotiating cooperation under uncertainty: Communication in noisy, indefinitely repeated interactions. *American Economic Journal: Microeconomics*, 16(3):232–58.
- Ellingsen, T. and Östling, R. (2010). When does communication improve coordination? *American Economic Review*, 100(4):1695–1724.
- Enke, B. and Zimmermann, F. (2017). Correlation neglect in belief formation. *The Review of Economic Studies*, 86(1):313–332.

- Fehrler, S., Hochleitner, A., and Janas, M. (2024). Too good to be true: Individual and collective decision-making with misleading information. *Mimeo*.
- Fischbacher, U. (2007). Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.
- Fréchette, G. R. and Yuksel, S. (2017). Infinitely repeated games in the laboratory: four perspectives on discounting and random termination. *Experimental Economics*, 20(2):279–308.
- Fudenberg, D., Rand, D. G., and Dreber, A. (2012). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review*, 102(2):720–749.
- Green, E. J. and Porter, R. H. (1984). Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87.
- Kartal, M. and Müller, W. (2024). A new approach to the analysis of cooperation under the shadow of the future: Theory and experimental evidence. Available at SSRN: <https://ssrn.com/abstract=3222964>.
- Rabin, M. (1994). A model of pre-game communication. *Journal of Economic Theory*, 63(2):370–391.
- Romero, J. and Rosokha, Y. (2023). Mixed Strategies in the Indefinitely Repeated Prisoner’s Dilemma. *Econometrica*, 91(6):2295–2331.
- Sekiguchi, T. (1997). Efficiency in repeated prisoner’s dilemma with private monitoring. *Journal of Economic Theory*, 76:345–361.
- Stigler, G. J. (1964). A theory of oligopoly. *Journal of Political Economy*, 72(1):44–61.

A PPE with Correlated Grim-trigger Strategies

We construct the PPE δ -threshold of a grim-trigger strategy under *Cor*. Players start with C and never deviate to D as long as public signals match. To show that cooperation is optimal, we show that the gain from cooperation is at least as high as the gain from defection.

Denote by V_g the value in reward state where both players opt for C , and V_d the value in punishment state where both deviate to D . Denote further by $u(C)$ and $u(D)$ the expected payoffs from playing C or D . Cooperation is optimal if the following inequality holds

$$u(C) + \delta V_g \geq u(D) + \delta (2\epsilon(1 - \epsilon)V_g + (1 - 2\epsilon(1 - \epsilon))V_d) \quad (3)$$

The biggest difference between correlated and non-correlated grim-trigger strategy is that if both cooperate, signals always match under *Cor* and players stick to cooperative path. We represent the non-expected stage-game cooperation, betrayal, sucker and defection payoffs by c, b, s, d . $u(C)$ and $u(D)$ is thus explicitly given by

$$\begin{aligned} u(C) &= (1 - \epsilon)c + \epsilon s \\ u(D) &= (1 - \epsilon)b + \epsilon d \end{aligned}$$

The continuation value depends on own continuation strategy and has the following recursive form

$$\begin{aligned} V_g &= (1 - \epsilon)c + \epsilon s + \delta V_g \\ V_d &= (1 - \epsilon)d + \epsilon b + \delta V_d \end{aligned}$$

Solving for V_g , V_d and plugging them and $u(C)$ and $u(D)$ back into (3) yields

$$(1 - \epsilon)(c - b) + \epsilon(s - d) + \frac{\delta}{1 - \delta}(1 - 2\epsilon + 2\epsilon^2)((1 - \epsilon)(c - d) + \epsilon(s - b)) \geq 0 \quad (4)$$

We can solve for δ and rewrite the expression by replacing c, b, s, d with expected payoffs and normalize them into a function of g :

$$\delta \geq \frac{g}{1 - 2\epsilon + 2\epsilon^2 + g}.$$

Correlated grim-trigger strategy is a perfect public equilibrium if the above inequality holds.

B Overview of Strategies

Tables B1-B3 list the 20 strategies taken from Fudenberg et al. (2012), and are reprinted from Appendix B of Dvorak and Fehrler (2024). The 20 strategies together with the four correlation-based strategies CGRIM, CTFT, CWSLS and CT2 are used to produce Table C1. Circles are strategy states and arrows indicate state transitions. Tables B1-B3 summarize the basic scenario of only five states. In our strategy frequency estimation, transition is assumed to condition on own action and both signals, thus nine states.

Table B1: Strategies 1-7

Acronym	Description	Automaton
ALLD	Always play D.	
ALLC	Always play C.	
DC	Start with D, then alternate between C and D.	
FC	Play C in the first round, then D forever.	
Grim	Play C until either player plays D, then play D forever.	
TFT	Play C unless partner played D last round.	
PTFT (WSLS)	Play C if both players chose the same move last round, otherwise play D.	

Notes: Circles represent the states of an automaton. The first state from the left is the start state. The labels *C* and *D* indicate whether the automaton prescribes cooperation or defection in the state. Arrows represent deterministic state transitions. The labels indicate the information profiles of the previous periods which trigger the transitions. An unlabeled arrow indicates an unconditional transition that occurs independent of the observed profile.

Table B2: Strategies 8-15

Acronym	Description	Automaton
T2	Play C until either player plays D, then play D twice and return to C (regardless of all actions during the punishment rounds).	
TF2T	Play C unless partner played D in both of the last 2 rounds.	
TF3T	Play C unless partner played D in all of the last 3 rounds.	
T2FT	Play C unless partner played D in either of the last 2 rounds (2 rounds of punishment if partner plays D).	
T2F2T	Play C unless partner played 2 consecutive Ds in the last 3 rounds (2 rounds of punishment if partner plays D twice in a row).	
GRIM2	Play C until 2 consecutive rounds occur in which either player played D, then play D forever.	
GRIM3	Play C until 3 consecutive rounds occur in which either player played D, then play D forever.	
PT2FT	Play C if both players played C in the last 2 rounds, both players played D in the last 2 rounds, or both players played D 2 rounds ago and C last round. Otherwise play D.	

Notes: Circles represent the states of an automaton. The first state from the left is the start state. The labels *C* and *D* indicate whether the automaton prescribes cooperation or defection in the state. Arrows represent deterministic state transitions. The labels indicate the information profiles of the previous periods which trigger the transitions. An unlabeled arrow indicates an unconditional transition that occurs independent of the observed profile.

Table B3: Strategies 16-20

Acronym	Description	Automaton
DTFT	Play D in the first round, then play TFT.	
DTF2T	Play D in the first round, then play TF2T.	
DTF3T	Play D in the first round, then play TF3T.	
DGRIM2	Play D in the first round, then play GRIM2.	
DGRIM3	Play D in the first round, then play GRIM3.	

Notes: Circles represent the states of an automaton. The first state from the left is the start state. The labels C and D indicate whether the automaton prescribes cooperation or defection in the state. Arrows represent deterministic state transitions. The labels indicate the information profiles of the previous periods which trigger the transitions. An unlabeled arrow indicates an unconditional transition that occurs independent of the observed profile.

C Additional Results

Table C1: SFEM with Fudenberg et al. (2012) and Strategy Selection

	<i>Cor</i>	<i>NoCor</i>
ALLD	0.03 (0.02)	0.10 (0.04)
ALLC	-	0.52 (0.09)
DC	0.03 (0.02)	0.04 (0.03)
FC	0.02 (0.02)	-
TFT	0.05 (0.03)	-
WSLS	0.04 (0.03)	-
TF2T	-	0.14 (0.08)
GRIM2	0.31 (0.09)	0.16 (0.07)
GRIM3	0.20 (0.08)	-
DTFT	0.02 (0.02)	-
CGRIM	0.21 (0.08)	0.04 (0.03)
CTFT	0.11 (0.06)	-
γ	0.08	0.08
BIC	805.56	678.89
$\ln L$	-381.39	-326.97

Notes: The table reports the maximum-likelihood shares of the strategies selected from Tables B1-B3 of Appendix B with data from the first 5 rounds of the last 3 supergames. All strategies condition on action-public signal profile $\{a_i, \omega_{-i}, \omega_i\}$. The estimation procedure assumes constant strategy use. γ is the estimated tremble probability, which avoids likelihood shares of zero when subjects deviate from a choice pattern. Strategies are selected based on Bayesian Information Criterion. Strategies attracting zero shares are omitted (-). Standard errors are reported in parentheses. Values may not add up to 1 because of rounding.

D Communication Content

Table D1: Frequency of Valid Coding for Sub-Categories

#	Subcategory	Category	Freq.	Freq. in Treatment		$\bar{\kappa}$
				<i>Cor</i>	<i>NoCor</i>	
1	Proposal: both C	C	0.936	0.926	0.948	0.614
2	Proposal: both D	C	0.015	0.019	0.010	0.447
3	Proposal: alternate	C	0.078	0.111	0.042	0.756
4	Proposal: self D other C	C	0.054	0.074	0.031	0.872
5	Proposal: self C other D	C	0.025	0.046	-	0.762
6	Proposal: other coordination	C	0.025	0.037	0.010	0.479
7	Question: action of the other	C	0.005	0.009	-	0.164
8	Announcement: C	C	0.010	0.009	0.010	0.138
9	Announcement: D	C	0.005	0.009	-	0.328
10	Rejection of proposal	C	0.020	0.028	0.010	0.357
11	Acceptance proposal	C	0.711	0.639	0.792	0.368
12	Implicit punishment threat for D	C	-	-	-	-
13	Punishment threat grim	C	0.015	0.028	-	0.356
14	Punishment threat lenient grim	C	-	-	-	-
15	Approval of punishment threat	C	-	-	-	-
16	Ask for coordination	C	0.123	0.139	0.104	0.480
17	Benefits of C	D	0.142	0.157	0.125	0.310
18	Benefits of D	D	-	-	-	-
19	Benefits of asymmetric play	D	0.005	0.009	-	0.159
20	Related to fairness discussion	D	-	-	-	-
21	Related to strategic uncertainty	D	-	-	-	-
22	Related to payoffs	D	0.029	0.046	0.010	0.078
23	Related to Prisoner's dilemma	D	0.010	0.009	0.010	1.000
24	Related to game theory	D	0.005	0.009	-	1.000
25	Future benefit of C	D	-	-	-	-
26	Short term incentives of D	D	0.010	0.019	-	0.385
27	Related to signal comparison	D	0.005	0.009	-	0.187
28	Related to (un)matched signals	D	-	-	-	-
29	Promise	R	0.039	0.065	0.010	0.484
30	Distrust	R	0.025	0.046	-	0.289
31	Trust	R	0.005	-	0.010	0.235
32	Plea for trustworthy behavior	R	0.025	0.037	0.010	0.198
33	Implicit punishment threat when signals differ	C	-	-	-	-
34	Explicit punishment threat when signals differ grim	C	-	-	-	-
35	Explicit punishment threat when signals differ lenient grim	C	-	-	-	-
36	Small Talk	T	0.882	0.824	0.948	0.407
37	Off topic	T	-	-	-	-
38	Boredom	T	0.010	0.019	-	0.136
39	Confusion	D	0.010	0.019	-	0.035

Notes: Two raters identify whether a sub-category occurs in a subject-round observation. Frequency indicates the probability of occurrence of a sub-category whose classifications are the same between raters. The table shows overall frequency and frequency in treatments. Data is from last three supergames. The 39 sub-categories map into 4 main categories: *Coordination* (C), *Deliberation* (D), *Relationship* (R) and *Trivia* (T). Frequencies < 0.001 omitted (-). $\bar{\kappa}$ is the average Cohen's Kappa over all treatments. Mean $\bar{\kappa}$ of all subcategories with an overall frequency > 0.01 is 0.38.

E Experimental Instructions and Quiz

[Below is the instructions and quiz for the information correlation treatment. Instructions for the other treatment is very similar and therefore is omitted. The original instructions are in German. Instructions for both treatments can be obtained from the authors upon request.]

Overview

Welcome to this experiment. We ask you not to talk to the other participants during this experiment and to switch off your mobile devices.

At the end of the experiment, You will be paid in cash for today's participation. The amount of money you receive depends on your own decisions, the other participants' decisions, and pure chance. It is important that you understand the instructions before the experiment starts.

In this experiment, every interaction between the participants runs through the computers you are sitting in front of. They will interact with each other anonymously. Neither your name nor the names of other participants will be announced. Also, for the evaluations only the anonymized data are used.

Today's session consists of several rounds. Your payout amount will be the sum of the points earned in all rounds, converted into euros. The conversion of the points into euros is done as follows. Each point is worth 2 cents, so that applies: 50 Points = 1.00 EUR.

All participants are paid privately, so other participants cannot see how much you have earned.

Experiment

Interactions and Role Assignment

This experiment consists of 7 interactions that are identical in their sequence, each consisting of a randomly determined number of rounds.

At the very beginning, before the first interaction, you will be randomly placed in a group with other participants. In each of the 7 interactions you will interact with another participant of your group.

Specifically, this is what happens: Before the first interaction, you will be assigned to a person from your group with whom you will interact in all rounds of the first interaction. In the second interaction, you will then be assigned to a new person from your group with whom you will interact in all rounds of the second interaction, and so on. In this way, you will interact with each person assigned to you in only one interaction, but in all rounds of that interaction.

Length of an Interaction

The length of an interaction is determined randomly. After each round there is an 80% chance that there will be at least one more payout-relevant round.

You can imagine this as follows. After each round, a 100-sided dice is thrown. If the roll results in a number less than or equal to 20, there is no further payout relevant round. If the roll is a different number (21-100), the interaction continues. Note that the probability of another payout-relevant round does not depend on the round you are in. The probability of a third payout-relevant round if you are in round 2 is 80%, as is the probability of a tenth payout-relevant round if you are in round 9.

A special feature concerns the first 5 rounds of each interaction. These rounds are always run even if the interaction has already been completed by the random number generator. At the end of the fifth round, you will find out whether the interaction has already been completed and, if so, up to which round your decisions were relevant for the payout. If the interaction has not been completed by round five, it will continue round after round and the interaction will be ended immediately if there is no further round.

When an interaction is finished, a new person is assigned for the next interaction. After the seventh interaction, the experiment ends.

Interaction and Round Schedule

At the beginning of an interaction, that is before the first round of the interaction, you can chat with the other person on the screen. The chat takes place in an anonymous chat window. To protect your anonymity, it is important that you do not give any information about yourself or your seat number while communicating. Otherwise we reserve the right not to pay you in the end. The chat content will be displayed during the interaction and you can read it.

Then the first round of interaction begins.

In each round you choose one of two possible options, A or B. The other person also chooses one of two possible options, A or B.

For each option, a signal is randomly determined in each round, which corresponds to the option with 80% probability. With 20% probability the signal does not correspond to the option but shows the other option. At the end of a round, you and the other person do not learn what the other person has chosen, but receive the signals determined for the chosen options. Your signal corresponds to the signal of the option chosen by the other person. The other person's signal corresponds to the signal of the option you have chosen. That is, if both persons choose the same option, both receive the same signal. If the two people have chosen different options, both can get different signals. This results in the following.

Important: Since exactly one signal is randomly determined for each option in each round, it may be possible to draw conclusions based on the two signals as to what the other person has chosen.

If two different signals occur, the other person has certainly chosen a different option than you. If you and the other person had chosen the same option, you and the other person would also receive the same signal (the signal determined for the option you both chose).

If two identical signals occur, the other person has either chosen the same option (case 1) or another option (case 2) and the two signals for the different options correspond randomly (probability for case 2: $0.8 * 0.2 + 0.2 * 0.8 = 0.32$).

Your income depends on the option you choose and the signal you receive. Similarly, the payout of the other person depends on the option you choose and the signal you receive.

Figure E1: Round Income [Figure 1 from Instructions]

Ihre Optionen	Ihr Einkommen bei Signal		Erwartetes Einkommen, wenn die andere Person	
	A	B	Option A wählt	Option B wählt
Option A	32	2	26	8
Option B	40	10	34	16

In Figure 1, the four fields on the left indicate the lap income resulting from the combinations of the chosen option and your signal. The same table applies to the other person. For example, your lap income is 10 points if you chose option B and received signal B, and the other person's income is 2 points if he or she chose option A and received signal B.


Once you and the other person have chosen an option, chance determines the signals for the options with the probabilities given above. The signals for the chosen options are then used to determine your lap earnings and those of the other person.

The four fields on the right in Figure 1 show the earnings you can expect depending on your option and the option of the other person. For example, if you choose option B and the other person chooses option A, you are 80% likely to receive signal A and 20% likely to receive signal B. Therefore, you will receive 40 points with 80% probability and 10 points with 20% probability, which means that your expected earnings in this case are: $0.8 * 40 + 0.2 * 10 = 34$ points.

At the end of the round you will receive a short feedback regarding your chosen option, the signal you received, the signal the other person received and your own round earnings (see Figure 2). You will not be informed of the other person's choice of option.

All possible subsequent rounds are identical in terms of the sequence of events. However, you can only chat with the other person before the first round of interaction. This is not possible before later rounds. The progress of the current interaction, that is the feedback that you received at the end of previous rounds, is displayed in a tabular view.

Figure E2: Part of Feedback Screen (Example) [Figure 2 from Instructions]



Rundeneinkommen	
Ihre Wahl:	Option B
Ihr Signal:	B
Signal der anderen Person:	B
Ihre Punkte aus dieser Runde:	10

End and Final Payoff

As soon as chance ends the last interaction, the experiment is over.

At the end of the experiment all interactions are paid off. The total amount of points from all rounds will be converted into Euros and paid out privately.

On the last screen of the last round of the last interaction, you can see how much you have earned in Euros.

Any Question?

If you have any questions, please contact us. An experimenter will then come to your place.

If you think you have understood everything well, you may start the quiz on the screen. This quiz is only to make sure that everyone has understood the instructions well. The answers to the quiz will not affect your payoff.

Quiz [on screen]

[After completing the quiz, correct answers will appear on the next screen.]

1. How many interactions are there?

[1,7, it is by chance]

2. What is the minimum number of rounds in an interaction (payment-relevant or not)?

[1, 3, 5]

3. What is the probability that there will be another payout-relevant round of interaction when you are in round six of an interaction?

[20%, 80%, 100%]

4. What is the probability that the signal corresponds to the actual option?

[20%, 80%, 100%]

5. Suppose you choose option A, the other person receives the signal A.

(a) You receive the signal B, which option did the other person choose?

[certain option A, certain option B, cannot be stated with certainty]

(b) You also get the signal A, which option did the other person choose?

[certain option A, certain option B, cannot be stated with certainty]

(c) If the other person chose option B, how likely was it for you to receive signal A?

[20%, 80%, 100%]

6. You choose option B. Suppose the other person also chooses option B.

(a) With what probability will the other person receive signal B?

[20%, 80%, 100%]

(b) You receive signal A, which signal does the other person receive?

[signal A, signal B, cannot be stated with certainty]

(c) What is your round income when you receive signal A?

[10, 34, 40]

(d) What is the expected round income of the other person?

[16, 34, 40]

F Pre-Analysis Plan

As registered on the AEA RCT Registry (Bao et al., 2020, AEARCTR-0005369).

Experimental Design

We implement different variants of a Prisoners' Dilemma game with imperfect monitoring in a laboratory experiment. In every round, two players choose their actions $a_i \in \{C, D\}$ simultaneously. Payoffs depend on the player's own action a_i and the received signal about the other player's action ω_{-i} . Under public monitoring, subjects are informed about $(a_i, \omega_i, \omega_{-i})$ at the end of every round. Under private monitoring, subjects are informed about (a_i, ω_{-i}) at the end of every round. The continuation probability δ of the repeated game is 0.8. In the treatments without correlation, signals are drawn independently for each of the chosen actions. Signals are noisy and indicate the wrong action with probability $\epsilon = 0.2$. Therefore, if both players play C , the probability that the two signals differ is 0.32. In the treatments with correlation, both players receive the same signals if they choose the same action. The signals are correct, that is: they indicate cooperation (defection) when both choose cooperation (defection), with probability $1 - \epsilon = 0.8$. However, if their actions differ, signals are drawn independently, as in the treatments without correlation.

In all treatments, subjects engage in a pre-play communication-stage before the first round of every supergame. In this stage, subjects can communicate via a chat-box interface for 120 seconds.

Under private monitoring, two treatments have a reporting stage, which is implemented in the form of a structured communication stage after every round. In this stage, subjects can report the received signal from the current round to their partner (or misreport it).

In every session of every treatment, subjects are randomly divided into 3 matching groups, with 8 each. Subjects play 7 supergames with pre-determined lengths. At the beginning of every supergame, each subject is matched with a new partner from his/her matching group using perfect stranger matching, so that they do not play with the same partner for a second time. To keep the length of supergames constant across treatments, we generated 3 sequences of random numbers beforehand, and used them to determine the length L_i of each supergame.¹¹ To increase the number of observations per supergame, we adapt the block-random-termination method (Fréchette and Yuksel, 2017). Subjects play a block of 5 rounds at the beginning of every supergame. If the true length L_i is smaller or equal than 5, the supergame ends at the end of round 5 and only the first L_i rounds are payoff relevant. If L_i is larger than 5, the supergame continues until round L_i has been reached and all rounds are payoff relevant. Before the end of round 5, subjects are not informed about whether the supergame ends or not. Subjects are required to answer control questions before the game starts. At the end of the experiment, subjects answer a short survey to elicit basic socio-economic characteristics, such as age and gender.

¹¹We used Stata to generate 3 sequences of uniformly distributed random numbers between 0 and 1 with seeds 3, 4, and 5 (we used seeds 1 and 2 in: Dvorak and Fehrler, 2024). Denote the 3 sequences as $\{r_n\}_i = \{r_1, r_2, \dots, r_x\}_i$, where $i = 3, 4, 5$ indicates the seed underlying the sequence and $n \in \mathbb{N}$. The first supergame has x_1 rounds if $r_{x_1} \leq 0.2$ and for all $n < x_1$, $r_n > 0.2$. The second supergame has $x_2 - x_1$ rounds if $r_{x_2} \leq 0.2$ and for all $x_1 < n < x_2$, $r_n > 0.2$. And so forth. The resulting (lengths of the) sequences are SQ1 (2, 8, 1, 5, 7, 1, 7), SQ2 (4, 2, 2, 21, 4, 3, 5) and SQ3 (2, 3, 1, 1, 4, 6, 6).

Experimental Parameters

Figure 1 shows a screenshot of the decision interface.

Figure F1: Stage Game Parameters [Figure 1 from the pre-analysis plan]

Ihre Optionen	Ihr Einkommen bei Signal		Erwartetes Einkommen, wenn die andere Person	
	A	B	Option A wählt	Option B wählt
Option A	32	2	26	8
Option B	40	10	34	16

Notes: Screenshot from the experiment. Payoffs are in experimental currency units with an exchange rate of 50 ECU = 1 EUR.

The left two columns depict the stage-game payoff in experimental currency units. The payoff parameters do not vary across treatments. The last two columns show the expected stage-game payoffs and are calculated given a fixed error rate of 0.2 among all treatments. The parameters are chosen such that two conditions are satisfied:

- 1) Under imperfect private monitoring with signal correlation, there is a quasi-perfect public (truth-telling) equilibrium (QPPE) if reporting is allowed, in which both players play a “reporting grim-trigger” strategy. The reporting grim-trigger strategy prescribes the following behavior: Start with C and report your received signals truthfully, continue cooperating as long as both reports in the previous round are the same, otherwise defect for all subsequent rounds. The reporting mechanism translates the private monitoring into (quasi) public monitoring. An analogous cooperative PPE exists under imperfect public monitoring, in which subject play the same grim-trigger strategy as the one sketched above (but without reporting).
- 2) No cooperative (Q)PPE exists if there is either no correlation or there is correlation but it cannot be detected due to the absence of a reporting stage.

Treatments

We implement up to five different treatments. We begin with collecting data for two treatments with imperfect-public monitoring: one with correlation (T1) and one without (T2). In case we find a statistically significant treatment difference (see next section for details on the test), we continue with two private-monitoring treatments with correlated signals: one with reports (T3) and one without (T4). In case, we find a statistically significant treatment difference between T3 and T4, we continue with the final treatment T5, which is a private-monitoring treatment without correlation but with a reporting stage.

T1 Signals are public and independent.

T2 Signals are public and perfectly correlated if both actions are the same.

T3 Signals are private and perfectly correlated if both actions are the same. Participants can publicly report their private signal after each round.

T4 Signals are private and perfectly correlated if both actions are the same. Participants cannot report signals.

T5 Signals are private and independent. Participants can publicly report their private signal after each round.

Hypotheses Tests, Power, and Further Analyses

In a previous study of communication and cooperation in a noisy, indefinitely repeated Prisoner's Dilemma with uncorrelated signals, we saw high cooperation rates in the first rounds of the supergame with pre-play communication but then a steady and strong decline over the subsequent rounds (Dvorak and Fehrler, 2024). In a pretest session of treatment T1, we again saw high cooperation rates in the first round and a decline afterwards. However, the decline was much weaker.

Based on these observations and the existence (or absence) of cooperative (Q)PPEs in the different treatments, we formulate our main hypotheses:

H1a: *The cooperation rate will be higher in T1 than in T2.*¹²

H1b: *The cooperation rate will be higher in T3 than in T4.*

¹²The correct statement is "The cooperation rate will be higher in T2 than in T1".

H1c: *The cooperation rate will be higher in T3 than in T5.*

We test the corresponding H0s by comparing the cooperation rates in the first 5 rounds of the last 3 supergames between the treatments. We run one-sided t-tests, for which we average the cooperation rates within each matching group and then take these averages as our independent observations.

Our simulations (see next paragraph) suggest that we will have enough power ($> 80\%$) to detect effect sizes of 10 percentage points with 9 matching groups with 8 participants each per treatment.

Simulations for Assessing the Statistical Power

In the simulations, we iterate the following process 10,000 times for various effect sizes Δ :

1. Create a data set of 8 (subjects per matching group) * 9 (number of matching groups per treatment) * 2 (treatments) observations and an indicator variable for treatment 2.
2. Draw random numbers from the Bernoulli distribution with a success probability that starts at 1 in round 1 and then linearly declines to 0.9 in round 5 for treatment 1.¹³
3. Draw random numbers from the Bernoulli distribution with a success probability that starts at 1 in round one and then linearly declines to $0.9 - \Delta$ in round 5 for treatment 2.
4. Average the random draws from rounds 1-5 within each matching group. These are the simulated average cooperation rates.
5. Run a one-sided t-test on the matching-group averages. Return the p -value.

Finally, we compute the share of the p -values smaller than 0.05, which gives us the statistical power. We check the accuracy of the procedure by running it 10'000 times with the same success probabilities in both treatments, which results in a share of p -values smaller than 0.05 of 0.049, which is close to the 5% that we would expect for this scenario. The simulation results indicate that the power increases in Δ and is 80.1% with $\Delta = 0.1$. The power remains similar if we introduce matching-group-specific variation in the slopes of the declining cooperation probabilities.

¹³0.9 was the cooperation rate we observed in round 5 of the pretest of the T1 treatment.

Further Analyses

In addition to testing our three hypotheses, we explore subjects' strategies across the treatments to better understand the aggregate findings. These analyses are explorative in nature and we, therefore, refrain from specifying further hypotheses. The questions we are interested to explore are the following:

- How are participants' choices in the private treatments with reports (T3 and T5) influenced by the reports of the previous period?
- Are the strategies used in the private treatment with reports (T3 and T5) similar to the strategies used in the the public treatment with correlated signals (T1 and T2)?

For the treatments T2 and T3, we are particularly interested to assess how many participants use a reporting grim-trigger strategy. It will further be interesting to compare the estimated strategies in T2 and T3 to the strategies in T1 and T5 for which the reporting grim-trigger strategy is theoretically not supported.

For the analysis of strategies, we build on the strategy frequency estimation method (SFEM) introduced by Dal Bó and Fréchette (2011), and use the R package `stratEst` (Dvorak, 2023), which was first used in Dvorak and Fehrler (2024). The SFEM is frequently used to obtain maximum-likelihood estimates of the shares of a candidate set of strategies in experimental data. However, the results of the SFEM are specific to this set and it is hard to know ex-ante which strategies should be included. To circumvent this problem, we will compute Maximum Likelihood estimates for an endogenously determined number of strategies where the structure of each strategy is the result of a model-selection process. Thus we will infer the strategies from the data rather than imposing a predefined set of strategies. The process will always start with a large number of such generic strategies, which will then be reduced step-by-step using the integrated-completed-likelihood criterion (ICL-BIC, Biernacki et al. (2000)). The ICL-BIC is an entropy-based selection criterion for mixture-models which has been used to estimate the dimensionality of the strategy space in other settings before (Breitmoser, 2015).

To assess whether strategies differ in two treatments, we fit a model on the pooled data of the two treatments and bootstrap the likelihood-ratio test statistic. If the distribution indicates that the likelihood-ratio statistic is sufficiently extreme, we conclude that the strategies differ between the two treatments.