Negotiating Cooperation under Uncertainty: Communication in Noisy, Indefinitely Repeated Interactions*

Fabian Dvořák
University of Konstanz
fabian.dvorak@uni.kn

Sebastian Fehrler
University of Konstanz
sebastian.fehrler@uni.kn

June 2017

Abstract

Case studies of cartels and recent theory suggest that repeated communication is key for stable cooperation in environments where signals about others' actions are noisy. However, empirically the exact role of communication is not well understood. We study cooperation under different monitoring and communication structures in the lab. Under all monitoring structures - perfect, imperfect public, and imperfect private - communication boosts efficiency. However, under imperfect monitoring, where actions can only be observed with noise, cooperation is only stable when subjects can communicate before every round of the game. Beyond improving coordination, communication increases efficiency by making subjects' play more lenient and forgiving. We further find clear evidence for the exchange of private information - the central role ascribed to communication in recent theoretical contributions.

Keywords: infinitely repeated games, monitoring, communication, cooperation, strategic uncertainty, prisoner's dilemma

JEL Classification: C72, C73, C92, D83

*We would like to thank Sam Bowles, Marco Casari, Simon Gächter, Susanne Goldlücke, Ariel Rubinstein, Martin Sefton, Susumu Shikano, Bauke Visser, Nick Zubanov participants at several conferences and seminars, and the TWI group, Urs Fischbacher, Moritz Janas, and Irenaeus Wolff in particular, for helpful comments. We are also very grateful to Guillaume Fréchette and Yves Breitmoser for sharing their programming insights and code of the Strategy Frequency Estimation Method and to Felix Beck and Malin Pimper for excellent research assistance. All remaining errors are our own.

1 Introduction

Many social and economic relationships are characterized by repeated interactions, in which the behavior of partners is only observable with noise. Two examples from very different domains of society are friendships with repeated gift-exchange, and cartels with repeated price-setting of its members. How much effort a friend exerts to find a suitable birthday present cannot be observed directly but only inferred from the present itself - a noisy signal (Compte and Postlewaite, 2015). Likewise, whether or not other cartel members stick to a collusive agreement cannot be observed directly but only inferred from noisy signals like own sales in Stigler's (1964) or the market price in Green and Porter's (1984) seminal treatments of oligopolies. The former is the classic example for imperfect private monitoring - own sales can only be observed by the firm itself - the latter is the classic example for imperfect public monitoring - the market price being publicly observable.

How cooperation can be sustained under imperfect monitoring has been the central topic in the theory of infinitely repeated games for the last three decades. Communication in combination with private monitoring has played a key role in this literature. It has been shown that repeated communication opportunities can enlarge the set of achievable equilibria (Matsushima, 1991; Ben-Porath and Kahneman, 1996; Compte, 1998; Kandori and Matsushima, 1998; Obara, 2009; Awaya and Krishna, 2016), and that there are *truthful communication* equilibria that are stable, while the equilibria without communication that have been analyzed in the literature are not (Heller, 2017).

A number of case studies of cartels suggest that communication is, indeed, crucial for stable cooperation and point to different roles for communication. Genesove and Mullin (2001) note in the abstract of their account of the sugar-refining cartel that its weekly *"[m]eetings were used to interpret and adapt the agreement, coordinate on jointly profitable actions, and determine whether cheating had occured."* Levenstein and Suslow (2006) review the empirical literature on cartels and identify communication as a key ingredient of successful cartel organizations - communication *"not only to provide flexibility in the details of the agreement, but to build trust as well"* (p.67). Finally, Harrington and Skrzypacz (2011), who study various cartel agreements, conclude that truthful communication of sales is an important property of all of them. These accounts suggest (i) that communication is crucial for cooperation in noisy environments, and (ii) that, while the exchange of private information is an important part of it, the role of communication is broader. However, relying on field data has its limitations for the study of cooperation under imperfect monitoring - in particular if signals are private and thus not observable for researchers. Likewise, most cartel communication is not documented as such documents could be used as evidence in legal cases.[1] To overcome

[1]See Genesove and Mullin (2001), Andersson and Wengström (2007), and Cooper and Kühn (2014a) for

these limitations, researchers have begun to design lab experiments. Several studies have explored the effects of communication and tested predictions from renegotiation-proofness refinements (Pearce, 1987; Farrell and Maskin, 1989) in experiments without noise (Andersson and Wengström, 2012; Fonseca and Normann, 2012; Cooper and Kühn, 2014a) or imperfect public monitoring (Embrey et al., 2013). While they offer important insights that we discuss further in Section 2.2, these experiments were not designed to address a number of important questions that arise from recent theoretical results and our own extensions thereof. These include questions concerning the role of repeated communication under private monitoring, of strategic uncertainty under different monitoring structures, and of the complexity of coordination via pre-play communication that we derive below.[2]

We employ an experimental (3 × 3)-design varying both the communication and the monitoring structure of the game. We study the following communication structures: (i) *no communication*, (ii) *pre-play communication*, where subjects can chat with their partner before the first round of an indefinitely repeated game, henceforth called a supergame, and (iii) *repeated communication*, where subjects can chat with their partner before every round of the supergame. The second treatment variable is the monitoring structure. We study (i) *perfect monitoring*, (ii) *imperfect public montoring*, and (iii) *imperfect private monitoring*. The game that we let subjects play is a noisy prisoner's dilemma (PD), similar to that studied theoretically by Sekiguchi (1997) and Compte and Postlewaite (2015), and experimentally, but without communication, by Aoyagi et al. (2017). In this variant of the PD, signals are independent conditional on actions. Payoffs depend on own actions and received signals, which are noisy reflections of the other player's actions. Under perfect monitoring, signals and actions are observed; under imperfect public monitoring, sent and received signals are observed; and under private monitoring, only the received signals are observed. For the case of private monitoring, Heller (2017) shows that only defection can be sustained, by any of the mechanisms discussed in the literature, as an equilibrium that survives his weak stability criterion. He further shows that if players can communicate repeatedly, there typically are cooperative *truthful communication* equilibria, which are weakly stable and even survive the stronger criterion of evolutionary stability. In our parametrization this is the case, which allows us to address our first key question: (1) is there a benefit of repeated communication as compared to no or pre-play communication, and if so, does it differ between monitoring structures, as Heller's (2017) results suggest?[3]

further discussion and examples of cartel cases.

[2]The only study that we are aware of that also studies communication under private monitoring is by Arechar et al. (2017), who limit the message space in a way that allows subjects to report their action but rules out any other form of communication. Camera et al. (2013) study communication in a setting with random re-matching within groups after every round of the repeated game. Vespa and Wilson (2016) study an indefinitely repeated version of a sender-receiver game (Crawford and Sobel, 1982).

[3]Heller (2017) himself laments the lack of experimental studies on communication under private monitoring

More generally, indefinitely repeated games that allow for cooperative equilibria usually feature many such equilibria and coordination on one of them is a major challenge. Depending on the parameters of the game, mis-coordination can be costly and make the choice of a cooperative strategy risky. Blonski et al. (2011) and Dal Bó and Fréchette (2011) adapt Harsanyi and Selten's (1988) concept of risk dominance to infinitely repeated games and show that cooperation is much more frequent if it is risk dominant (see also Breitmoser, 2015; Dal Bó and Fréchette, 2017). However, these concepts are only defined for games with perfect monitoring and cannot be directly applied to imperfect-monitoring settings. To derive predictors for these structures, we extend the theoretical work by adapting Blonski et al.'s (2011) and Dal Bó and Fréchette's (2011) measures to the cases of imperfect public and private monitoring, and by constructing a third measure in the spirit of Breitmoser's (2015) strategic interpretation of Blonski et al. (2011). These extensions allow us to address our next empirical question concerning the no-communication scenarios: (2) will subjects cooperate more under perfect monitoring where strategic uncertainty is much lower than under the imperfect-monitoring structures in our parametrization of the stage game. This step complements Aoyagi et al.'s (2017) study, who compare cooperation under the three monitoring conditions in an experiment without communication and stage game parameters that lead to very similar and low measures of strategic uncertainty across treatments. They find no significant differences in the cooperation rates.

While strategic uncertainty has been shown to matter under perfect monitoring without communication, it has also long been recognized that communication can help coordination (e.g., Cooper et al., 1992; Rabin, 1994; Ellingsen and Östling, 2010), and coordination on a cooperative equilibrium would decrease strategic uncertainty. Communication might thus increase cooperation in settings where it is not risk dominant. Pre-play communication should, in principle, be sufficient for coordination on a cooperative equilibrium. However, while efficient equilibria are easy to find in the perfect-monitoring case, this task becomes a lot more difficult under imperfect public monitoring. Even if players cooperate, bad signals occur with positive probability and thus players will likely have to enter a phase of punishment at some point. For this reason, simple punishments, such as *defect forever* after a bad signal, are inefficient and players have to coordinate on lenient or forgiving strategy profiles to reap a greater share of the potential gains of cooperation. With private monitoring it gets even worse. The equilibria that have been found and analyzed in the literature are all mixed (or behavioral) strategy profiles, which are extremely hard to find, and coordination on these equilibria seems highly unlikely (Compte and Postlewaite, 2015). For these reasons, it is unclear whether pre-play communication will be sufficient for coordination on efficient equilibria when monitoring is imperfect. Repeated communication could, on the

which would allow a test of his predictions.

one hand, help to overcome insufficient coordination by allowing subjects to revisit incomplete agreements. On the other hand, renegotiation concerns could undermine coordination on cooperation before the start of the game, as assumed in renegotiation-proofness refinements (Pearce, 1987; Farrell and Maskin, 1989). These considerations bring us back to our first question regarding the additional benefits of repeated communication as compared to pre-play communication, and raise the new question (3) in how far pre-play communication helps subjects to coordinate and increase cooperation as compared to the no communication case.

Further roles for communication have been discussed in the behavioral economics literature. Promises, apologies, and threats of punishment have been shown to influence cooperative behavior by building or restoring trust in settings without noise (Charness and Dufwenberg, 2006; Utikal, 2012; Fischbacher and Utikal, 2013; Camera et al., 2013; Cooper and Kühn, 2014a). If such effects played a role in our setting, communication might entail more cooperation, leniency, and forgivingness. However, it is not obvious that these findings will carry over to scenarios with uncertainty. Several studies of settings with *moral wiggle room* show that subjects become more selfish when their behavior is only observable with noise (e.g., Dana et al., 2007; Larson and Capra, 2009). The imperfect monitoring of their actions gives them a disguise for less moral behavior and might mute the role of image concerns (Bénabou and Tirole, 2006). Coordination on efficient equilibria under imperfect monitoring and behavioral effects could both affect strategy choices, which leads us to the next research question: (4) Will subjects play more lenient and/or forgiving strategies with communication?

To summarize the discussion so far: many different roles for communication are plausible. To avoid pushing it toward one of these roles, we chose open chat as the mode of communication. Free-form communication is also the most natural form and allows us to study its use and content, and thereby to shed light on the following final questions: (5) What will subjects talk about? (6) Will subjects share private information under private monitoring? (7) Which content will be correlated with cooperative behavior in the game?

Our main results are the following: (1) With repeated communication cooperation rates are high and stable under all monitoring conditions, whereas they start high but decline much more rapidly with pre-play communication if monitoring is imperfect. Under perfect monitoring the additional benefit of repeated communication is much smaller. (2) Without communication cooperation rates are indistinguishable between the monitoring structures. (3) Cooperation rates in the pre-play communication treatments are much higher under all monitoring structures than in the *no communication* treatments. This suggests that coordination via pre-play communication effectively reduces strategic uncertainty. Moreover, we find (4) that subjects' play becomes both more lenient and more forgiving with communication, (5) that communication is mainly used to share information, to coordinate behavior and to talk about issues not directly related to the experiment, (6) that subjects share private information under private monitoring but that less than half of the private signals are reported, and (7) that communication falling into any of the three main categories of communication content - information sharing, coordination, and trivia - is positively correlated with cooperation.

The rest of the paper is structured as follows. In the next section, we present the game and its theoretical properties, thereby reviewing the related literature and extending the theoretical predictors of cooperation to the imperfect monitoring cases. In Section 3, we present the experimental design and state our research questions, which we address in Section 4. The paper ends with a discussion of our key findings.

2 The Repeated Prisoner's Dilemma with Noise

Two players interact with each other in indefinitely many rounds of a supergame. Let δ denote the fixed continuation probability after any given round. In every round, each of the two players can choose between two actions C or D. After both players have chosen an action $a \in \{C, D\}$, a noisy process translates both actions into conditionally independent signals. Each signal $\omega \in \{c, d\}$ corresponds to the chosen action with probability $(1 - \epsilon)$. With probability ϵ, an error occurs and the action is translated into the wrong signal (C to d and D to c). All aspects of this process, the conditional independence of signals as well as the probability of an error are common knowledge. The payoff π_i of player i from the current round is defined by player i's own action a_i and the signal of the other player's action ω_{-i}.[4] We consider the following normalized expected stage-game payoffs of action profiles which take the noise into account:

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

Since $g > 0$ and $l > 0$ the stage game has the form of a prisoner's dilemma. We consider three different monitoring structures. Under perfect monitoring each player i is informed about the actions (a_i, a_{-i}) and the signals (ω_i, ω_{-i}). Under imperfect *public* monitoring (Green and Porter, 1984), players cannot observe the action of the other player a_{-i} as the information set reduces to $(a_i, \omega_i, \omega_{-i})$. Under imperfect *private* monitoring (Stigler, 1964), players also remain uninformed about ω_i, the signal the other player receives, as the information set

[4] One could interpret this as a gift-exchange game, in which players exert low or high effort into finding a suitable present for the other partner, and whether the present is suitable or not is a noisy signal of effort (Compte and Postlewaite, 2015). For an alternative but similar interpretation see Sekiguchi (1997).

reduces to (a_i, ω_{-i}). The conditions for cooperative subgame-perfect equilibria under the three different monitoring structures are well-known results of the theoretical literature (see, e.g., Mailath and Samuelson, 2006). With perfect monitoring, players can condition on the intended actions and support full cooperation using pure strategies if the continuation probability δ is greater or equal to $\delta^{SPE} = \frac{g}{1+g}$. With public monitoring and strategies conditioning only on the public signals the stricter condition $\delta^{SPE} = \frac{g}{1-\epsilon+(1-\epsilon)^2 g}$ applies with reduced efficiency.[5] With private monitoring, cooperation cannot be supported by a SPE based on pure strategies and players have to rely on mixed (Bhaskar and Obara, 2002; Sekiguchi, 1997) or behavior strategies (Ely and Välimäki, 2002; Piccione, 2002).[6]

2.1 Predictors of Cooperation

Experimental evidence suggests that the SPE conditions are necessary but insufficient to observe high levels of cooperation in the laboratory (for a survey see Dal Bó and Fréchette, 2017). More accurate predictors of cooperation exist for the case of perfect monitoring. We highlight the two most prominent predictors and provide their extensions to public and private monitoring. We also extend the existence condition for equilibria in memory-one belief-free strategies, henceforth m1BF (Ely and Välimäki, 2002; Piccione, 2002), a class of strategies which is studied in both the theoretical and experimental literature on repeated games (Heller, 2017; Aoyagi et al., 2017; Breitmoser, 2015). We extend Breitmoser's (2012) existence condition of equilibria in m1BF strategies which condition on actions to equilibria in m1BF strategies which condition on public signals and on action-signal combinations, respectively.

Dal Bó and Fréchette (2011) develop the basin of attraction of defection, henceforth *BAD*, as a predictor of cooperation and show that it explains cooperation levels under perfect monitoring. In a mixed population of grim-trigger (GRIM) and always defecting players (ALLD), the *BAD* is defined as the share of GRIM which makes players indifferent between the two strategies. Let π^{DF} denote the probability of playing GRIM. Under perfect monitoring, since $\pi/(1-\delta) - (1-\pi)l = \pi(1+g)$ the BAD is defined as:

$$\pi^{DF} = \frac{l}{l - g + \frac{\delta}{1-\delta}}$$

[5] The continuation probability must be high enough to make symmetric grim-trigger a SPE. The crucial subgame is the state where both players cooperate. In this state, the long-run incentives of cooperation must be as least as big as the immediate gains from defection which requires $1 + \frac{\delta}{1-\delta} \geq 1 + g$ under perfect monitoring and $\frac{1}{1-\delta(1-\epsilon)^2} \geq \frac{1+g}{1-\delta\epsilon(1-\epsilon)}$ under public monitoring.

[6] Under private monitoring, players lack the public signal to coordinate behavior in such a way that defection after a defection signal is a mutual best-response. Instead, players believing that the other player cooperated and saw a cooperation signal with a high probability would want to ignore the bad signal. This incentive to ignore bad signals undermines the necessary responsiveness of the strategy to defer defection.

In contrast to the SPE condition the BAD also takes the *sucker* pay-off $-l$ into account. Dal Bó and Fréchette (2011) interpret the BAD as the degree of strategic uncertainty associated with cooperation, and note that $\pi^{DF} = 0.5$ is the δ threshold where cooperation becomes risk dominant in the spirit of Harsanyi and Selten (1988). The BAD is inversely related to the frequency of cooperation observed in the laboratory (Dal Bó and Fréchette, 2017). The extension of the BAD to the imperfect monitoring structures is straightforward and reveals that the strategic uncertainty of cooperation increases with noise.[7] Whether the size of the BAD also predicts cooperation levels under imperfect monitoring is an open question.

Blonski et al. (2011) use an axiomatic approach to develop a predictor of cooperation: the BOS threshold. They use five axioms to identify a critical δ as a function of the three incentives: the long-run incentive to cooperate ($\frac{\delta}{1-\delta}$), the short-run incentive to defect if the opponent cooperates (g), and the short-run incentive to defect if the opponent defects as well (l). The BOS threshold is

$$\delta^{BOS} = \frac{g+l}{1+g+l}$$

and corresponds to the risk-dominance condition based on $\pi^{DF} = 0.5$. Breitmoser (2015) shows that the BOS threshold has implications for the set of cooperative equilibrium strategies. It is the existence condition of a sub-class of memory-one belief-free equilibria which are in semi-grim strategies, henceforth semi-GRIM. He provides empirical evidence that behavior on both the aggregate and the individual level is well summarized by these strategies (Breitmoser, 2015). The extensions of the BOS threshold to public and private monitoring based on Breitmoser's interpretation show that δ^{BOS} increases with noise.[8]

We complement these predictors of cooperation by deriving an existence condition for equilibria in memory-one belief-free strategies. Breitmoser (2012) defines the existence condition for the case where the strategies condition on actions. We extend the condition to the case where the strategies condition on the signals (ω_i, ω_{-i}) and the case where the strategies condition on the action signal combination (a_i, ω_{-i}). The last two cases can be used to support cooperation under imperfect monitoring. We show that there exists a δ for each of the three cases where all m1BF equilibrium strategies show the same response pattern

[7] With public monitoring, since $\pi\frac{1}{1-\delta(1-\epsilon)^2} - (1-\pi)\frac{l}{1-\delta\epsilon(1-\epsilon)} = \pi\frac{(1+g)}{1-\delta\epsilon(1-\epsilon)}$ the BAD is $\pi^{DF} = \frac{l}{l-g+\frac{\delta((1-\epsilon)^2-\epsilon(1-\epsilon))}{1-\delta(1-\epsilon)^2}}$. With private monitoring, since $\pi\frac{1+\delta\epsilon(1-\epsilon)(1+g-l)/(1-\delta\epsilon)}{1-\delta(1-\epsilon)^2} - (1-\pi)\frac{l}{1-\delta\epsilon} = \pi\frac{(1+g)}{1-\delta\epsilon}$ it is $\pi^{DF} = \frac{l}{l-g+\frac{\delta((1-2\epsilon)-\epsilon(1-\epsilon)(l-g))}{1-\delta(1-\epsilon)^2}}$
.

[8] The extensions are $\frac{g+l}{(1-2\epsilon)(1+g+l)}$ for semi-GRIM strategies which condition on the public signals and $\frac{g+l}{(1-2\epsilon)1+(1-\epsilon)(g+l)}$ for semi-GRIM strategies which condition on action-signal combinations. Since the threshold for action-signal combinations is always lower, we refer to it as the BOS threshold for public and private monitoring (see Appendix A for details).

to all possible histories after round one. The m1BF strategies can be represented by a vector of five cooperation probabilities which correspond to five possible memory-one histories. Let this vector be $\sigma = (\sigma_\emptyset, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$. The history (or state) cd, for instance, prevails if the player in question cooperated in the last round and the other player defected. $\sigma_\emptyset$ represents the initial cooperation probability where $\emptyset$ represents the empty history before round one. If the history is defined by the actions chosen in the last period (which is only possible under perfect monitoring), the existence condition for m1BF equilibria is

$$\delta^{BF} = \frac{\phi}{1 + \phi}$$

where ϕ denotes the larger of the two values g and l. Note that if $g \geq l$ the δ^{BF} threshold corresponds to δ^{SPE} under perfect monitoring. If $l > g$, both conditions differ with $\delta^{BF} > \delta^{SPE}$. The extensions to the cases where the strategies condition on the public signals action-signal combinations reveal higher thresholds which increase in the level of noise.[9] For all three cases it can be shown that all m1BF equilibria exhibit the same pattern of cooperation probabilities after the four non-empty memory-one histories (cc,cd,dc,dd) if $\delta = \delta^{BF}$. We call this response behavior the threshold memory-one belief-free response (T1BF). At $\delta = \delta^{BF}$ there exists a continuum $\sigma_\emptyset \in (0, 1)$ of equilibrium strategies with T1BF since the initial cooperation probability is a free parameter. The pattern of T1BF is defined by the stage-game parameters in the same way for all three cases (but usually occurs at different values of δ). If $l > g$, the response is a lenient form of Tit-for-Tat. If $g > l$, it is a forgiving form of GRIM. In the frequently studied case $g = l$, it is the Tit-for-Tat response (see Appendix A for the proof).

2.2 Communication

Renegotiation-proofness refinements (Pearce, 1987; Farrell and Maskin, 1989) are the most widely used tools to restrict the usually large set of equilibria in repeated games that allow for cooperation. Weak renegotiation proofness (Farrell and Maskin, 1989) requires that an equilibrium strategy profile must not have continuation values in any subgame that are Pareto-dominated by the continuation values in another subgame - the idea being that subjects would otherwise renegotiate away from the former to the latter. Equilibria that support cooperation with strongly symmetric strategies, such as equilibria where both players defect in the punishment state, do not survive this refinement because players would

[9]For public signals, the condition is $\delta^{BF} = \frac{\phi}{1-2\epsilon-\epsilon\psi+(1-\epsilon)\phi}$ and for action-signal combinations it is $\delta^{BF} = \frac{(1-\epsilon)\phi-\epsilon\psi}{(1-2\epsilon)(1-2\epsilon+(1-\epsilon)\phi-\epsilon\psi)}$ where ψ stands for the smaller of the two values g and l. Note that with noise, the conditions are equivalent only if $g = l$.

otherwise renegotiate and restart the game. However, weakly renegotiation-proof cooperative equilibria often exist in the indefinitely repeated prisoner's dilemma (van Damme, 1989). They require more complex behavior in the punishment phase, where players have to play asymmetrically. The player that deviated must play C while the punisher plays D. After a certain number of rounds the punishment phase ends and play restarts with mutual cooperation.[10] Such an equilibrium is arguably more complicated to coordinate on, which has led some authors to restrict attention to strongly-symmetric strategies. Embrey et al. (2013), for example, adapt Pearce's (1987) slightly different renegotiation-proofness concept to derive predictions for a game with imperfect monitoring. In their variant of renegotiation-proofness *"a candidate equilibrium would survive potential renegotiation if there is no other perfect public equilibrium, using strongly symmetric two-state automata, that has a larger expected value in the punishment state"* (p.11). In addition to considering only strongly symmetric strategies, they restrict their attention to perfect public equilibria. The reason for this additional restriction is that renegotiation concepts rely on the existence of multiple subgames, which requires that subjects condition their play on the public history. If subjects instead condition on private histories, as they do in the belief-based equilibrium construction by Sekiguchi (1997) or in belief-free equilibria (Piccione, 2002; Ely and Välimäki, 2002), the only subgame is the entire game.

Predictions that are based on these refinements have been tested in a number of experimental studies with repeated communication. Results are mixed. Cooper and Kühn (2014a), who study two-stage games, and Embrey et al. (2013) find no reduction in cooperation when renegotiation-proofness predicts less cooperation.[11] Andersson and Wengström (2012), also studying a two-stage game with structured communication, find that pre-play messages are more effective if renegotiation between the two periods is not possible. They observe slightly lower cooperation rates with repeated as compared to pre-play communication. Cooper and Kühn (2014a and 2014b) compare treatments with structured and free-form communication via a chat interface, and find that cooperation rates are higher with the more natural free-form communication.

As discussed above, in the absence of communication there are only complicated equilibrium constructions under private monitoring. However, when players can communicate repeatedly, private signals can be reported, which creates a public history and thereby allows for simpler and more stable equilibria (Heller, 2017). Such *truthful communication* equilibria

[10]For the existence such equilibria l must not be too large. Otherwise, the return to CC is not attractive enough for the punished to agree on playing C and receiving $-l$ in the punishment phase.

[11]Fonseca and Normann (2012), also studying a two stage game, and Camera et al. (2013), studying a game with random rematching in groups after every round, find a positive rather than a negative effect of repeated communication on cooperation. Neither of the two studies explicitly tests renegotiation-proofness predictions.

can exist if certain revelation constraints are fulfilled (Compte, 1998). The punishment stage is constructed in a way that makes every player indifferent between truthfully reporting her private signal and misreporting or staying silent. This requires that no player benefits or suffers from entering the punishment phase in which the other player is punished. The stability of these equilibria stems from the fact that they provide strict incentives for cooperation, whereas the other equilibrium constructions by Sekiguchi (1997) or Piccione (2002) and Ely and Välimäki (2002) do not (see Heller, 2017).[12]

3 Experimental Design

Our experiments follows a 3 (monitoring structure: perfect, imperfect public, imperfect private) × 3 (communication: none, pre-play, repeated) between-subject design with 9 experimental treatments that we henceforth abbreviate in the following way: PerNo, PerPre, PerRep, PubNo, PubPre, PubRep, PriNo, PriPre, PriRep. We implement the three different monitoring conditions like Aoyagi et al. (2017). Under *perfect* monitoring both players are informed about the intended actions (a_i, a_{-i}) and the signals (ω_i, ω_{-i}). Under *public* monitoring, players are given the reduced information set $(a_i, \omega_i, \omega_{-i})$. Under *private* monitoring, players are only informed about (a_i, ω_{-i}).

In addition to the three different monitoring conditions, we implement three different communication conditions. The benchmark case is that of *no communication* (as in Aoyagi et al., 2017). In the *pre-play communication* condition, subjects enter an open-chat communication stage before the first round of a supergame. The chat can be used by both players of the current match to exchange messages for 120 seconds. In the *repeated communication* condition, players additionally enter a communication stage before each of the following rounds where they can exchange messages for 40 seconds.

To keep the length of the supergames constant between treatments, we generate two sequences of supergames beforehand using a series of random numbers to determine the length of each supergame.[13] Both sequences are implemented for all treatments in different sessions.

[12] If signals are correlated, which is not the case in our set-up, *truthful communication* equilibria with strict revelation constraints can be constructed (Kandori and Matsushima, 1998), and higher levels of efficiency might be achievable by exploiting the informational content of the correlation (Awaya and Krishna, 2016). Awaya and Krishna (2016) study a set-up with a fixed discount rate, whereas other studies have focused on proving Folk theorems (Ben-Porath and Kahneman, 1996; Compte, 1998; Kandori and Matsushima, 1998; Obara, 2009).

[13] We use the Stata random number generator with seeds 1 and 2 to create two series of uniformly distributed random numbers between zero and one. The first supergame had x rounds if the xth random number was less than or equal to 0.2 and all previous numbers were greater than 0.2. Then the first x random numbers were deleted and the following numbers determined the length of the second supergame, and so forth. We used the two series to determine the lengths of seven supergames each. The length of the two resulting sequences of supergames are: SQ1 (11 3 5 1 5 2 11) and SQ2 (2 5 5 7 13 4 4). Average supergame lengths were moderately

At the end of every round of a supergame, subjects receive feedback about their earnings and additional information which allows them to (imperfectly) monitor others' decisions. The realized random number, which determined whether the supergame continued or not, is also displayed at the end of each round, and could thus be used as a public randomization device. To allow for learning, each participant in our experiment plays 7 super-games with different partners. The matching proceeds as follows: we divide the subjects of an experimental session into matching groups of 8 to 12 subjects. For the first supergame, each subject is then randomly matched with another participant from their matching group. After termination of a supergame, participants are rematched with a new partner from their matching group who they did not interact with before. Subjects were informed about this matching procedure. Before the start of the treatment, participants had to answer control questions to check their understanding of the instructions.

We collected data from three matching groups per sequence-treatment combination, that is from six matching groups per treatment. A total of 458 participants participated in the 24 sessions of our experiment at the LakeLab of the University of Konstanz.[14] Table 1 below summarizes the distribution of sessions, subjects and the average size of the matching groups across experimental treatments.

Table 1: Summary Statistics for the Experimental Treatments

	Perfect			Public			Private		
	No	Pre	Rep	No	Pre	Rep	No	Pre	Rep
Sessions	6	6	6	6	6	6	6	6	6
Subjects	52	54	54	48	52	50	48	50	50
Mean group size	8.7	9.0	9.0	8.0	8.7	8.3	8.0	8.3	8.3

Notes: Mean group size indicates the average number of subjects who formed a matching group. The modal size of a matching group was 8 (44 groups). Seven groups were of size 10 and three of size 12.

longer than the expected length of 5 of the underlying geometric distribution (SQ1: 5.4; SQ2: 5.7). Random termination is the most widely used form of implementing infinitely repeated games in the lab. See Fréchette and Yuksel (2017) for a study on other implementation methods.

[14]The experiment was programmed in z-Tree (Fischbacher, 2007) and subjects were recruited via ORSEE (Greiner, 2015).

3.1 Experimental Parameters

The upper part of Figure 2 shows the payoff a subject receives in a round of a supergame as a function of the action and the received signal about the other player's action. We use the same payoff structure, the same continuation probability of $\delta = 0.8$ and the same error probability of $\epsilon = 0.1$ in all our treatments. These values translate into expected stage-game payoffs profiles for actions depicted in the lower part[15].

Table 2: Stage-Game Parameters

	c	d
C	30	0
D	37	17

	C	D
C	27,27	3,35
D	35,3	19,19

Table 3: Predictors of Cooperation

	Perfect	Public	Private
δ^{SPE}	0.5	0.65	–
π^{DF}	0.4	0.76	0.77
δ^{BOS}	0.75	0.86	0.86
δ^{BF}	0.67	0.85	0.8

Table 3 shows the values of the predictors of cooperation which result from the parameters. We choose the parameters for the following reasons: First, the parameters are such that without communication we expect low levels of cooperation based on an analysis of the BAD under imperfect monitoring and a somewhat higher level under perfect monitoring (but still far below full cooperation). This leaves scope for higher cooperation levels in the communication treatments.[16] Second, we want to focus on the main difference between the public and private monitoring treatments identified in the theoretical literature. This is the possibility to support cooperation based on pure strategies with public signals. We chose parameters, which lead to a similar BAD with public and private monitoring. The parameters also rule out that the set of memory-one belief-free strategies is different between public and private monitoring since no m1BF equilibria where strategies condition on the public signals exist. Third, to facilitate an equilibrium prediction for the use of m1BF strategies, the parameters predict the same T1BF pattern for m1BF equilibria under perfect and imperfect monitoring. This

[15]Payoffs are in experimental currency units. The exchange rate was 50 ECU = 1 Euro. Subjects saw both representations of the stage-game at all times when making their decisions.

[16]Our no-communication treatments complement the treatments of Aoyagi et al. (2017) where the BAD takes the values of 0.03, 0.15, 0.13 for perfect, public, private monitoring.

pattern is a lenient version of Tit-for Tat with $\sigma = (\sigma_\emptyset, 1, 0.5, 1, 0)$ conditional on action-signal combinations. Finally, we are interested in whether subjects use communication to transform the game with private monitoring into one with public signals. The parameters assure that equilibria exist where players truthfully reveal their private signals under private monitoring. The parameters also assure the existence of renegotiation-proof cooperative equilibria under all three monitoring structures.

3.2 Research Questions and Methods

Question 1: [Repeated vs. Pre-play] *Is there a benefit of repeated communication? Does the benefit of repeated communication differ across monitoring structures?*

According to Heller's (2017) stability analysis, we would expect a large positive effect of repeated communication on cooperation under private monitoring, whereas high cooperation is already achievable in stable equilibria without communication under public and perfect monitoring. In the latter two monitoring structures, weak renegotiation-proofness (Farrell and Maskin, 1989) eliminates cooperative equilibria in strongly symmetric strategies and repeated communication might thus have a negative effect. However, there are cooperative equilibria that are weakly renegotiation-proof. Finally, repeated communication might have an additional benefit under imperfect monitoring where coordination on an efficient equilibrium is difficult and players might need to revisit incomplete agreements after round 1. To address *Question 1*, we compare the average frequency of cooperation and the average stability of cooperation over rounds between repeated and pre-play communication within the same monitoring structure. We subsequently compare the differences between repeated and pre-play communication across different monitoring structures.

Question 2: [No Communication] *Without communication, is cooperation less frequent with imperfect monitoring as predicted by strategic uncertainty?*

As strategic uncertainty is substantially higher under imperfect monitoring than under perfect monitoring, we expect lower cooperation rates in these treatments. To answer *Question 2* we compare the level and the stability of cooperation without communication across monitoring structures.

Question 3: [Pre-play vs. No Communication] *What is the benefit of pre-play communication? Does the benefit of pre-play communication differ across monitoring structures?*

We expect pre-play communication to facilitate coordination and thereby to lower strategic uncertainty. This should have a positive effect on cooperation rates. However, as already mentioned above, coordination on an efficient equilibrium is much more difficult under imperfect monitoring. As a consequence, the effect might be larger under perfect monitoring. To answer *Question 3*, we execute the same test sequence outlined for *Question 1* for the differences between pre-play and no communication treatments.

Question 4: [Strategies] *Does communication change the use of strategies? Does anticipated verbal punishment or more lenient or more forgiving behavior arise?*

Coordination on efficient equilibria and behavioral effects such as trust building, apologies, threats of puishment and anticipation of verbal punishment, could all lead to more cooperativeness, leniency and forgivingness. To answer *Question 3*, we adapt and use the strategy frequency estimation method (SFEM) by Dal Bó and Fréchette (2011) to explore the use of strategies in our treatments (see Appendix B for details). We begin by characterizing the behavior of subjects in terms of an average memory-one Markov strategy.[17] To this end, we first estimate the strategy shares and parameters of five memory-one Markov strategies of the form $\sigma = (\sigma_\emptyset, \sigma_{cc}, \sigma_{cd}, \sigma_{dc}, \sigma_{dd})$, which generate the highest likelihood in our experimental treatment. We integrate the five strategies of each treatment to the average Markov strategy by calculating the average of each strategy parameter from the five estimates weighted by their shares. We interpret the probability of cooperation in the five possible memory-one states $(\emptyset, cc, cd, dc, dd)$ in the following way. $1 - \sigma_{cc}$ is interpreted as an estimate for the frequency of unjustified defection since most cooperative strategies require cooperation in this state. If anticipation of verbal punishment played a role, we would expect unjustified defection to be less frequent in treatments with repeated communication. σ_{cd} and σ_{dc} give an estimate of the frequency of lenient behavior and forgiving behavior while σ_{dd} indicates how frequently subjects return to cooperation after mutual defection. $\sigma_\emptyset$ is the probability of initial cooperation. Besides characterizing the average strategy, we also want estimates for the heterogeneity of strategies in our treatments. To this end, we use the SFEM and estimate the strategy shares of a standard set of pure strategies (Fudenberg et al., 2012) in our treatments. We add one strategy to this set with the T1BF pattern and starting with cooperation ($\sigma_\emptyset = 1$).

Question 5: [Communication Content] *What do subjects communicate about?*

[17]We exploit the fact that constructing an average strategy based on different strategies is straightforward if we restrict the estimation to memory-one Markov strategies and generates an easy to interpret proxy for subjects' average behavior. Note that estimating one memory-one Markov strategy yields different results except for the initial cooperation probability $\sigma_\emptyset$ than estimating the average reaction for each state. This would not account for the fact that subjects select into states conditional on their strategy. Of course, the approach we use also has its limitations. If we estimate subjects' responses in a certain state the quality of the estimate depends on the number of observations in that state which might be low due to the same reason and produce noisy averages.

We expect pre-play communication to be used for coordination. Many other uses are conceivable (see the discussion in the introduction) and this analysis will therefore be of exploratory nature. However, under imperfect private monitoring, we expect a very specific and important role of communication, as highlighted by the next question.

Question 6: [Private Information Sharing] *Do subjects share private information?*

This is the key role ascribed to communication under private monitoring in the recent theoretical literature. Only *truthful communication* equilibria, in which private signals are reported, are stable (Heller, 2017). To give an answer to *Questions 5 and 6*, two research assistants coded the content of communication based on 72 sub-categories from which we created five main categories (Table 14 in Appendix C summarizes the categories). These five categories are Coordination, Deliberation, Relationship, Information and Trivia. The coding was done on the sub-category level for subject-round observations and multiple coding was possible. We only consider a coding as valid if both raters independently indicated the same sub-category for a subject-round observation. When comparing differences between the use of communication across treatments we limit our attention to the category level. For the more detailed analyses on information sharing and verbal punishment we use the respective sub-categories.

Question 7: [Content-Cooperation Correlations] *What are the correlations between different forms of communication content and cooperation?*

Given the different potential roles for communication discussed so far, we expect coordination attempts, sharing of private information and communication with the intention to build or restore trust to be positively related to cooperation in the game. We employ standard regression techniques to answer *Question 7*. As communication content is endogenous, the results obtained in this step of the analysis can only be suggestive of causal effects.

4 Experimental Results

A common result in the experimental literature is that subjects need a few supergames to adapt their behavior to the experimental environment (e.g., Dal Bó, 2005). We observe a considerable amount of learning in supergames 1-4, too. In the following, we will present results based on the last three supergames where subjects' behavior has stabilized.[18] Unless explicitly indicated otherwise, the findings are robust if we consider data from all 7 supergames.

[18]Figure 4 in Appendix D displays the evolution of cooperation over the supergames.

4.1 Cooperation

Figures 1 and 2 present two correlated but not necessarily congruent measures of cooperation: the average frequency of cooperation and the average stability of cooperation over rounds. We provide answers to Questions 1-3 based on these two figures.

Question 1: [Repeated vs. Pre-play] *Is there a benefit of repeated communication? Does the benefit of repeated communication differ across monitoring structures?*

Figure 1 shows the average frequency of cooperation across the nine experimental treatments for the last three supergames. The depicted levels of cooperation mostly reflect the amount of cooperation observed in the four rounds where each subjects contributes two or three observations depending on the sequence. The bars show that the mean cooperation level in treatments with repeated communication is higher compared to the treatments with pre-play cooperation under all three monitoring structures (Permutation Tests: perfect, $p = 0.03$; public $p = 0.01$; private $p = 0.07$). The size of the effect is largest under public monitoring where the mean cooperation level is 19 percentage points higher with repeated communication. Differences in differences tests which compare the size of the effect under perfect monitoring to the effect sizes under public and private monitoring are insignificant (Permutation Tests: perfect vs. public, $p = 0.17$; perfect vs. private, $p = 0.36$).[19] Figure 2 shows the mean cooperation level over rounds for the last three supergames. Data is shown until round 11 to assure that each subject contributes at least one data point for every round. The lines illustrate where differences between the repeated and pre-play communication treatments arise. Cooperation levels in round one are generally above 90% with communication and do not differ much between treatments. With repeated communication, cooperation levels are more stable over rounds compared to pre-play communication. The effect is much bigger in the imperfect monitoring treatments where the average cooperation level reduces by 36 percentage points over 11 rounds with pre-play communication but only by 13 percentage points with repeated communication. In contrast if monitoring is perfect, the average cooperation level only reduces by 9 percentage points with pre-play communication and does not decline at all with repeated communication. To assess whether the differences are significant we fit a linear OLS model on the data from each matching group and compare the distributions of the coefficients across treatments. The results indicate that the differences in the stability of cooperation are significant in the treatments with imperfect monitoring (Wilcoxon rank-sum test: perfect, $p = 0.15$; public, $p = 0.02$; private, $p = 0.07$).

[19]Figure 5 in Appendix D displays efficiency levels in comparison to different equilibrium benchmark efficiency rates.

Figure 1: Average Frequency of Cooperation Across Treatments

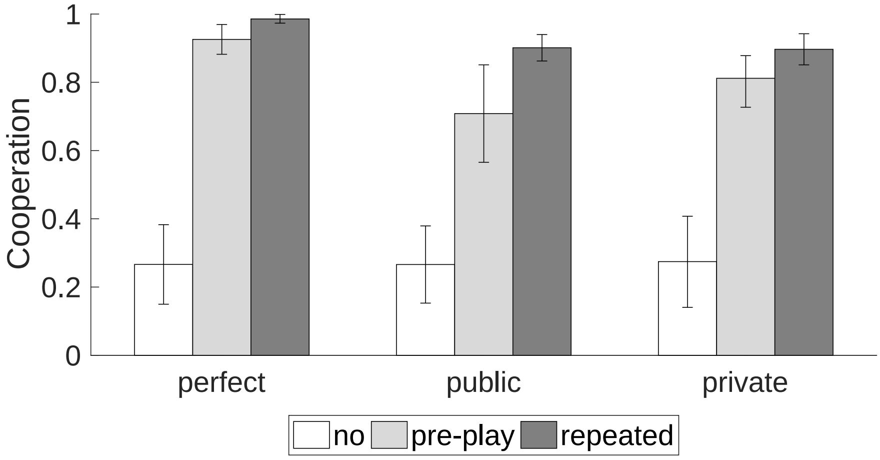

Notes: Bars show relative frequency of cooperation in last 3 supergames. Whiskers depict two-way clustered SE of mean (clustered on subject and match).

Figure 2: Stability of Cooperation over Rounds

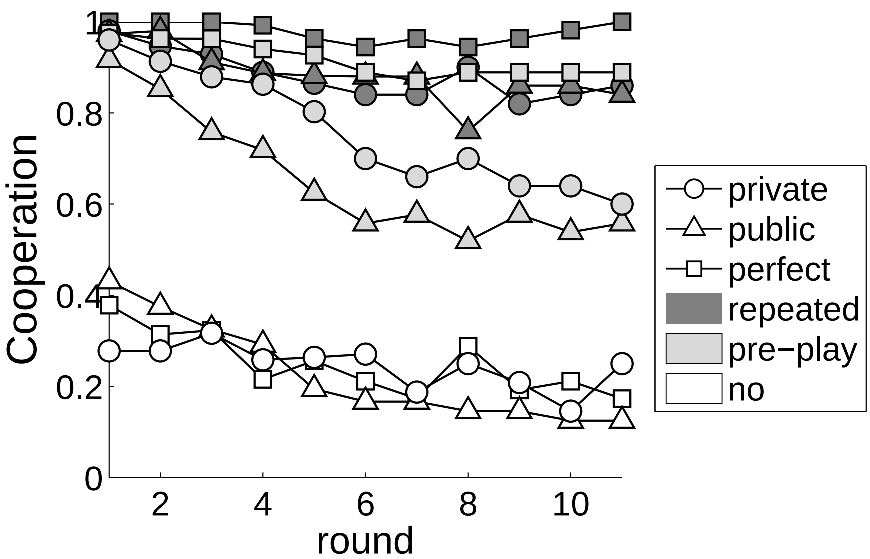

Notes: Graph depicts the frequency of cooperation over rounds averaged over the last last 3 supergames. Averages until round 2 are based on 3, until round 4, 2-3 observations from each subject. All other averages are based on at least one observation per subject.

Question 2: [No Communication] *What is the benefit of pre-play communication? Does the benefit of pre-play communication depend on monitoring structure?*

Figure 1 shows that the mean cooperation level in treatments with pre-play communication is substantially higher compared to the treatments without cooperation under all three monitoring structures (Permutation Tests: all monitoring treatments, $p < 0.01$). The size of the effect is largest under perfect monitoring where the mean cooperation level increases by 66 percentage points with pre-play communication (public, 44 pp; private 54 pp). The differences in differences tests reveal that the benefit of pre-play communication does not differ between the perfect monitoring treatment and the imperfect monitoring treatments. (Permutation Tests: perfect vs. public, $p = 0.21$; perfect vs. private, $p = 0.34$). Figure 2 shows that the average cooperation level also declines over rounds without communication.

Question 3: [Pre-play vs. No Communication] *Without communication, are cooperation levels lower with imperfect monitoring as predicted by strategic uncertainty?*

There are no differences between perfect and imperfect monitoring structures without communication. Neither the average cooperation level nor the stability of cooperation over rounds is affected by the monitoring structure.

4.2 Strategy Choice

Question 4: [Strategies] *Does communication change the use of strategies? Does anticipated verbal punishment or more lenient or more forgiving behavior arise?*

In this section we report results on the use of strategies in our treatments. Table 4 shows the average memory-one Markov strategy which results from the estimation generating the highest likelihood.[20] The parameter that shows the most variability within the treatments of a monitoring structure is the initial cooperation probability $\sigma_{\emptyset}$. This finding corresponds to the higher level of cooperation in round one with communication. The differences in $\sigma_{\emptyset}$ between the treatments with and without communication are the only statistically significant differences between communication treatments on the five percent level after a Bonferroni correction for 30 comparisons. The differences in $\sigma_{\emptyset}$ between pre-play and repeated communication are small which might be due to a ceiling effect. The measure for unjustified

[20]We assume that the 5 strategies we estimate in a specific treatment condition on the same type of information and report the results for the type which generates the highest likelihood. In all treatments with perfect monitoring this is the version of the Markov strategy which conditions on (a_i, a_{-i}). In all treatments with imperfect monitoring this is the version which conditions on (a_i, ω_{-i}). This makes it possible to compare strategy shares across the imperfect monitoring treatments but not between treatments with perfect and imperfect monitoring. The 5 underlying strategies for each treatment can be found in Table 10 of Appendix B.

defection $(1 - \sigma_{cc})$ is generally lower with communication but differences are not significant. Under public monitoring $(1 - \sigma_{cc})$ is 11 percentage points lower with repeated as compared to pre-play communication. It should be noted that *cc* is by far the most frequent state in treatments with communication. Therefore the difference under public monitoring might be meaningful even though we cannot statistical differentiate $(1 - \sigma_{cc})$ between PubPre and PubRep. We do not observe a difference under private monitoring where $(1 - \sigma_{cc})$ is similar in the pre-play and repeated communication treatments.

Table 4: Average Memory-One Markov Strategies

	Perfect			Public			Private		
	No	Pre	Rep	No	Pre	Rep	No	Pre	Rep
$\sigma_{\emptyset}$	0.38 <	0.98	1.00	0.28 <	0.92	0.97	0.43 <	0.96	0.98
σ_{cc}	0.79	0.98	0.99	0.73	0.82	0.93	0.85	0.94	0.96
σ_{cd}	0.19	0.39	0.71	0.30	0.69	0.81	0.46	0.64	0.77
σ_{dc}	0.43	0.96	0.37	0.22	0.31	0.64	0.11	0.54	0.34
σ_{dd}	0.05	0.05	0.09	0.10	0.11	0.62	0.04	0.04	0.37
LL	-341	-67	-46	-339	-356	-209	-281	-221	-218

Notes: Estimates for non-constant strategy choice over the last 3 supergames. Strategies condition on action profiles in perfect treatments, and on action-signal profiles in public and private treatments. Significance based on t-test using bootstrapped two-way clustered standard errors (subject and match, 5000 repetitions). > indicates significant differences at $\alpha = 0.05$ (Boneferroni corrected for 30 comparisons). Note that mixed states (*cd* and *dc*) are observed 18 times in PerPre and 14 times in PerRep which does not produce reliable estimates.

The estimates for leniency σ_{cd} indicate more leniency as one moves from no to pre-play to repeated communication. Since the state (*cd*) is frequently observed in the communication treatments with imperfect monitoring the cooperation rate is sensitive to this parameter. At the same time, the leniency estimates for the communication treatments under perfect monitoring are not reliable due to too few observations. Leniency increases under imperfect monitoring as one moves from no to pre-play communication and further as one moves to repeated communication. For forgiving behavior we see a similar pattern with the exception that the average strategy is by 20 percentage points less forgiving with repeated communication under private monitoring compared to pre-play communication. Finally, the willingness to return to cooperation in state *dd* is substantially higher with repeated communication under imperfect monitoring but not under perfect monitoring (public, 52 pp; private 33 pp). Comparing the average memory-one Markov strategies to T1BF we can identify substantial differences. Most importantly σ_{dc} is much smaller than what can be expected based on T1BF. This does not rule out that the behavior of some subjects is well summarized by T1BF. Some strategies in Table 10 of Appendix B do exhibit more similarities (s_1 of PerNo and PerPre; s_1 of PubNo; s_1 of PrivPre, s_5 of PrivRep).

Table 5: Strategy Use across Treatments

	Perfect			Public			Private		
	No	Pre	Rep	No	Pre	Rep	No	Pre	Rep
ALLD	**0.44**	-	-	**0.64**	0.02	-	**0.52**	-	-
FC	-	-	-	-	0.08	0.02	-	-	-
GRIM	0.15	0.25	-	-	0.05	-	0.11	0.04	-
L/F	**0.36**	**0.75**	**0.98**	**0.30**	**0.56**	**0.58**	**0.23**	**0.72**	**0.67**
T1BF	0.05	-	-	0.06	0.13	0.04	**0.14**	**0.22**	**0.21**
ALLC	-	-	-	-	0.15	0.36	-	-	0.12
γ	0.07	0.01	0.01	0.09	0.10	0.05	0.06	0.05	0.05
LL	-329	-80	-59	-345	-377	-213	-284	-240	-236

Notes: Estimates for constant strategy use over the last 3 supergames. Strategies condition on action profiles in perfect treatments, and on action-signal profiles in public and private treatments. L/F contains the shares of TFT, PTFT, T2, TF2T, TF3T, T2F1T, T2F2T, LGRIM2, LGRIM3, PTFT2, DTFT, DTF2T, DTF3T, DLGRIM2, DLGRIM3, DC. T1BF is a strategy which starts with cooperation and shows the threshold memory-one belief-free response under imperfect monitoring with $\sigma_{as} = (1, 1, 0.5, 1, 0)$. Shares ≤ 0.02 not displayed (-) to increase readability. Significance based on t-test using block-bootstrapped standard errors (subject, 5000 repetitions). Boldface indicates shares significantly different from zero at the 5 percent level. γ indicates the probability of trembles. Values might not add up as expected due to rounding.

Table 5 depicts the results of a treatment-wise estimation of a standard candidate set of pure strategies (Fudenberg et al., 2012) augmented by a strategy which starts with cooperation and shows T1BF.[21] The strategies are ordered according to their potential for eliciting cooperation beginning with the least cooperative strategy ALLD. The false-cooperator (FC) is listed next as the strategy cooperates once and then defects forever. We

[21]The set of strategies by Fudenberg et al. (2012) has been used in a number of studies (reviewed in Dal Bó and Fréchette, 2017). In our implementation, these strategies condition on the same information as in the estimation of Markov strategies. Strategies which condition on actions generate the highest likelihood in the perfect monitoring treatments. Strategies which condition on the action-signal combination generate the highest likelihoods in the imperfect monitoring treatments. Note that the likelihood between the Markov and pure strategy estimations cannot be compared directly. In particular one cannot conclude from the lower likelihoods in Table 4 that these strategies are more accurate descriptions of subjects' behavior.

sum the shares of all lenient and/or forgiving strategies in the category L/F even though we estimate these strategies separately. The reason is that the actual distinct shares of these strategies are hard to identify in treatments with communication where cooperation is the norm. The only lenient and forgiving strategy which is listed separately is T1BF with initial cooperation which is not to say that it is more cooperative than the strategies summarized in L/F. In all treatments without communication, ALLD receives the mode of the shares. The substantial density attributed to ALLD explains why cooperation cannot prevail in treatments without communication. The remaining density is attributed to cooperative strategies such as GRIM or lenient or forgiving strategies which occasionally form cooperative relationships. With pre-play communication, subjects switch to cooperative strategies. Under perfect monitoring the combinations of these strategies are sufficient to generate high levels of cooperation since the estimated strategies condition on actions and are not affected by the noise. Under imperfect monitoring, we see that only little density is attributed to GRIM. If communication is repeated, virtually all strategies are lenient and/or forgiving and even ALLC receives positive shares. The strategy with T1BF and $\sigma_{\emptyset} = 1$ receives some density under imperfect monitoring. The shares are significant in all communication treatments under private monitoring.

One drawback of the SFEM using a candidate set of pure strategies is that the estimated strategy shares might not be robust to the inclusion of other strategies to the set. We check the robustness of the SFEM results based on the classification method proposed by Camera et al. (2012). The approach estimates the total number frequency of supergames a strategy predicts well. Assume the a supergame has x rounds and y errors are observed. A strategy does predict the supergame well if the probability of observing n or more errors is smaller than p. The classification method with $p = 0.05$ validates the results of the SFEM and shows that the N of lenient and/or forgiving strategies increases as one moves from no to pre-play to repeated communication (see Table 13 in Appendix B).

4.3 Communication

Our two raters made an average of 2.65 classifications per subject-round observation resulting in 18678 and 18984 classifications in total. With 72 sub-categories the probability of agreement with random classification based on the average number of classifications is approximately 0.15. The true probability of agreement of our raters is 0.58 (Cohen's $\kappa = 0.5$). This is remarkable if one considers the size of the task and the fact that the scope of some sub-categories clearly overlap.

Question 5: [Communication Content] *What do subjects communicate about?*

Table 6 depicts the relative frequency of five categories we use to summarize the content of participants' communication across treatments. The category coordination includes all attempts being made by subjects to coordinate behavior in future rounds. The category also includes implicit or explicit announcements of choices since such announcements could also be used to coordinate behavior. According to our two independent raters, the category applies in the vast majority of subject-round observations of the pre-play phase. Its role in the repeated communication treatments is rather minor which suggests that coordination is predominantly important before the first round.[22] The category Deliberation includes all instances where subjects discuss choices or strategies. Our raters indicate content related to deliberation in roughly every second subject-round observation with pre-play communication. Content related to deliberation also becomes less frequent after round one. All content which concerns the relationship of a matched pair of subjects is included in the category Relationship. The category also covers motivational talk and positive feedback which we find to be quite common among paired subjects. Content related to this category is more frequent under imperfect monitoring. In contrast to the categories Coordination and Deliberation, the category Relationship does not become less frequent after round one. The category Information contains all statements which contain reports of action, signals or payoffs from the current supergame. It is not possible to report such information before round one. The category ranges among the most frequent with repeated communication. Note that content related to the Information category does not imply that subjects reveal information. In order to asses whether the report private information we will look at data from sub-category level in the following. Our last category *Trivia* contains content which is off-topic or classified as small talk by our raters. In contrast to the Relationship category the content does not have a obvious relation to the game. The Trivia category is always among the most frequent in all treatments. Looking at the sub-category level, we observe only four instances of verbal punishment. Two are from the same subject who punishes in two consecutive rounds under public monitoring. The other two are from different subjects under private monitoring.

Question 6: [Private Information Sharing] *Do subjects share private information?*

Table 7 documents the degree to which private information is exchanged under repeated communication. Under public monitoring this concerns the actions which cannot be observed by the other player. The right columns show that an action is reported in only eight percent of all subject-round observations after round one. The vast majority of reports indicate cooperation in the last round which is true 94% of all cases. Table 7 also lists the frequency

[22] In fact, the frequencies of the categories in round one of the repeated communication treatments exhibit a similar pattern as those of pre-play communication.

Table 6: Frequency of Codings per Individual-Round Observation

	Perfect		Public		Private	
	Pre	Rep	Pre	Rep	Pre	Rep
Coordination	0.96	0.09	0.95	0.19	0.94	0.24
Deliberation	0.47	0.05	0.57	0.08	0.47	0.07
Relationship	0.11	0.19	0.21	0.31	0.27	0.29
Information	–	0.20	–	0.36	–	0.37
Trivia	0.94	0.47	0.86	0.33	0.73	0.46

Notes: Level of the analysis are individual-round observations. Data based on the last 3 supergames. A coding is considered as valid if both raters indicated the same sub-category for a subject-round observation.

Table 7: Exchange of Private Information

	Private		Public	
	p(report)	p(true)	p	p(true)
Actions				
Report of action	0.15	0.93	0.08	0.94
Report of C	0.15	0.93	0.08	0.95
Report of D	0.00	1	0.01	0.86
Report of C if $\omega_i = d$	0.31	0.75	0.15	0.84
Signals				
Report of signal	0.37	0.96	-	-
Report of c	0.27	0.99	-	-
Report of d	0.10	0.86	-	-
Report of d if $\omega_{-i} = d$	0.45	-	-	-

Notes: Frequencies of coding in all subject-round observations after round one of the last three supergames. A coding is considered as valid if both raters indicated the same sub-category for a subject-round observation. Values might not add up as expected due to rounding.

of C reports if the signal was defection. In 15% of these cases this is followed by a report of C (truthful in 84 %). The left columns of Table 7 show that a similar pattern exists under private monitoring but the frequency of action reports double. One important difference concerns the interpretation of reporting C if the signal says differently. Under private monitoring the difference compared to the baseline frequency of C reports suggests that players do report the d signal in the first place. This indirect evidence is supported by the values in the lower part of the table. A signal is reported in 37% of all subject-round interaction after round one. Most of the reports truthfully reveal a c signal. Defection signals are reported in 10% of all subject round interactions. To put this value into perspective the last line shows the frequency of d reports in response to an d signal which is 0.45. To sum up the results of Table 7, we can say that subjects make use of repeated communication to exchange private information. Actions are communicated less often than signals but both reports are usually credible.

Question 7: [Content-Cooperation Correlations] *What are the correlations between different forms of communication content and cooperation?*

Table 8: Communication and First Round Cooperation

	all Pre	all Rep
Coordination	0.22***	0.07*
Deliberation	0.02	0.03*
Relationship	0.10	0.06
Information	–	–
Trivia	0.05*	0.07***
Supergame	0.05***	0.04***

Notes: Marginal effects from logistic regressions for cooperation in round 1 including all supergames. All models control for socio-demographic and other subject related characteristics. Significance based on t-test using bootstrapped standard errors, two-way clustered on subject and match (1000 repetitions). $***$ ($**$,$*$) indicates significance on the 1 (5,10) percent level.

Table 8 shows the marginal effects from logistic regressions, in which cooperation in the first round of a supergame is regressed on dummies indicating whether pre-play communication included messages falling into any of 4 of our 5 main categories and controlling for the supergame played (all seven supergames are considered here). The Information category is left out as there is no information to share before round 1. Note that the coefficients reflect correlations as the communication content is endogenous. Thus they can only be suggestive of causal effects. Both communication on trivia and efforts to coordinate behavior are positively correlated with cooperation in round 1 in all pre-play and all repeated communication treatments. Table 9 shows the marginal effects from logistic regressions from the last three supergames, in which cooperation in later rounds is regressed on dummies indicating whether pre-play communication included messages falling into any of the 5 main categories and controlling for the supergame played, the round of the supergame, the last own action and the last signal received. Interestingly, most marginal effects are not statistically different from zero, with the exceptions of Coordination and Information under private monitoring.

Table 9: Communication Content and Cooperation in Rounds > 1

	Repeated Treatments		
	Perfect	Public	Private
Coordination	−0.006	0.003	0.030*
Deliberation	−0.032	−0.005	−0.005
Relationship	0.001	−0.021	0.007
Information	−0.001	0.016	0.027*
Trivia	0.002	0.066	0.033
Supergame	−0.003	0.015	0.017*
Round	0.000	−0.004	−0.002
Last Action	0.049***	0.226***	0.159***
Last Signal	0.014***	0.105***	0.149***

Notes: Marginal effects from logistic regressions for cooperation in rounds > 1. Last 3 supergames. All models control for socio-demographic and other subject related characteristics. Significance based on t-test using bootstrapped standard errors, two-way clustered on subject and match (1000 repetitions). $***$ ($**$,$*$) indicates significance on the 1 (5,10) percent level.

5 Discussion

Our results demonstrate that communication can have an enormous impact on cooperation and its stability in indefinitely repeated interactions. They further highlight the different roles communication plays and how this depends on the monitoring structure. In the following, we summarize our main theoretical and empirical results and discuss their implications.

We study communication in the indefinitely repeated prisoner's dilemma with imperfect monitoring. To characterize its theoretical properties with respect to strategic uncertainty, we extend existing work by Dal Bó and Fréchette (2011), Blonski et al. (2011), and Breitmoser (2015), which only applies to the perfect monitoring case. We derive two new measures of strategic uncertainty for the imperfect public and private monitoring cases and characterize the threshold above which memory-one belief-free equilibria (m1BF) exist. For the experiment, we choose parameters that make cooperation more risky under imperfect monitoring than under perfect monitoring and under which cooperation seems to be unlikely without communication. Our continuation probability $\delta = 0.8$ coincides with the m1BF-threshold at which a unique m1BF equilibrium exists under imperfect monitoring.[23] Moreover, *truthful communication* equilibria under private monitoring also exist with these parameters.

These design choices allow us to address a number of important open empirical questions. First, we test whether repeated communication helps cooperation under private monitoring, as suggested by Heller's (2017) stability criterion and recent case studies of cartels. Indeed, cooperation is more stable with repeated communication. However, this result can only partially be explained by private information sharing. Only 37% of all private signals are reported. In addition, coordination and trust-building via talk about trivia appear to play a role. Second, we find that pre-play communication is also very effective in increasing cooperation rates. Apparently, pre-play communication decreases strategic uncertainty and makes subjects choose cooperative strategies more frequently. Most subjects use the communication opportunity to coordinate on CC in the first period of the game. However, subjects do not coordinate on complex efficient equilibria. This might to some extent explain the stronger decrease in cooperation rates over time as compared to the repeated communication treatments under imperfect public and private monitoring, and the stronger difference in this decrease between the imperfect and the perfect monitoring treatments. Under perfect monitoring cooperation rates stay on a high level even in the absence of repeated communication opportunities. The high cooperation rate under private monitoring with pre-play communication casts some doubt on the predictive power of Heller's (2017) stability criterion. Third, we complement Aoyagi et al.'s (2017) results on cooperation rates without communication under the three monitoring structures. While their parametrization is characterized by low and similar levels of strategic uncertainty, our levels are higher and differ substantially between the perfect and imperfect monitoring cases. While our cooperation rates are indeed lower than theirs, somewhat surprisingly we do not find a difference between the monitoring structures either. This suggests that the levels of strategic uncertainty above

[23] It is unique up to the probability of cooperation in the first period, which is always a free parameter.

which cooperation becomes unlikely might be lower under perfect than under imperfect monitoring. Further studies that test the predictive power of our new measures could help to shed more light on this.

Figure 3: Simulated Cooperation Levels over 20 Rounds

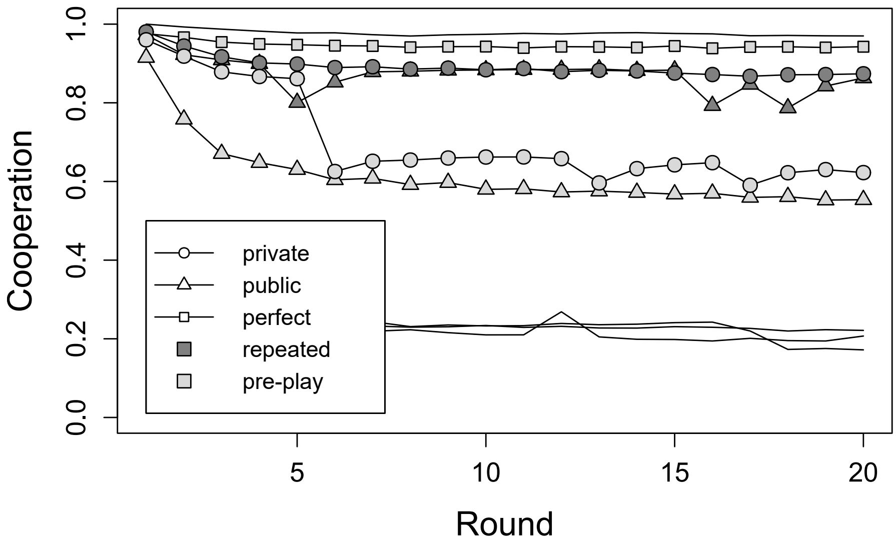

Notes: Graph depicts the results of a simulation of cooperation levels over rounds based on the memory-one Markov estimations. Lines without dots indicate the cooperation level in the treatments without communication and in PerRep.

To get an idea of how cooperation rates might evolve in longer interactions than those we observe in the laboratory, we conduct the following simulation exercise. We explore the stability of cooperation in our treatments up to 20 rounds which would in expectancy be observed in approximately 1 out of 100 supergames. To predict the evolution of cooperation in each treatment, we take the corresponding memory-one Markov estimates, draw a random sample of 100 individuals based on the shares, and simulate 50 interactions over 20 rounds. We repeat this procedure 1000 times to average out the variability in cooperation due to the random sampling. Figure 3 depicts the average level of cooperation over the total of 1000×50 simulated interactions. Cooperation levels up to round 11 are similar to those actually observed in the laboratory. The estimated strategies reproduce the decay of cooperation with pre-play communication in the imperfect monitoring treatments. Even the delayed decline in the private treatment is reflected in the memory-one Markov strategies. The main finding from the simulated data is that cooperation levels do not change much after round 11. This suggests that very high cooperation rates can be expected with repeated communication, even in supergames with up to twenty rounds.

To gain a better understanding of the mechanisms behind these results, we study strategy choices using and extending Dal Bó and Fréchette's (2011) strategy frequency estimation method (SFEM), and analysing the content of communication and its correlations with behavior in the game. We find that subjects' play becomes both more lenient and more forgiving with communication, and that the m1BF-strategy which is a best response to itself in the most-efficient m1BF-equilibrium of our game appears to be played by a substantial number of subjects under private monitoring. Further, we find that communication falling into any of the three main categories of communication content - information sharing, coordination, and trivia - is positively correlated with cooperation. This final result underscores that the role of communication is broader than suggested by recent theory - just as the case studies reviewed in the introduction suggested. However, if communication was limited to mere sales reports, the information sharing role would be more salient and things might look different. Future studies could shed more light onto such environments with structured communication.

References

Andersson, O. and E. Wengström (2007): "Do Antitrust Laws Facilitate Collusion? Experimental Evidence on Costly Communication in Duopolies," *Scandinavian Journal of Economics*, 109, 321–339.

——— (2012): "Credible communication and cooperation: Experimental evidence from multi-stage Games," *Journal of Economic Behavior and Organization*, 81, 207–219.

Aoyagi, M., V. Bhaskar, and G. R. Frechette (2017): "The Impact of Monitoring in Infinitely Repeated Games: Perfect, Public, and Private," *Mimeo.*

Arechar, A., A. Dreber, D. Fudenberg, and D. Rand (2017): "I'm just a soul whose intentions are good: The role of communication in noisy repeated games," *Mimeo*, 06511.

Awaya, Y. and V. Krishna (2016): "On Communication and Collusion," *American Economic Review*, 106, 285–315.

Ben-Porath, E. and M. Kahneman (1996): "Communication in Repeated Games with Private Monitoring," *Journal of Economic Theory*, 70, 281–297.

Bénabou, R. and J. Tirole (2006): "Incentives and prosocial behavior," *American Economic Review*, 96, 1652–1678.

Bhaskar, V., G. J. Mailath, and S. Morris (2008): "Purification in the infinitely-repeated prisoners' dilemma," *Review of Economic Dynamics*, 11, 515–528.

Bhaskar, V. and I. Obara (2002): "Belief-based equilibria in the repeated prisoners' dilemma with private monitoring," *Journal of Economic Theory*, 102, 40–69.

Blonski, M., P. Ockenfels, and G. Spagnolo (2011): "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence," *American Economic Journal: Microeconomics*, 3, 164–192.

Breitmoser, Y. (2012): "Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma," *Working Paper.*

——— (2015): "Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma," *American Economic Review*, 105, 2882–2910.

Camera, G., M. Casari, and M. Bigoni (2012): "Cooperative strategies in anonymous economies: An experiment," *Games and Economic Behavior*, 75, 570–586.

——— (2013): "Experimental markets with frictions," *Journal of Economic Surveys*, 27, 536–553.

Charness, G. and M. Dufwenberg (2006): "Promises and partnership," *Econometrica*, 74, 1579–1601.

Compte, O. (1998): "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66, 587–626.

Compte, O. and A. Postlewaite (2015): "Plausible cooperation," *Games and Economic Behavior*, 91, 45–59.

Cooper, D. J. and K.-U. Kühn (2014a): "Communication, renegotiation, and the scope for collusion," *American Economic Journal: Microeconomics*, 6, 247–278.

——— (2014b): "Cruel to be Kind: An Experimental Study of Communication, Collusion, and Threats to Punish," *Mimeo.*

Cooper, R. W., D. V. DeJong, R. Forsythe, and T. W. Ross (1992): "Communication in Coordination Games," *Quarterly Journal of Economics*, 107, 739–771.

Crawford, V. and J. Sobel (1982): "Strategic information transmission," *Econometrica*, 50, 1431–1451.

Dal Bó, P. (2005): "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games," *The American Economic Review*, 95, 1591–1604.

Dal Bó, P. and G. Fréchette (2011): "The evolution of cooperation in infinitely repeated games: Experimental evidence," *American Economic Review*, 101, 411–429.

Dal Bó, P. and G. R. Fréchette (2017): "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey," *Journal of Economic Literature*, forthc.

Dana, J., R. A. Weber, and J. X. Kuang (2007): "Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness," *Economic Theory*, 33, 67–80.

Ellingsen, T. and R. Östling (2010): "When Does Communication Improve Coordination?" *American Economic Review*, 100, 1695–1724.

Ely, J. C. and J. Välimäki (2002): "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, 102, 84–105.

Embrey, M., G. Fréchette, and E. Stacchetti (2013): "An Experimental Study of Imperfect Public Monitoring: Efficiency Versus Renegotiation-Proofness," *Mimeo*.

Farrell, J. and E. Maskin (1989): "Renegotiation in repeated games," *Games and Economic Behavior*, 1, 327–360.

Fischbacher, U. (2007): "Z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10, 171–178.

Fischbacher, U. and V. Utikal (2013): "On the acceptance of apologies," *Games and Economic Behavior*, 82, 592–608.

Fonseca, M. a. and H. T. Normann (2012): "Explicit vs. tacit collusion-The impact of communication in oligopoly experiments," *European Economic Review*, 56, 1759–1772.

Fréchette, G. R. and S. Yuksel (2017): "Infinitely repeated games in the laboratory: four perspectives on discounting and random termination," *Experimental Economics*, forthc.

Fudenberg, D., D. G. Rand, and A. Dreber (2012): "Slow to Anger and Fast to Forgive: Cooperation in an Uncertain World," *American Economic Review*, 102, 720–749.

Genesove, D. and W. P. Mullin (2001): "Rules, communication, and collusion: Narrative evidence from the sugar institute case," *American Economic Review*, 91, 379–398.

Green, E. J. and R. H. Porter (1984): "Noncooperative Collusion under Imperfect Price Information," *Econometrica*, 52, 87.

Greiner, B. (2015): "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 1, 114–125.

Harrington, J. E. and A. Skrzypacz (2011): "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices," *American Economic Review*, 101, 2425–2449.

Harsanyi, J. C. and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*, Cambridge, MA: MIT Press.

Heller, Y. (2017): "Instability of belief-free equilibria," *Journal of Economic Theory*, 168, 261–286.

Kalai, E., D. Samet, and W. Stanford (1988): "A note on reactive equilibria in the discounted prisoner's dilemma and associated games," *International Journal of Game Theory*, 17, 177–186.

Kandori, M. and H. Matsushima (1998): "Private Observation, Communication and Collusion," *Econometrica*, 66, 627–652.

Larson, T. and C. M. Capra (2009): "Exploiting moral wiggle room: Illusory preference for fairness? A comment," *Judgment and Decision Making*, 4, 467–474.

Levenstein, M. C. and V. Y. Suslow (2006): "What Determines Cartel Success?" *Journal of Economic Literature*, 44, 43–95.

Mailath, G. J. and L. Samuelson (2006): *Repeated Games and Reputations*, Oxford University Press: Oxford.

Matsushima, H. (1991): "On the theory of repeated games with private information. Part II: revelation through communication," *Economics Letters*, 35, 257–261.

Obara, I. (2009): "Folk Theorem with Communication," *Journal of Economic Theory*, 144, 120–134.

Pearce, D. G. (1987): "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation," *Cowles Foundation Discussion Paper.*

Piccione, M. (2002): "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, 102, 70–83.

Rabin, M. (1994): "A Model of Pre-game Communication," *Journal of Economic Theory*, 63, 370–391.

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, 44–61.

Utikal, V. (2012): "A fault confessed is half redressed - Confessions and punishment," *Journal of Economic Behavior & Organization*, 81, 314–327.

van Damme, E. (1989): "Renegotiation-proof equilibria in repeated prisoners' dilemma," *Journal of Economic Theory*, 47, 206–217.

Vespa, E. and A. J. Wilson (2016): "Information Transmission under the shadow of the future: An experiment," *Mimeo.*

Ye, Y. (1987): "Interior algorithms for linear, quadratic, and linearly constrained non linear programming," *Ph.D. thesis, Department of EES Stanford University, Stanford CA.*

Appendix A: Belief-free Equilibria

Depending on the monitoring structure, different versions of memory-one belief-free strategies exist. We consider three cases: (1) m1BF strategies which condition on (a_i, a_{-i}), (2) m1BF strategies which condition on (ω_i, ω_{-i}), and (3) m1BF strategies which condition on (a_i, ω_{-i}). Under perfect monitoring, all three cases are possible. Under public monitoring, only cases 2 and 3 are possible while case 3 is the only possible case under private monitoring. Proposition 0 states that there exists a δ for all cases such that all equilibrium strategies of that case exhibit the same T1BF pattern. The extensions of the BOS threshold to pubic signals and action-signal combinations are defined in Propositions 2.2 and 3.2.

Proposition 0. *For all three cases there exists a value of δ such that all memory-one belief-free equilibrium strategies of that particular case have the same cooperation probabilities for all nonempty memory-one histories. The strategies are characterized by*

$$\sigma = \left(\sigma_{\emptyset}, 1, max\left\{0, \left(1 - \frac{g}{l}\right)\right\}, min\left\{1, \frac{l}{g}\right\}, 0\right)$$

in cases one and three, and

$$\sigma = \left(\sigma_{\emptyset}, 1, max\left\{0, \left(1 - \frac{g - \epsilon(g+l)}{l - \epsilon(g+l)}\right)\right\}, min\left\{1, \frac{l - \epsilon(g+l)}{g - \epsilon(g+l)}\right\}, 0\right)$$

in case two.

Proof of Proposition 0. The proof results from the proofs of Propositions 1.1, 2.1 and 3.1. □

Actions

Lemma 1.1. *If strategies condition on actions, the existence condition for symmetric memory-one belief-free equilibria depends on the larger of the two values g and l. Let ϕ denote the larger of the two values. The existence condition is:*

$$\delta^{BF} = \frac{\phi}{1+\phi} \tag{1}$$

Since g and l are both positive values these equilibria exist for high enough values of δ. Note that if $g \geq l$ the δ threshold corresponds to the one for cooperative subgame-perfect equilibria of the repeated game with perfect monitoring. However, if $l > g$ as in our case, the conditions differ with $\delta^{BF} > \delta^{SPE}$. The condition applies for belief-free equilibria in reactive strategies (Kalai et al., 1988) which condition on the other player's action and require $g = l$ which yields $\delta^{BF} = \delta^{SPE}$.

Lemma 1.2. *Above the threshold, a two-dimensional manifold of memory-one belief-free equilibria exists given by*

$$\sigma_{cd} = \sigma_{cc} + \left(\sigma_{cc} - \sigma_{dd} - \frac{1}{\delta}\right) g \tag{2}$$

and

$$\sigma_{dc} = \sigma_{dd} - \left(\sigma_{cc} - \sigma_{dd} - \frac{1}{\delta}\right) l \tag{3}$$

Proposition 1.1. *For $\delta = \delta^{BF}$ all memory-one belief-free equilibrium strategies have the same cooperation probabilities after nonempty memory-one histories and are $\sigma = (\sigma_\emptyset, 1, (1-g/l), 1, 0)$ if $l > g$, $\sigma = (\sigma_\emptyset, 1, 0, (l/g), 0)$ if $g > l$ and $\sigma = (\sigma_\emptyset, 1, 0, 1, 0)$ if $g = l$.*

Proof of Proposition 1. Let $V^{a_i}_{a_j a_i}$ denote player i's expected payoff for playing a_i if player j observed the action profile $\{a_j, a_i\}$ in the previous round (we say player j is in state $a_j a_i$). If $\sigma_{a_i a_j}$ denotes the probability to play c for any player i after $\{a_i, a_j\}$, we have:

$$V^c_{aa} = (1-\delta)(\sigma_{aa} - (1-\sigma_{aa})l) + \delta(\sigma_{aa} V_{cc} + (1-\sigma_{aa})V_{dc}) \tag{4}$$

$$V^d_{aa} = (1-\delta)(\sigma_{aa}(1+g) + (1-\sigma_{aa})0) + \delta(\sigma_{aa} V_{cd} + (1-\sigma_{aa})V_{dd}) \tag{5}$$

Following Bhaskar et al. (2008), we derive conditions for V_{cd} and V_{cc} which assure the strategies are belief-free, i.e. for any $\sigma_{aa} \in (0,1)$, player i is indifferent between playing c or d independent of player j's state. Subtracting (5) from (4) gives:

$$0 = \sigma_{aa}\left\{(1-\delta)(l-g) + \delta\left(V_{cc} - V_{cd} - V_{dc} + V_{dd}\right)\right\} - (1-\delta)l + \delta\left(V_{dc} - V_{dd}\right)$$

The equation holds independent of σ_{aa} if the terms in curly brackets and the last part are both zero. Solving the the condition resulting from the last part for $V_{dc} - V_{dd}$ and inserting the solution into the condition derived from the terms in curly brackets gives

$$V_{cc} = V_{cd} + \frac{(1-\delta)g}{\delta}$$

and

$$V_{dc} = V_{dd} + \frac{(1-\delta)l}{\delta}$$

Solving (4) for σ_{cc} using the condition on V_{dc} above and rearranging for V_{cc} yields

$$V_{cc} = \frac{(1-\delta)\sigma_{cc} + \delta(1-\sigma_{cc})V_{dd}}{1-\delta\sigma_{cc}}$$

Solving (4) for σ_{dd} using the condition on V_{cd} and V_{cc} above gives

$$V_{dd} = \frac{\sigma_{dd}}{1+\delta\sigma_{dd}-\delta\sigma_{cc}}$$

Now, all V_{aa} can be eliminated from (4) solved for σ_{dd} and σ_{dc} this yields (2) and (3) which proofs *Lemma 1.2.* Note that $\partial\sigma_{cd}/\partial\delta > 0$, $\partial\sigma_{cd}/\partial\sigma_{cc} > 0$ and $\partial\sigma_{cd}/\partial\sigma_{dd} < 0$. The question is, how big δ must be at least in order to assure that $\sigma_{cd} \geq 0$ if $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$. Inserting these values into (2) and rearranging gives $\delta > \delta^{BF}$ with $\phi = g$. Note that $\sigma_{cd} \leq 1$ is true even if $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ for all feasible values of δ, g and l. At the same time $\partial\sigma_{dc}/\partial\delta < 0$, $\partial\sigma_{dc}/\partial\sigma_{cc} < 0$ and $\partial\sigma_{dc}/\partial\sigma_{dd} > 0$. The question here is, how big δ must be at least in order to assure that $\sigma_{dc} \leq 1$ if $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$. Inserting these values into (3) and rearranging gives $\delta > \delta^{BF}$ with $\phi = l$. At the same time, $\sigma_{dc} \geq 0$ true even if $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ for all feasible values of δ, g and l. Hence, the larger of the values g and l imposes the stricter condition on δ which proofs Lemma 1.1. To complete the proof, insert (1) together with $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ into (2) and (3) to obtain the structure of the T1BF response defined by g and l. □

For the sake of completeness, we highlight the proof of the δ threshold for semi-GRIM equilibria which is due to Breitmoser (2015). We use (2) and (3) and set $\sigma_{cd} = \sigma_{dc}$ which gives

$$\sigma_{dd} = \sigma_{cc} - \frac{g+l}{\delta(1+g+l)} \tag{6}$$

and

$$\sigma_{cd} = \sigma_{dc} = \sigma_{cc} - \frac{g}{\delta(1+g+l)} \tag{7}$$

Note that $\sigma_{dd} < \sigma_{cd} < 1$ for $\sigma_{cc} \in (0,1)$ and any $\delta \in (0,1)$. For existence σ_{dd} must be positive solving which can be rearranged to yield the BOS threshold. N

$$\delta^{SG} = \frac{g+l}{1+g+l} \tag{8}$$

ote that the condition on δ is always stricter then the condition on δ which results from $\sigma_{cd} = \sigma_{dc} \geq 0$ which is $g/(1+g+l)$. Above the BOS threshold a continuum $\sigma_{cc} \in (\frac{g+l}{\delta(1+g+l)}, 1)$

of memory one belief-free equilibria in semi-grim strategies exists satisfying Note that the condition for semi grim equilibria is a mixture of the two possible conditions based on the different values of ϕ with equal weight on g and l as required by axiom 5 in Blonski et al. (2011) while (1) gives full weight on the larger of the two values.

Public signals

Lemma 2.1. *If strategies condition on the ϵ-noisy public signals, the existence condition for symmetric memory-one belief-free equilibria depends on the larger of the two values g and l. Let ϕ denote the larger and ψ the smaller of the two values. The existence condition is:*

$$\delta^{BF} = \frac{(1-\epsilon)\phi - \epsilon\psi}{(1-2\epsilon)(1-2\epsilon+(1-\epsilon)\phi - \epsilon\psi)} \tag{9}$$

In contrast to result for actions, combinations of the parameters g, l and ϵ exists for which $\delta^{BF} > 1$.

Lemma 2.2. *Above the threshold, a two-dimensional manifold of memory-one belief-free equilibria exists given by*

$$\sigma_{cd} = \sigma_{cc} + \frac{\sigma_{cc} - \sigma_{dd} - \frac{1}{\delta(1-2\epsilon)}}{1-2\epsilon}((1-\epsilon)g - \epsilon l) \tag{10}$$

and

$$\sigma_{dc} = \sigma_{dd} - \frac{\sigma_{cc} - \sigma_{dd} - \frac{1}{\delta(1-2\epsilon)}}{1-2\epsilon}((1-\epsilon)l - \epsilon g) \tag{11}$$

Proposition 2.1. *For $\delta = \delta^{BF}$ all memory-one belief-free equilibrium strategies have the same cooperation probabilities after nonempty memory-one histories and are $\sigma = (\sigma_\emptyset, 1, (1-g/l), 1, 0)$ if $l > g$, $\sigma = (\sigma_\emptyset, 1, 0, (l/g), 0)$ if $g > l$ and $\sigma = (\sigma_\emptyset, 1, 0, 1, 0)$ if $g = l$.*

Proof of Proposition 2.1. The proof follows the same steps as for actions. Let $V^{a_i}_{s_j s_i}$ denote player i's expected payoff for playing a_i if player j observed $\{s_j, s_i\}$ in the previous round (which means player j is in state $s_j s_i$). If $\sigma_{s_i s_j}$ denotes the (universal) probability of player i

to play c after $\{s_i, s_j\}$, we get:

$$\begin{aligned}V_{ss}^c &= (1-\delta)(\sigma_{ss} - (1-\sigma_{ss})l) + \delta[(1-\epsilon)(\sigma_{ss}(1-\epsilon) + (1-\sigma_{ss})\epsilon)V_{cc} \\ &\quad +\epsilon(\sigma_{ss}(1-\epsilon) + (1-\sigma_{ss})\epsilon)V_{cd} \\ &\quad +(1-\epsilon)(\sigma_{ss}\epsilon + (1-\sigma_{ss})(1-\epsilon))V_{dc} \\ &\quad +\epsilon(\sigma_{ss}\epsilon + (1-\sigma_{ss})(1-\epsilon))V_{dd}] \qquad (12)\\ V_{ss}^d &= (1-\delta)(\sigma_{ss}(1+g) + (1-\sigma_{ss})0) + \delta[\epsilon(\sigma_{ss}(1-\epsilon) + (1-\sigma_{ss})\epsilon)V_{cc} \\ &\quad +(1-\epsilon)(\sigma_{ss}(1-\epsilon) + (1-\sigma_{ss})\epsilon)V_{cd} \\ &\quad +\epsilon(\sigma_{ss}\epsilon + (1-\sigma_{ss})(1-\epsilon))V_{dc} \\ &\quad +(1-\epsilon)(\sigma_{ss}\epsilon + (1-\sigma_{ss})(1-\epsilon))V_{dd}] \qquad (13)\end{aligned}$$

Again we derive conditions for V_{cd} and V_{cc} which together assure the belief-free property following Following Bhaskar et al. (2008), i.e. for any $\sigma_{ss} \in (0,1)$, player i is indifferent between playing c or d independent of player j's state. First, subtracting (13) from (12) gives:

$$\begin{aligned}0 &= \sigma_{ss}\left\{(1-\delta)(l-g) + \delta\left((1-2\epsilon)^2V_{cc} - (1-2\epsilon)^2V_{cd} - (1-2\epsilon)^2V_{dc} + (1-2\epsilon)^2V_{dd}\right)\right\} \\ &\quad -(1-\delta)l + \delta\left((1-2\epsilon)\epsilon V_{cc} - (1-2\epsilon)\epsilon V_{cd} + (1-2\epsilon)(1-\epsilon)V_{dc} - (1-2\epsilon)(1-\epsilon)V_{dd}\right)\end{aligned}$$

Note that he expression holds independent of σ_{ss} if the terms in curly brackets and the terms in the second line are both zero. Solving the the condition on the second line for $V_{dc} - V_{dd}$ and inserting into the other condition gives

$$V_{cc} = V_{cd} + \frac{(1-\delta)((1-\epsilon)g - \epsilon l)}{\delta(1-2\epsilon)^2}$$

and

$$V_{dc} = V_{dd} + \frac{(1-\delta)((1-\epsilon)l - \epsilon g)}{\delta(1-2\epsilon)^2}$$

Solving (12) for σ_{cc} and rearranging for V_{cc} yields

$$V_{cc} = \frac{(1-\delta)(\sigma_{cc} - l) + \delta((1-\epsilon-\sigma_{cc}(1-2\epsilon))V_{dd} + \frac{(1-\delta)((1-\epsilon)l-\epsilon g)}{(1-2\epsilon)} + \frac{(1-\delta)(g-l)}{(1-2\epsilon)^2}}{1-\delta(\sigma_{cc}(1-2\epsilon)+\epsilon)}$$

and solving (12) for σ_{dd} and inserting V_{cc} yields

$$V_{dd} = \frac{\sigma_{dd}\left(1 - \delta\epsilon(1-\epsilon)\frac{g-l}{1-2\epsilon}\right) + (1 - \delta\sigma_{cc}(1-2\epsilon))\left(\epsilon(1-\epsilon)\frac{l-g}{(1-2\epsilon)^2}\right) + \delta\epsilon(\sigma_{cc} - \sigma_{dd})}{1 + (1-2\epsilon)\delta\sigma_{dd} - (1-2\epsilon)\delta\sigma_{cc}}$$

Now, all V_{as} can be eliminated from (12) solved for σ_{dd} and σ_{dc} which proofs Lemma 2.2. For existence we need to assure that $\sigma_{cd} \in (0,1)$ and $\sigma_{dc} \in (0,1)$ for a feasible combination of values σ_{cc}, σ_{dd} and δ. First assume $(1-\epsilon)\psi - \epsilon\phi > 0$ and consider σ_{cd} (note that $(1-\epsilon)\phi - \epsilon\psi > 0$ always holds for $\epsilon < 0.5$). In this case $\partial\sigma_{cd}/\partial\sigma_{cc} > 0$ and $\partial\sigma_{cd}/\partial\sigma_{dd} < 0$. Note that $\sigma_{cd} \leq 1$ for any $\delta \in (0,1)$ even if $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$. To establish $\sigma_{cd} \geq 0$ we use $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$. Solving for δ shows gives the condition $\delta > \delta^{BF}$ with $\phi = g$. Next, we consider σ_{dc} still assuming $(1-\epsilon)\psi - \epsilon\phi > 0$. Hence $\partial\sigma_{dc}/\partial\sigma_{cc} < 0$ and $\partial\sigma_{dc}/\partial\sigma_{dd} > 0$. Again $\sigma_{dc} \geq 0$ for any $\delta \in (0,1)$ even if $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$. To establish $\sigma_{dc} \leq 1$ we use $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ which gives $\delta > \delta^{BF}$ with $\phi = l$. Therefore, if $(1-\epsilon)\psi - \epsilon\phi > 0$ the stricter condition on δ results from the larger of the two values g or l as in (9). Note that $(1-\epsilon)\psi - \epsilon\phi < 0$ imply smaller conditions on δ. This proofs Lemma 2.1. To complete the proof, insert (9) together with $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ into (10) and (11) to obtain the structure of the T1BF response defined by g and l. □

Proposition 2.2. *If players condition on the ϵ-noisy public signals, the existence condition for semi-GRIM equilibria is:*

$$\delta^* = \frac{g+l}{(1-2\epsilon)(1+g+l)} \tag{14}$$

Above this threshold, a continuum $\sigma_{cc} \in (\frac{g+l}{\delta(1-2\epsilon)(1+g+l)}, 1)$ of semi-grim equilibria exists given by:

$$\sigma_{dd} = \sigma_{cc} - \frac{g+l}{\delta(1-2\epsilon)(1+g+l)} \tag{15}$$

and

$$\sigma_{cd} = \sigma_{dc} = \sigma_{cc} - \frac{g}{\delta(1-2\epsilon)(1+g+l)} \tag{16}$$

Proof of Proposition 2.2. Using the semi-grim property $\sigma_{cd} = \sigma_{dc}$ for (10) and (11) yields (15) and (16). Observe that $\sigma_{dd} < \sigma_{cd} < 1$ for $\sigma_{cc} \in (0,1)$ and for existence σ_{dd} must be positive which can be rearranged to yield (14). □

Action-signal combinations

Lemma 3.1. *If players condition on their own action and the ϵ-noisy signal of the other player's action, the existence condition for symmetric memory-one belief-free equilibria also depends on the larger of the two values g and l. Let ϕ denote the larger of the two values and ψ the smaller of the two. The existence condition is:*

$$\delta^{BF} = \frac{\phi}{1 - 2\epsilon - \epsilon\psi + (1-\epsilon)\phi} \tag{17}$$

If $g = l$ the condition is the same as for private signals.

Lemma 3.2. *Above the threshold, a two-dimensional manifold of memory-one belief-free equilibria exists given by*

$$\sigma_{cd} = \sigma_{cc} + \frac{\sigma_{cc} - \sigma_{dd} - \frac{1}{\delta}}{1 - 2\epsilon - \epsilon(g+l)} g \tag{18}$$

and

$$\sigma_{dc} = \sigma_{dd} - \frac{\sigma_{cc} - \sigma_{dd} - \frac{1}{\delta}}{1 - 2\epsilon - \epsilon(g+l)} l \tag{19}$$

Proposition 3.1. *For $\delta = \delta^{BF}$ all memory-one belief-free equilibrium strategies have the same cooperation probabilities after nonempty memory-one histories and are $\sigma = (\sigma_\emptyset, 1, (1-g/l), 1, 0)$ if $l > g$, $\sigma = (\sigma_\emptyset, 1, 0, (l/g), 0)$ if $g > l$ and $\sigma = (\sigma_\emptyset, 1, 0, 1, 0)$ if $g = l$.*

Proof of Proposition 3.1. Again the proof follows the same steps as for actions. Let $V^{a_i}_{a_j s_i}$ denote player i's expected payoff for playing a_i if player j played a_j and observed s_i in the previous round (which means player j is in state $a_j s_i$). If $\sigma_{a_i s_j}$ denotes the (universal) probability of player i to play c after $\{a_i, s_j\}$, we get:

$$\begin{aligned} V^c_{as} &= (1-\delta)(\sigma_{as} - (1-\sigma_{as})l) + \\ &\delta\left((1-\epsilon)\sigma_{as}V_{as} + \epsilon\sigma_{as}V_{cd} + (1-\epsilon)(1-\sigma_{as})V_{dc} + \epsilon(1-\sigma_{as})V_{dd}\right) \end{aligned} \tag{20}$$

$$\begin{aligned} V^d_{as} &= (1-\delta)\sigma_{as}(1+g) + \\ &\delta\left((1-\epsilon)\sigma_{as}V_{as} + \epsilon\sigma_{as}V_{cd} + (1-\epsilon)(1-\sigma_{as})V_{dc} + \epsilon(1-\sigma_{as})V_{dd}\right) \end{aligned} \tag{21}$$

Subtracting (21) from (20) gives:

$$\begin{aligned} 0 = \sigma_{as}\left\{(1-\delta)(l-g) + \delta\left((1-2\epsilon)V_{as} - (1-2\epsilon)V_{cd} - (1-2\epsilon)V_{dc} + (1-2\epsilon)V_{dd}\right)\right\} \\ -(1-\delta)l + \delta\left((1-2\epsilon)V_{dc} - (1-2\epsilon)V_{dd}\right) \end{aligned}$$

The conditions on V_{cd} and V_{cc} based on the belief-free property are now:

$$V_{dc} = V_{dd} + \frac{(1-\delta)l}{\delta(1-2\epsilon)}$$

$$V_{cc} = V_{cd} + \frac{(1-\delta)g}{\delta(1-2\epsilon)}$$

Solving (20) for σ_{cc} and rearranging for V_{cc} yields

$$V_{cc} = \frac{(1-\delta)(\sigma_{cc} - (1-\sigma_{cc})l) + \delta(1-\sigma_{cc})V_{dd} - \delta\sigma_{cc}\frac{(1-\delta)((1-\epsilon)l+\epsilon g)}{\delta(1-2\epsilon)} + \delta(1-\epsilon)\frac{(1-\delta)l}{\delta(1-2\epsilon)}}{1-\delta\sigma_{cc}}$$

Solving (20) for σ_{dd} and inserting the solution for V_{cc} gives

$$V_{dd} = \frac{\sigma_{dd}\left(1 - \frac{(1-\delta)\epsilon l + \epsilon g}{1-2\epsilon}\right) + (1-\delta\sigma_{cc})\frac{\epsilon l}{1-2\epsilon}}{1+\delta\sigma_{dd} - \delta\sigma_{cc}}$$

Next, all V_{as} can be eliminated from (20) solved for σ_{dd} and σ_{dc} proofs Lemma 3.2. For existence we need to assure that $\sigma_{cd} \in (0,1)$ and $\sigma_{dc} \in (0,1)$ for a feasible combination of values σ_{cc}, σ_{dd} and δ. First assume $1 - 2\epsilon - \epsilon(g+l) > 0$ and consider σ_{cd}. In this case $\partial\sigma_{cd}/\partial\sigma_{cc} > 0$ and $\partial\sigma_{cd}/\partial\sigma_{dd} < 0$. Note that $\sigma_{cd} \leq 1$ for any $\delta \in (0,1)$ even if $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$. To establish $\sigma_{cd} \geq 0$ we use $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$. Solving for δ shows gives the condition $\delta > \delta^{BF}$ with $\phi = g$. Next, we consider σ_{dc} still assuming $1 - 2\epsilon - \epsilon(g+l) > 0$. Hence $\partial\sigma_{dc}/\partial\sigma_{cc} < 0$ and $\partial\sigma_{dc}/\partial\sigma_{dd} > 0$. Again $\sigma_{dc} \geq 0$ for any $\delta \in (0,1)$ even if $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$. To establish $\sigma_{dc} \leq 1$ we use $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ which gives $\delta > \delta^{BF}$ with $\phi = l$. Therefore, if $1 - 2\epsilon - \epsilon(g+l) > 0$ the stricter condition on δ results from the larger of the two values g or l as in (17).

If $1 - 2\epsilon - \epsilon(g+l) < 0$, $\partial\sigma_{cd}/\partial\sigma_{cc} < 0$ and $\partial\sigma_{cd}/\partial\sigma_{dd} > 0$. Using $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ we establish that $\sigma_{cd} \leq 1$ only if $\delta \geq 1$ (and the same can be shown for $\sigma_{dc} \geq 0$ when using $\sigma_{cc} = 0$ and $\sigma_{dd} = 1$). Note that (17) also requires $\delta \geq 1$ in this case. For the last case $1 - 2\epsilon - \epsilon(g+l) = 0$, σ_{cd} and σ_{dc} are not defined and (17) also requires $\delta \geq 1$. This proofs Lemma 3.1. To complete the proof, insert (17) together with $\sigma_{cc} = 1$ and $\sigma_{dd} = 0$ into (18) and (19) to obtain the structure of the T1BF response defined by g and l. □

Proposition 3.2. *If players condition on their own action and the ϵ-noisy signal of the other player's action, the existence condition for symmetric memory one belief-free equilibria* *in semi grim strategies is:*

$$\delta^* = \frac{g+l}{(1-2\epsilon)1 + (1-\epsilon)(g+l)} \tag{22}$$

Above this threshold, a continuum $\sigma_{cc} \in (\frac{g+l}{\delta((1-2\epsilon)1+(1-\epsilon)(g+l))}, 1)$ *of semi-grim equilibria exists given by:*

$$\sigma_{dd} = \sigma_{cc} - \frac{g+l}{\delta((1-2\epsilon)1 + (1-\epsilon)(g+l))} \tag{23}$$

and

$$\sigma_{cd} = \sigma_{dc} = \sigma_{cc} - \frac{g}{\delta((1-2\epsilon)1 + (1-\epsilon)(g+l))} \tag{24}$$

Proof of Proposition 3.2. Using the semi-grim property $\sigma_{cd} = \sigma_{dc}$ for (18) and (19) yields (23) and (24). Observe that $\sigma_{dd} < \sigma_{cd} < 1$ for $\sigma_{cc} \in (0, 1)$ and for existence σ_{dd} must be positive which can be rearranged to yield (22). □

Appendix B: Strategy Frequency Estimation

We use the strategy frequency estimation method (SFEM) (Dal Bó and Fréchette, 2011) to obtain maximum-likelihood estimates of prevalence of strategies in our treatments. Consider $i \in I$ subjects with actions a_{ikr} in round $r \in R_k$ of interaction $k \in K$. Our goal is to find $s \in S$ strategies with parameters θ_s and weights ρ_s with the highest likelihood given the data. The probability of action a_{ikr} based on strategy s is:

$$Pr(a_{ikr}|\theta_s) = a_{ikr}s_{ikr}(\theta_s) + (1 - a_{ikr})(1 - s_{ikr}(\theta_s))$$

where $s_{ikr}(\theta_s)$ denotes the probability to play C defined by the parameters of strategy s for subject i in round r of interaction k. For pure strategies $s_{ikr}(\theta_s) \in \{0, 1\}$ and for behavior strategies $s_{ikr}(\theta_s) \in (0, 1)$. Let $\gamma \in (0, 0.5)$ be an error probability constant over i, k, r. We define the likelihood that subject i uses strategy s in interaction k as

$$L(a_{ikr}|\theta_s, \gamma) = (1 - \gamma)Pr(a_{ikr}|\theta_s) + \gamma(1 - Pr(a_{ikr}|\theta_s))$$

The error probability γ avoids that a single mismatch between subject i's actual action and the action predicted by a pure strategy s results in a likelihood of zero that i uses s.

In the estimation of memory-one Markov strategies we assume that for each interaction, subject i independently draws a strategy s from S. The probability to draw strategy s is defined by the weight ρ_s. Based on this assumption, the log likelihood for the data is:

$$LL(\rho, \theta, \gamma) = \sum_{i \in I} \sum_{k \in K} ln \left(\sum_{s \in S} \rho_s \prod_{r \in R_k} L(a_{ikr}|\theta_s, \gamma) \right)$$

For the pure strategy estimation, the strategy parameters θ are defined by the candidate set of pure strategies. To allow for a better comparison we use the standard assumption in the experimental literature that each subject i once draws a strategy s for all interactions k. Based on this assumption, the log likelihood for the data is:

$$LL(\rho, \gamma) = \sum_{i \in I} ln \left(\sum_{s \in S} \rho_s \prod_{k \in K} \prod_{r \in R_k} L(a_{ikr}|\theta_s, \gamma) \right)$$

To obtain parameter which produce global maxima of these functions, we rely the general nonlinear augmented Lagrange multiplier method (Ye, 1987, executed with R package Rsolnp by Alexios Ghalanos and Stefan Theussl). We let the solver generate 10,0000 randomly chosen parameter vectors and sort according to the resulting log-likelihood. The algorithm then performs 100 restarts from the best 100 random parameter vectors as starting points. The reported results come from solutions with the highest likelihood from these 100 restarts of the solver. For the bootstrap samples, we reduce the number of restarts to 10 to reduce the computation time.

Under private monitoring, strategies can only condition on action-signal combination. Under perfect and public monitoring, we do not know ex ante which information profiles subjects take into account. Which strategies best describe subjects behavior under these monitoring structures is an empirical question. Strategies in the perfect treatments can potentially condition on action profiles, the public signals or action signal combinations. In the public monitoring treatments the possibilities reduce to public signals or the action-signal combination. To explore which sets of strategies describe the data from these treatments best, we assume that all strategies of the estimation condition on the same information and run the estimation procedure for all 3 (2) possibilities. We compare the likelihood resulting from estimations and present the result with the highest likelihood.

Table 10: Memory-one Markov Strategies

		Perfect						Public						Private					
		share	$\sigma_\emptyset$	σ_{cc}	σ_{cd}	σ_{dc}	σ_{dd}	share	$\sigma_\emptyset$	σ_{cc}	σ_{cd}	σ_{dc}	σ_{dd}	share	$\sigma_\emptyset$	σ_{cc}	σ_{cd}	σ_{dc}	σ_{dd}
	s_1	0.39	0.65	0.98	0.30	0.82	0.00	0.60	0.00	0.63	0.22	0.02	0.05	0.44	0.37	1.00	0.58	0.00	0.01
	s_2	0.35	0.00	0.51	0.00	0.04	0.03	0.19	1.00	0.96	0.69	0.00	0.10	0.39	0.52	1.00	0.49	0.24	0.00
No	s_3	0.13	0.00	1.00	0.23	0.59	0.24	0.10	0.00	0.97	0.40	1.00	0.30	0.10	0.00	0.00	0.00	0.00	0.21
	s_4	0.07	1.00	1.00	0.35	0.00	0.03	0.09	1.00	0.79	0.00	1.00	0.17	0.04	1.00	0.51	0.00	0.44	0.30
	s_5	0.05	1.00	0.43	0.22	0.40	0.11	0.02	0.00	0.00	0.00	1.00	0.29	0.03	1.00	0.00	0.45	0.00	0.06
	s_1	0.81	1.00	1.00	0.30	1.00	0.00	0.55	0.97	1.00	0.78	0.10	0.10	0.34	1.00	0.97	0.82	1.00	0.00
	s_2	0.15	1.00	0.96	1.00	1.00	0.36	0.22	1.00	0.84	0.83	0.75	0.00	0.30	0.91	1.00	0.88	0.00	0.00
Pre	s_3	0.04	0.38	0.62	0.02	0.00	0.00	0.12	0.61	0.32	0.10	0.50	0.37	0.26	1.00	1.00	0.26	0.52	0.07
	s_4	-	-	-	-	-	-	0.07	0.70	0.42	0.60	0.00	0.00	0.08	1.00	0.66	0.42	0.64	0.00
	s_5	-	-	-	-	-	-	0.04	1.00	0.40	0.74	1.00	0.30	0.03	0.47	0.36	0.00	0.80	1.00
	s_1	0.91	1.00	1.00	0.75	0.33	0.00	0.61	0.98	0.96	0.83	1.00	1.00	0.48	1.00	1.00	0.98	0.00	0.00
	s_2	0.09	1.00	0.91	0.34	0.75	1.00	0.30	1.00	1.00	0.89	0.00	0.00	0.18	1.00	0.99	0.74	0.49	1.00
Rep	s_3	-	-	-	-	-	-	0.08	1.00	0.62	0.41	0.36	0.10	0.17	1.00	0.84	0.56	0.83	1.00
	s_4	-	-	-	-	-	-	0.01	0.00	0.00	0.00	0.17	0.26	0.10	1.00	0.88	0.46	0.38	0.09
	s_5	-	-	-	-	-	-	0.01	0.00	0.00	0.24	1.00	1.00	0.06	0.69	1.00	0.31	1.00	0.12

Notes: Estimates for non-constant strategy use over the last 3 supergames. Shares of zero omitted (-) to increase readability. The error probability γ is zero in all 9 estimations. Values might not add up as expected due to rounding.

Table 11: Strategy set of Fudenberg et al. (2012)

Strategy	Name	Description
ALLD	Always Defect	Always play D
DC	Alternator	Start with D, then alternate between C and D
FC	False cooperator	Play C in the first round, then D forever
DTFT	Exploitative TFT	Play D in the first round, then play TFT
DTF2T	Exploitative TF2T	Play D in the first round, then play TF2T
DTF3T	Exploitative TF3T	Play D in the first round, then play TF3T
DLGrim2	Exploitative Grim2	Play D in the first round, then play Grim2
DLGrim3	Exploitative Grim3	Play D in the first round, then play Grim3
Grim	Grim	Play C until either player plays D, then play D forever
TFT	Tit-for-Tat	Play C unless partner played D last round
PTFT	Perfect TFT/WSLS	Play C if both players chose the same move last round, otherwise play D
T2	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	Tit-for-2-Tats	Play C unless partner played D in both of the last 2 rounds
TF3T	Tit-for-3-Tats	Play C unless partner played D in all of the last 3 rounds
2TFT	2-Tits-for-1-Tat	Play C unless partner played D in either of the last 2 rounds (2 rounds of punishment if partner plays D)
2TF2T	2-Tits-for-2-Tats	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)
LGrim2	Lenient Grim 2	Play C until 2 consecutive rounds occur in which either player played D, then play D forever
LGrim3	Lenient Grim 3	Play C until 3 consecutive rounds occur in which either player played D, then play D forever
PTFT2	Perfect 2TFT	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
ALLC	Always Cooperate	Always play C

Table 12: Estimation of Pure Strategies

	Perfect			Public			Private		
	No	Pre	Rep	No	Pre	Rep	No	Pre	Rep
ALLD	0.44	-	-	0.65	0.02	0.02	0.52	0.02	-
ALLC	-	-	-	-	0.13	0.48	-	-	0.10
GRIM	0.15	0.25	-	-	0.06	-	0.11	0.04	0.01
TFT	0.12	0.40	-	0.06	-	0.04	-	0.03	-
PTFT	-	-	0.18	-	-	-	-	-	-
T2	-	-	-	-	-	-	-	-	-
TF2T	0.04	0.31	-	-	0.02	-	-	0.11	-
TF3T	-	-	-	-	-	-	-	-	-
T2F1T	-	-	-	-	0.04	-	-	0.14	-
T2F2T	-	0.03	-	-	0.18	0.07	-	0.07	0.23
LGRIM2	-	-	0.80	0.06	0.20	0.09	0.22	0.24	0.16
LGRIM3	-	-	-	0.05	0.16	0.24	-	0.14	0.28
PTFT2	-	-	-	-	-	-	-	-	-
FC	-	-	-	-	0.08	0.02	-	-	-
DTFT	0.17	-	-	-	-	-	-	-	-
DTF2T	0.02	-	-	0.09	-	-	-	-	-
DTF3T	-	-	-	-	-	-	-	-	-
DLGRIM2	-	-	-	-	-	-	-	-	-
DLGRIM3	-	-	-	-	-	-	0.01	-	-
DC	-	-	-	-	-	-	-	-	-
BF	0.05	-	0.02	0.09	0.12	0.04	0.14	0.22	0.21
γ	0.07	0.01	0.01	0.10	0.10	0.05	0.06	0.05	0.05
LL	329	80	59	348	385	218	284	40	236

Notes: Estimates for constant strategy use over the last 3 supergames. Strategies condition on action profiles in perfect treatments, and on action-signal profiles in public and private treatments. L/F contains the shares of TFT, PTFT, T2, TF2T, TF3T, T2F1T, T2F2T, LGRIM2, LGRIM3, PTFT2, DTFT, DTF2T, DTF3T, DLGRIM2, DLGRIM3, DC. T1BF is a strategy which starts with cooperation and shows the threshold memory-one belief-free response under imperfect monitoring with $\sigma_{as} = (1, 1, 0.5, 1, 0)$. Shares ≤ 0.02 not displayed (-) to increase readability. γ indicates the probability of trembles.

Table 13: Classification of strategies (Camera et al., 2012)

	Perfect			Public			Private		
	No	Pre	Rep	No	Pre	Rep	No	Pre	Rep
ALLD	71	4	0	84	7	1	71	4	1
ALLC	19	152	160	17	101	128	32	116	127
GRIM	39	149	154	14	49	58	34	80	73
TFT	41	151	156	17	66	87	30	89	83
PTFT	22	147	156	10	44	57	26	74	74
T2	22	147	154	6	45	58	27	77	73
TF2T	27	153	160	16	99	125	39	118	122
TF3T	21	152	160	14	98	126	34	116	125
T2F1T	37	150	156	15	70	84	30	87	75
T2F2T	27	153	160	16	102	126	38	118	124
LGRIM2	27	152	161	15	93	119	42	116	121
LGRIM3	22	152	160	14	101	126	33	115	125
PTFT2	23	147	156	7	43	57	27	76	72
FC	24	1	-	12	16	3	14	3	2
DTFT	51	4	-	34	2	2	34	-	3
DTF2T	13	-	-	13	1	2	6	-	1
DTF3T	14	-	-	8	1	2	4	-	1
DLGRIM2	13	-	-	11	1	2	6	-	1
DLGRIM3	17	-	-	10	1	2	5	-	1
DC	5	-	-	2	4	1	3	2	1

Notes: Classifications of the last 3 supergames. Table shows the number of supergames a strategy predicts well. If a supergame has x rounds and y errors are observed, a strategy does predict the supergame well if the probability of observing x or more errors is smaller than $p = 0.05$.

Appendix C: Communication Content Categories

Table 14: Sub-categories

#	Sub-Category	Category	#	Sub-Category	Category
1	Proposal: both C	coordination	37	Good past experience with DD	deliberation
2	Proposal: both D	coordination	38	Bad past experience with CC	deliberation
3	Proposal: alternate	coordination	39	Bad past experience with CC	deliberation
4	Proposal: self D other C	coordination	40	Good past experience with asym. play	deliberation
5	Proposal: self C other D	coordination	41	Bad past experience with asym. play	deliberation
6	Proposal: other coordination	coordination	42	Positive feedback after CC	relationship
7	Question: what action other	coordination	43	Positive feedback after DD	relationship
8	Announcement: C	coordination	44	Positive feedback after asym. play	relationship
9	Announcement: D	coordination	45	Empathy	relationship
10	Rejection of proposal	coordination	46	Confess D	info
11	Acceptance proposal	coordination	47	Apology	relationship
12	Implicit punishment threat for D	coordination	48	Justification of play	relationship
13	Punishment threat grim	coordination	49	Accusation of cheating	relationship
14	Punishment threat lenient grim	coordination	50	Verbal punishment	relationship
15	Approval of punishment threat	coordination	51	Renegotiation	coordination
16	Ask for coordination	coordination	52	Argument against punishment	coordination
17	Benefits of C	deliberation	53	Small talk	trivia
18	Benefits of D	deliberation	54	Off topic	trivia
19	Benefits of asymmetric play	deliberation	55	Boredom	trivia
20	Related to fairness discussion	deliberation	56	Disappointed after b signal	info
21	Related to strategic uncertainty	deliberation	57	Confusion	deliberation
22	Related to payoffs	deliberation	58	Motivational talk	relationship
23	Related to Prisoner's dilemma	deliberation	59	Report: own signal c	info
24	Related to game theory/economics	deliberation	60	Report: own signal d	info
25	Future benefit of C	deliberation	61	Report: own action C	info
26	Short term incentives of D	deliberation	62	Report: own action D	info
27	Attribute other d to randomness	info	63	Ask for others payoff	info
28	Attribute own d to randomness	info	64	Ask for others signal	info
29	Assurance to have played C	info	65	Ask for others action	info
30	Promise	relationship	66	Report: own payoff 0	info
31	Distrust	relationship	67	Report: own payoff 17	info
32	Trust	relationship	68	Report: own payoff 30	info
33	Argue for trustworthy behavior	relationship	69	Report: own payoff 37	info
34	Report payoff from past games	deliberation	70	Being cheated on in past games	deliberation
35	Report signals of past games	deliberation	71	Counter-proposal	coordination
36	Good past experience with CC	deliberation	72	Rejection of punishment	coordination

Appendix D: Miscellaneous Further Results

Figure 4: Evolution of Cooperation over Supergames

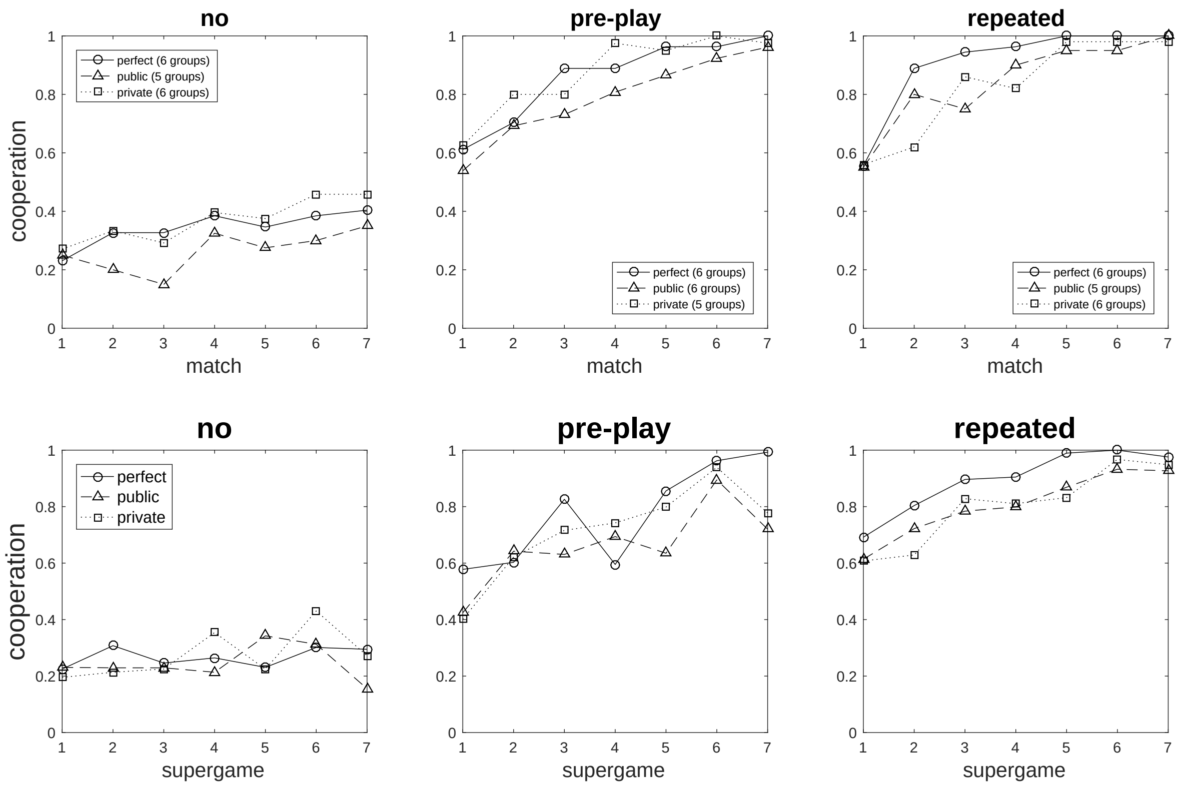

Notes: The upper three panels display average round 1 cooperation rates over the seven supergames. The lower three panels display overall average cooperation rates in the seven supergames.

Figure 5: Efficiency

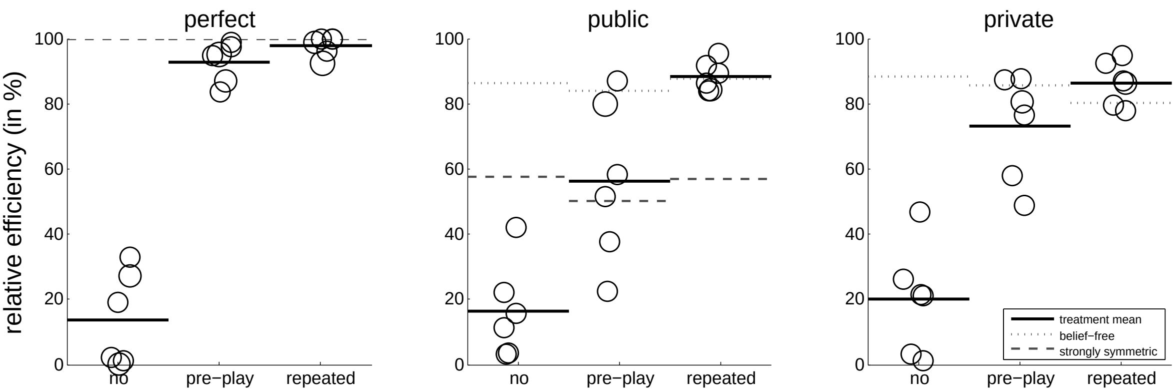

Notes: Efficiency is measured in realized payoffs relative to difference between hypothetical payoffs from 0% and 100% mutual cooperation. Simulated efficiency of the most efficient memory-one belief-free equilibrium and the most efficient strongly symmetric (for public and perfect) are displayed for comparison.