

Bruttel, Lisa V.; Güth, Werner; Kamecke, Ulrich

Working Paper

Time to defect: repeated prisoners' dilemma experiments with uncertain horizon

Jena Economic Research Papers, No. 2007,098

Provided in Cooperation with:

Max Planck Institute of Economics

Suggested Citation: Bruttel, Lisa V.; Güth, Werner; Kamecke, Ulrich (2007) : Time to defect: repeated prisoners' dilemma experiments with uncertain horizon, Jena Economic Research Papers, No. 2007,098, Friedrich Schiller University Jena and Max Planck Institute of Economics, Jena

This Version is available at:

<https://hdl.handle.net/10419/25672>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



JENA ECONOMIC RESEARCH PAPERS



2007 – 098

Time to Defect: Repeated Prisoners' Dilemma Experiments with Uncertain Horizon

by

**Lisa Bruttel
Werner Güth
Ulrich Kamecke**

www.jenecon.de

ISSN 1864-7057

The JENA ECONOMIC RESEARCH PAPERS is a joint publication of the Friedrich Schiller University and the Max Planck Institute of Economics, Jena, Germany. For editorial correspondence please contact m.pasche@wiwi.uni-jena.de.

Impressum:

Friedrich Schiller University Jena
Carl-Zeiss-Str. 3
D-07743 Jena
www.uni-jena.de

Max Planck Institute of Economics
Kahlaische Str. 10
D-07745 Jena
www.econ.mpg.de

© by the author.

Time to Defect: Repeated Prisoners' Dilemma Experiments with Uncertain Horizon

Lisa Bruttel* Werner Güth[†] Ulrich Kamecke[‡]

October 2007

Abstract

Using a symmetric 2-person prisoners' dilemma as the base game, each player receives a signal for the number of rounds to be played with the same partner. The actual number of rounds (the length of the supergame) is determined by the maximal signal where each player expects the other's signal to be smaller, respectively larger, by a fixed number of rounds with 50% probability. In the tradition of Folk Theorems we show that both, mutual defection and mutual cooperation until the individually perceived last round, are subgame perfect equilibrium outcomes. We find experimental evidence that many players do in fact cooperate beyond their individual signal period.

JEL classification: C91, D82, D84.

Keywords: Prisoners' dilemma, Continuation probability, Uncertainty, Experiment

*Humboldt-Universität zu Berlin, Department of Business and Economics, Spandauer Str. 1, 10099 Berlin (Germany), Phone: +49-(0)30-20935786, Fax: +49-(0)30-20935787, e-mail: lisacwoj@wiwi.hu-berlin.de (corresponding author)

[†]Max Planck Institut für Ökonomik, Strategic Interaction Group, Kahlaische Str. 10, 07745 Jena (Germany), email: gueth@econ.mpg.de

[‡]Humboldt-Universität zu Berlin, Department of Business and Economics, Spandauer Str. 1, 10099 Berlin (Germany), e-mail: kamecke@wiwi.hu-berlin.de

1 Introduction

Many strategic situations exhibit a conflict between efficiency and individual rationality. Examples for this include the well-known prisoners' dilemma as well as all standard market games in which individual rationality prevents players from reaching an efficient outcome. In such situations it is often possible to establish cooperation as a non-cooperative equilibrium outcome if the game is repeated infinitely often since the threat to play another equilibrium strategy¹ is sufficient to stabilize cooperation, at least if the players are patient enough. For this logic to hold it is necessary that the game have no commonly known, predetermined end. When rational players know the last decision period, they must play competitively in this period, and a simple backward induction argument implies that rational players will compete from the very first period. To reach cooperative outcomes with rational players, it is therefore necessary to allow players to believe that the game continues (at least with a positive probability).

Experiments, however, have a limited duration. Thus, whenever we want to test experimentally whether cooperative behavior emerges in repeated games, we have to rely on some approximation of an open-ended strategic situation. The most common way to do so is a long finite horizon in which human players seem to ignore the end as long as it is sufficiently far away. This, however, does not avoid endgame effects since cooperation typically collapses as the end of the game comes close. Selten and Stoecker (1986) have already shown that people learn to anticipate the endgame effect, but the point of departure from mutual cooperation is moving cyclically rather than monotonically earlier when the supergame is repeated several times.² Game theoretically, such a process should unravel all the way until there is no cooperation from the very beginning,³ but even limited unraveling questions the approximation of cooperative behavior throughout.

¹We will embed such equilibrium stabilization of voluntary cooperation more generally in the literature on punishment behavior in our conclusion.

²For a learning direction explanation of such cyclical behavior see Selten (1991).

³Similar effects were confirmed in Brosig et al. (2007), and in Bruttel and Kamecke (2007).

There have been several approaches to overcome the problem how to approximate infinity in the lab. The most popular method is the random stopping rule which poses two basic problems. First, it is not easy to implement because the experimenter has no control over the duration of the experiment. Second, credibility of the infinite repetition is difficult to sustain, because participants do not believe that there is a realistic chance that they will spend more than a couple of hours in the lab (see Selten et al., 1997). As a consequence, participants may lose confidence in the experimental setup, becoming less willing to cooperate as the game proceeds. An alternative implementation of Markov strategies (see Bruttel and Kamecke, 2007) turned out to solve this problem, but the procedure was so complicated that individual behavior was considerably less systematic than in comparable repeated games. The approach closest to the one presented here is a public good experiment of Gonzáles et al. (2005) who only provide interval information about the actual duration and find that asymmetric information about the number of rounds reduces the frequency of endgame effects.

Here we propose another procedure to establish an open-ended interaction by considering a finitely repeated prisoners' dilemma game with uncertainty about the duration of the repeated game. Both players receive a signal each about the number of rounds played in the supergame, but only one of these two signals indicates the correct number of rounds R , while the other signal is $R-5$. Both players receive the correct signal with 50% probability so that both are aware of a minimum duration at the beginning of the supergame. At some point in time, when the supergame lasted longer than the smaller signal, one participant learns that the other must have received the correct signal. Being aware of the distance between signals, this participant now also knows the actual duration of the supergame. The opponent with the larger signal, however, does not know the actual duration until the supergame stops.

With this setup we combine elements of the long, but finite horizon – as one participant can know the end of the game – with elements of a random stopping rule as both players must take an infinite regress on larger round numbers into account. We expect human players to rely on the 50% chance that the game stops when their own signal is reached. They will thus compare the expected payoff if they defect in this round with that if they continue to cooperate, yielding further cooperative outcomes also with an expected 50% chance. Although this is not in line with Bayesian rationality, we will show in section 2 that cooperation (except for the last round of the game) is indeed an equilibrium outcome if the prior uncertainty about of the final round is sufficiently diffuse.

Our experimental data indeed reveal the expected effect. We often observe joint cooperation until the second last round of the game. Initially, some participants defect already when their signal is reached, but this endgame effect does not unravel to earlier periods when the supergame is repeated and disappears completely when we repeat the experiment with experienced participants. Unfortunately, we cannot entirely avoid the disadvantages of the random treatment as there exists a robust trend to a lower willingness to cooperate in longer supergames. It seems that participants become increasingly suspicious whether a very large number of rounds will actually be performed. Needless to say that the very long plays were actually implemented as described in the instructions.

2 The experimental setup with constant continuation probability and almost diffuse priors

Consider a symmetric 2-person prisoners' dilemma with stage game strategies "cooperate" and "defect" and the payoff matrix from Table 1,

	Cooperate	Defect
Cooperate	c, c	l, d
Defect	d, l	n, n

Table 1: Payoff matrix

with cooperation as the upper, respectively left, strategy and defection as the bottom row, respectively right, column, and with the usual conditions $l < n < c < d$ (dominance of defect) and $d + l < 2c$ (efficiency of (cooperate, cooperate)). We rely on this game as the stage game for our supergame analysis.

In a Bayesian world one captures a random end of the game by a probability distribution for this end. If the number of rounds $R \in [\underline{R}; \infty)$ is determined by a random variable with a constant continuation probability ψ ,⁴ the (conditional) probability $\rho(R)$ for a last round $R \geq \underline{R}$ (after round $\underline{R} - 1$) is given by $\varrho(R = \underline{R}) = 1 - \psi$, $\varrho(R = \underline{R} + 1) = \psi(1 - \psi)$, and similarly $\varrho(R) = \psi^{R - \underline{R}}(1 - \psi)$ for $R > \underline{R} + 1$ (with $\rho(R) = 0$ for $R < \underline{R}$).

At the beginning of the game, let one player i receive signal $s_i = R$ while the other player $j \neq i$ receives signal $s_j = R - Q$ with $Q \leq \underline{R}$ such that both possible signal assignments, $(s_1, s_2) = (R - Q, R)$ and $(s_1, s_2) = (R, R - Q)$, have the same probability. This captures the intuition that the largest possible number of rounds is determined by the experimenter and that participants have no (objective) information how this is done. We will show more formally that the distribution of termination rounds converges to a uniform distribution as $\psi \rightarrow 1$. This limit (which is not a distribution) can be interpreted as an approximation of a “diffuse prior” in the usual sense (better approximations seem possible but are more complicated).

⁴This, of course, raises the question whether experimenters can actually (believe to be able to) determine the final round accordingly. We do not claim this but admit how we have actually determined R in ways guaranteeing that no excessively large R has to be implemented (see section 3 below).

If player i receives signal $s_i \geq \underline{R}$, either s_i or $s_i + Q$ is the true number R of rounds. The posterior distribution of these two events is readily computed with Bayes' rule as follows. The ex ante probability for $R = s_i$ is $\rho(s_i) = \psi^{s_i - \underline{R}}(1 - \psi)$ while the ex ante probability for $R = s_i + Q$ is $\rho(s_i + Q) = \psi^{s_i + Q - \underline{R}}(1 - \psi)$. Hence the posteriors are

$$\begin{aligned} P(s_i = R) &= \frac{\psi^{s_i - \underline{R}}(1 - \psi)}{\psi^{s_i - \underline{R}}(1 - \psi) + \psi^{s_i + Q - \underline{R}}(1 - \psi)} = \frac{1}{1 + \psi^Q} \\ P(s_i = R + Q) &= \frac{\psi^{s_i + Q - \underline{R}}(1 - \psi)}{\psi^{s_i - \underline{R}}(1 - \psi) + \psi^{s_i + Q - \underline{R}}(1 - \psi)} = \frac{\psi^Q}{1 + \psi^Q}. \end{aligned}$$

Obviously, $P(s_i = R) > 0.5$ and $\lim_{\psi \rightarrow 1} P(s_i = R) = 0.5$. In other words, signal s_i is, with more than 50% probability, the final round of the game, but the difference vanishes for diffuse priors.

Next, we derive the condition for both players i cooperating in round r as long as both have cooperated in the past and as long as $r < s_i + Q$, and defecting otherwise. If i is the first and only one to defect, he gets d once and n for the rest of the game instead of d or l in the last round and c before. When i has received signal s_i , player i knows that the game cannot last longer than $s_i + Q$ rounds. In a round r with $s_i < r < s_i + Q$ he therefore knows that the game will continue for $Q - r$ periods. Hence a defection in r gives $d + (Q - r - 1)n$ instead of $(Q - r - 1)c + d$. Thus, it does not pay to deviate from the cooperative strategy.

A defection in round $r = s_i$ has two possible consequences: if s_i is the correct number of rounds (with probability $1/(1 + \psi^Q)$), player i gets n instead of l because the opponent defects. However, if the game lasts for $s_i + Q$ rounds (with probability $\psi^Q/(1 + \psi^Q)$), he receives d once and Q times n instead of Q times c and d once. Stability therefore requires that

$$\begin{aligned}
 \frac{1}{1 + \psi^Q} (n - l) + \frac{\psi^Q}{1 + \psi^Q} ((d + Qn) - (Qc + d)) &\leq 0 \\
 (n - l) + \psi^Q (Qn - Qc) &\leq 0 \\
 \psi^Q &\geq \frac{n - l}{Q(c - n)}.
 \end{aligned}$$

Defecting one round earlier, in $r = s_i - 1$, yields d instead of c once with certainty, n instead of l once if s_i is the correct number of rounds, and $Q + 1$ times n instead of Q times c and one defection payoff d if the game lasts for $s_i + Q$ rounds. Stability therefore requires that

$$\begin{aligned}
 d - c + \frac{1}{1 + \psi^Q} (n - l) + \frac{\psi^Q}{1 + \psi^Q} ((Q + 1)n - (Qc + d)) &\leq 0 \\
 (1 + \psi^Q) (d - c) + (n - l) + \psi^Q (Q(n - c) + (n - d)) &\leq 0 \\
 \psi^Q (d - c + Q(n - c) + (n - d)) + n - l + d - c &\leq 0 \\
 \psi^Q &\geq \frac{n - l + d - c}{(Q + 1)(c - n)}, \quad (1)
 \end{aligned}$$

a condition which is more restrictive than the former. In earlier rounds a defection yields n instead of c once more without further gains. This establishes that requirement (1) is the most restrictive stability condition. It is easy to see from (1) that the difference Q in the signals weakens the requirement for equilibrium cooperation considerably.

Of course, this new round structure does not alter anything about the conclusion that universal defection by both players, i.e., the play of (defect, defect) in all subgames, constitutes an equilibrium benchmark solution with low payoffs $R \cdot n$. More interesting, however, is the symmetric cooperative solution in which both participants select their strategies $v = (v_1, v_2)$ according to the rule

- *defect* in round $s_i + Q$.
- *cooperate* in all earlier rounds if so far only $(cooperate, cooperate)$ has been used; *defect* otherwise.

Proposition If condition (1) is satisfied, $v = (v_1, v_2)$ is a subgame perfect equilibrium of the supergame, yielding the payoff of $(R - 1)c + l$ for the player with the larger signal and $(R - 1)c + d$ for the player with the smaller signal.

3 Experimental design

In the experiment, we use the payoff matrix from Table 2.

	Cooperate	Defect
Cooperate	80, 80	0, 140
Defect	140, 0	10, 10

Table 2: Parameterization of the payoff matrix

The payoff matrix is designed to facilitate initial cooperative play since we want to study timing of defection. Initial cooperation hopefully helps us to detect whether in fact a considerable share of subjects follows our cooperative prediction and defects only in the last round of a supergame. The payoff structure tries to induce this in two ways: First, we have a considerable cooperation incentive $c - n$ due to the large difference of 70 between the payoffs from joint cooperation and joint defection; second, the relative gain $d - l$ from being first in defecting from mutual cooperation is, with 60, quite large. Note that (1) can also be expressed as

$$1 + \psi^Q (Q + 1) \geq \frac{d - l}{c - n} \left(= \frac{6}{7} \right), \quad (2)$$

revealing how the large cooperation incentive helps to satisfy (1), whereas increasing the relative gain from defection would question (1).

In each of our supergames, the difference between the two signals is equal to $Q = 5$. For the above payoff matrix the binding stability condition is therefore satisfied if

$$\psi \geq \sqrt[Q]{\frac{n - l + d - c}{(Q + 1)(c - n)}} = \sqrt[5]{\frac{1}{6}} = 0.6988.$$

Both players $i = 1, 2$ receive an integer signal s_i before playing the first round of a supergame where it is commonly known that $|s_1 - s_2| = Q$ and that the actual number of rounds for successively playing the base game with the same partner is

$$R = \max\{s_1, s_2\}. \quad (3)$$

Payoffs in the supergame are simply the sum of periodic payoffs u_i in the $t = 1, \dots, R$ rounds for $i = 1, 2$. We assume that all previous moves in the history of the supergame are commonly known.

Of course, our setup does not avoid the usual inconsistency of a random stopping rule as our argument only works if there is no certain end to the game. In reality we have to impose such an end, and for the experiment we generate signals s_1 and s_2 according to the following process:

- We determine one signal $R \in [10; 30]$.
- We generate $s = R - Q$ (where $Q = 5$) as the other signal.
- We set either $(s_1, s_2) = (s, R)$ or $(s_1, s_2) = (R, s)$ with 50% probability each.

We tried to select the true number of rounds R approximating a random continuation rule with a sufficiently high continuation probability. For this purpose we started with a set of R 's drawn with an underlying continuation probability of 95%, yielding an expected duration of 20 periods. However, to keep the game within reasonable time limits, we excluded all values of R smaller than 10 and larger than 30. This, of course, was not told to the participants. Instead, the instructions said that the signals R were determined arbitrarily without any further specification about the specific mechanism. We expect that this procedure is a sufficiently good approximation to a situation in which both players $i = 1, 2$ expect ex ante $s_i > s_j$ as well as $s_i < s_j$ for $j \neq i$ with nearly equal probability. Of course, this ignorance is dissolved if round $s_i + 1$ is reached, because then player i knows that $R = s_i + Q$ is the last round and that his partner is not aware of this since $s_j > s_i$ must hold.

To provide a reference point with which we can compare our data, we conduct a baseline treatment where the actual duration of each supergame is known to all participants. Furthermore, we repeat the treatment with uncertain horizon with experienced participants to check the robustness of our results. From three different experimental studies (Selten and Stoecker, 1986; Brosig et al., 2007; Bruttel and Kamecke, 2007) with a recall of participants for a second session we know that a drastic change of behavior often occurs after sufficient time of reflection, e.g., when participants had time to reconsider their decisions at home. In particular, the experiment of Selten and Stoecker was conducted in two sessions on two different days, and the forward shift in defection time was only observed when we compared the first and second session. Simply repeating a

supergame within one session does not trigger such learning since participants have no time to thoroughly consider alternative strategies. Thus, inviting experienced players for the second time, we have a stronger test how robust our results are. Finally, we conduct one session in the video lab with two participants in each cabin assuming the role of one player. This experiment contains fewer supergames than the other sessions and is not included in the general data analysis. Its main purpose is to listen to the interaction of the matched participants to learn about their motivation of choices. Table 3 provides an overview of the treatments.

Treatment	Description	Experienced Subjects?	# Subjects	# Ind. Obs.
KNOWN	Known duration	No	30	4
UNCERTAIN	Duration signal	No	72	9
UNCERTAIN2	Duration signal	Yes	32	4
UNCERTAIN_VIDEO	Duration signal	No	12	1

Table 3: Treatments

The actual duration of the supergames selected by our random procedure are listed in Table 4. To facilitate the comparison between treatments with known and uncertain horizons, we maintained the duration of supergames in those sessions. Only for the repetition of the treatment with uncertain horizon we did generate new random sequences.

Participants were undergraduates recruited from different disciplines. After the experiment participants answered an ex-post questionnaire. Sessions took place in the computer laboratory at the Humboldt-University Berlin (pilot session) and the Max Planck Institute Jena (main sessions). The experiment was conducted using z-tree software (Fischbacher, 2007). All in all, 146 subjects participated in the experiment, 32 per main session (8 in the pilot and 12 in the video lab). In treatment UNCERTAIN2, we reinvited 32 participants, who had already gained experience in an earlier main session of

Horizon	UNCERTAIN			KNOWN
	Inexperienced	Experienced	Video	Inexperienced
	12	11	23	12
	16	19	26	16
	14	28	21	14
	30	14	12	30
	11	15	13	11
	23	10	12	23
	12	22	17	12
	25	24	24	25
	10	14	25	10
	14	18	16	14
	13	17	-	13
	26	23	-	26
	14	11	-	14
	24	15	-	24
	26	17	-	26
	14	14	-	14
	29	18	-	29
	24	28	-	24
	10	19	-	10
	18	10	-	18

Table 4: Duration of the supergames

this treatment. On average, participants earned 15.88 euros for their participation in one session. All sessions lasted on average two hours, including the time for reading the neutrally framed instructions (see Appendix for an English translation) and answering the post-experimental questionnaire. In all treatments participants were randomly rematched with another opponent after each supergame. They participated in 20 supergames (10 supergames in the video lab) of different lengths. The 32 participants of one main session were partitioned into four matching groups of eight participants each so that we collected four independent group observations per session.⁵ Participants were not aware of this restricted rematching. The instructions allude to different opponents in different supergames.

⁵In one matching group of treatment KNOWN we had only six participants instead of eight.

4 Results

Figure 1 depicts the cooperation rates in the last 10 periods of all three treatments.⁶ After the initial five supergames high cooperation rates of more than 90% persist in all sessions. In treatment UNCERTAIN2, we even find 100% cooperation most of the time. Thus, our payoff parameterization successfully encourages initial cooperation of experienced participants. In our data analysis we therefore concentrate on the long run reliability of these cooperation rates. What we mean by this will be demonstrated for the benchmark treatment with a known end in the following subsection.

4.1 Benchmark treatment

The first diagram of Figure 1 depicts the cooperation rates in the last 10 periods of treatment KNOWN. In this baseline treatment, we find the usual endgame effect with cooperation rates decreasing to about 15% in the final period of a supergame. Even more important is the observation that this endgame effect begins earlier in later supergames. In particular, we find that cooperation rates in the third and second last rounds of the final five supergames of a session are significantly lower than in the same rounds of the first five supergames (Wilcoxon signed rank test, one-sided, $p\text{-value} = 0.0625$). In the final round, this difference is not significant. On average, defection occurs 1.5 periods earlier in the final five than in the first five supergames. Note that this relatively small forward shift of the endgame effect occurs within one session.

From Selten and Stoecker (1986) as well as from Bruttel and Kamecke (2007) we know that unraveling of cooperation becomes stronger when reinviting participants for a second session. Illustrating this crucial effect, Figure 2 shows the average cooperation

⁶We decided to consider 10 periods which is the minimum duration of the supergames in order to facilitate the comparison of supergames having different lengths.




Figure 1: Average cooperation rate in each treatment

rates of Bruttel and Kamecke's experiment concerning a related repeated prisoners' dilemma supergame.⁷ Performing two successive sessions with the same participants on different days strongly shifts forward the endgame effect in the second session. Thus, learning under a known finite horizon qualitatively confirms backward induction, indicating that limited cooperation could be caused by the unraveling endgame effect.




Figure 2: Cooperation rates in Bruttel and Kamecke (2007), treatment with known horizon

⁷The main differences of their framework to the experiment presented here are a shorter duration of the supergames with constantly seven periods and another payoff table.

4.2 Reliable cooperation

We expect our treatment with uncertain horizon to overcome the unraveling problem of the finite benchmark treatment. Consequently, we are mainly interested in the time when cooperation collapses during a supergame. In the two treatments with uncertain horizon, the two possible periods where one might expect cooperation rates to drop are the two different signal periods. According to our cooperative equilibrium prediction, only the participants with the smaller signals should defect in the final period R of a supergame (they know that this is the final period), whereas we expect high cooperation rates close to 100% in all earlier periods. However, since the cooperative equilibrium requires some sophisticated and complex thoughts, we expect this pattern to show up especially in treatment UNCERTAIN2 with experienced participants. Inexperienced participants may pursue a more natural strategy by simply deviating in their own signal period. Therefore, we expect cooperation rates to drop five periods before the end when the lower signal period is reached in treatment UNCERTAIN, in particular in the initial supergames of this treatment. These expectations are expressed in the following hypotheses:

1. **Cooperation:** Cooperation collapses mainly in the final period of a supergame. Participants with the smaller signal (less than the actual number of periods) defect in the final period while participants with the larger signal continue to cooperate.
2. **Learning:** Cooperation rates increase as the experiment proceeds. They are significantly higher in UNCERTAIN2 than in UNCERTAIN.
3. **No unraveling:** The unraveling observed in treatment KNOWN does not appear, especially not in UNCERTAIN2.

The second diagram in Figure 1 illustrates the cooperation rates over the last 10 rounds of treatment UNCERTAIN. The curve drops in the fifth last period when reaching the first signal, which does not trigger a monotonic decline of cooperation rates, however.

In contrast, average cooperation increases again although the actual duration should now be clear also to the participants having received the higher signal. In our view, this indicates some willingness to forgive similar to tit-for-tat strategies rather than grim strategies. This pattern of cooperation rates is very stable over time. Although in the first five supergames participants cooperate significantly⁸ less than in later supergames, there is no difference in the behavior in later (than the fifth) supergames.

In contrast to treatment KNOWN, in UNCERTAIN2 (the third diagram of Figure 1) there is a slight decline of cooperation rates only five periods before the end. Nearly 100% cooperation persists until R-2 with a sharp drop in cooperation to about 30% in the final period. We think that subjects reconsidered the game at home after their first participation and concluded, as suggested by our cooperative equilibrium prediction, that it would be best to cooperate four periods longer than their own signal. There is significantly more cooperation in the final 10 periods (except for the last period) in UNCERTAIN2 than in UNCERTAIN (Wilcoxon-Mann-Whitney test, one-sided, p-values: $p(R0) = 0.1301$, $p(R - 1) = 0.0220$, $p(R - 2) = 0.0067$, $p(R - 3) = 0.0034$, $p(R - 4) = 0.0042$). This supports the hypothesis that the endgame effect in treatment UNCERTAIN2 is more concentrated in the final period. However, across all supergames the decline of the average cooperation rates from periods 7 to 6, 6 to 5, 3 to 2, 2 to 1, 1 to 0 (counted from the final round) is also in this treatment significant at the 10% level of significance.⁹

The difference between the cooperation rates in treatments KNOWN and UNCERTAIN2 is also reflected in the average profits. Table 5 presents the average profits in each of our treatments in the final seven periods. Average profits in UNCERTAIN2 are higher (but not significantly) than in KNOWN. Thus, the sharper endgame effect in UNCERTAIN2,

⁸We tested the difference between earlier (first to fifth) and later (sixth to twentieth) supergames in a Wilcoxon signed rank test for each of the final 10 periods. The difference is significant at the 5% or even at the 1% level (one-sided test) in all except the final two periods, where cooperation drops in all 20 supergames.

⁹Note that 6.25% is the maximum possible significance with only four independent observations.

which does – in contrast to the case of a known finite horizon – not move to earlier periods in later supergames, seems also to pay off. Again, the difference is only visible between KNOWN and UNCERTAIN2 (with experienced participants) but not between KNOWN and UNCERTAIN.

Matching Group	Treatment		
	UNCERTAIN	UNCERTAIN2	KNOWN
1	62.32	67.00	67.25
2	61.96	70.07	63.98
3	59.48	71.21	60.39
4	57.23	69.55	57.55
5	60.46	-	-
6	62.27	-	-
7	62.09	-	-
8	68.05	-	-
9	65.45	-	-
Average	62.15	69.46	62.61

Table 5: Average profits

4.3 Consequences of early deviation

Participants having received the smaller signal know the actual number of rounds R as soon as the supergame continues after their signal period. Thus, five periods later, they can defect in the final period without having to fear later punishment by their opponent. The other participants (larger signal) do not have this information and, in contrast, might expect that they could be punished during further five periods when deviating in their signal period. Therefore, we predict more participants who received the larger signal to cooperate in the final period than those having received the smaller signal. In the data we find that in treatment UNCERTAIN2 (UNCERTAIN) 53.13% (36.25%) of the participants having received the larger signal cooperate in the final period compared to only 9.37% (14.03%) of those having received the smaller signal. The difference is much more pronounced in the treatment with experienced participants.

However, quite naturally the recipients of the small signal are tempted to defect when their own signal period is reached. In fact, 21.9% of participants who received a small signal in treatment UNCERTAIN reveal their information by deviating in R-5 (after cooperating in R-6).¹⁰ Yielding to this temptation has severe consequences for cooperation in the successive last five periods of the supergame. The usual backward induction argument applies, and the only rational continuation after this move is defection for the rest of the game. As in KNOWN, we would not expect this to occur immediately, but again we expect an unraveling of cooperation to begin. Indeed, 64.6% of the opponents who received the high signal defect already in R-1 after a deviation of their opponent in R-5 in order not to be exploited in R, while only 21.8% of the opponents who received the high signal defect in R-1 when the recipient of the small signal cooperated in R-5.

In order to better understand the drop in cooperation in the fifth last period in treatments UNCERTAIN and UNCERTAIN2, we distinguish pairs according to the behavior of recipients of the small signal in the fifth last period (see Figure 3). According to the equilibrium prediction, the recipients of the small signal should continue cooperating until the second last round and defect only in the last period. Of the recipients of the small signal in treatment UNCERTAIN (UNCERTAIN2) 63.1% (83.4%) indeed cooperate in R-5. In these cases, cooperation remains on a very high level until the second last period. The remaining 36.9% (16.6%) of the recipients of the small signal defect in R-5. Deviating in the fifth last period causes cooperation to drop – after a moderate recovery to about 50% in periods R-4 and R-3 – below 10% by the end of the supergame. In treatment UNCERTAIN2, behavior is again more systematic than in UNCERTAIN. In particular, more recipients of the small signal follow the predicted cooperative strategy, cooperation mainly collapses in the fifth last period and recovers to a lesser degree in the final periods, indicating that unraveling is better understood by more experienced recipients.

¹⁰Here, we do not consider the (14.9% of the) cases where the recipients of the small signal already defected in R-6.

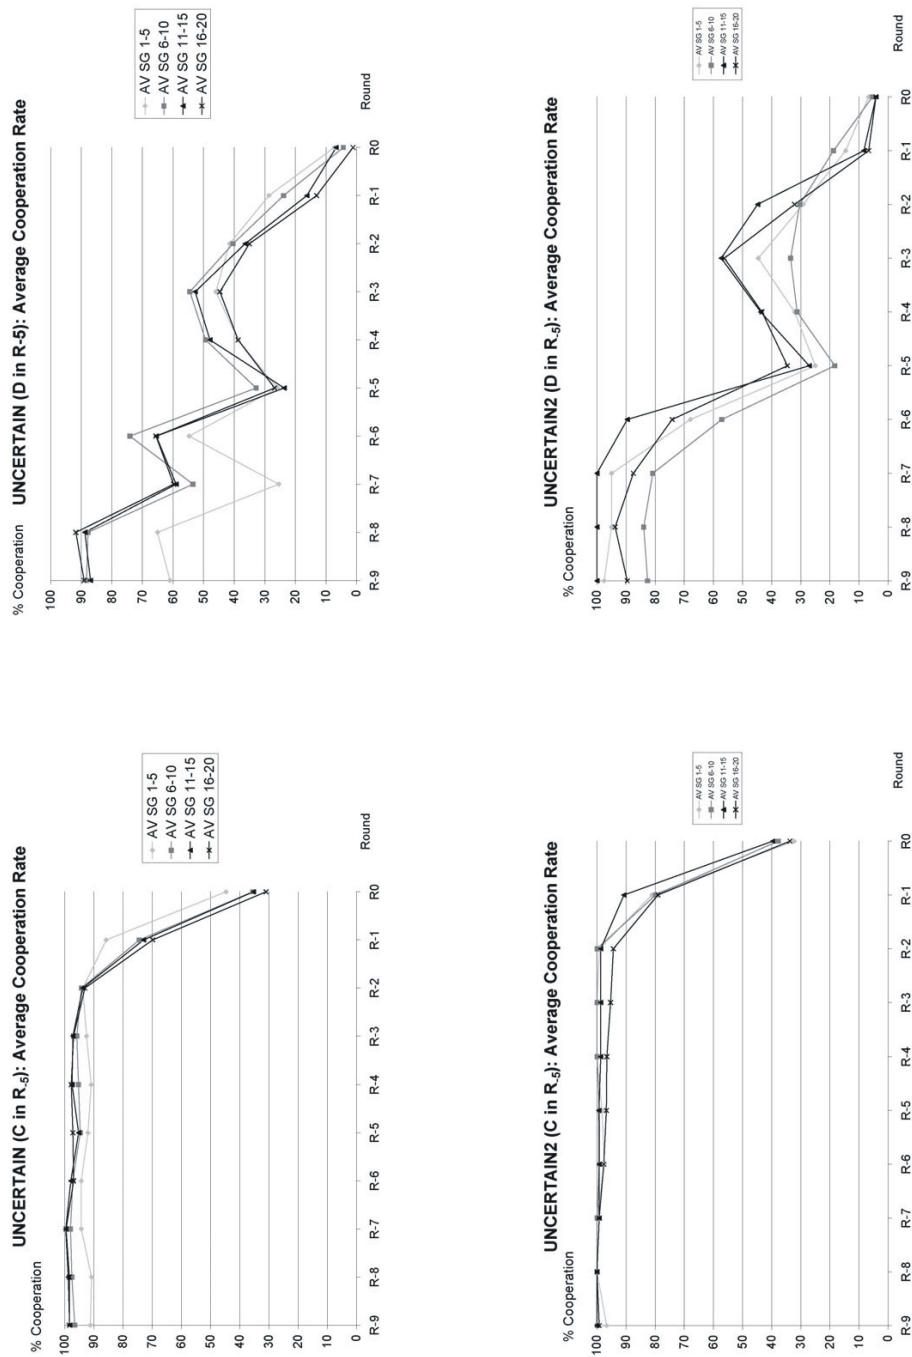


Figure 3: Average cooperation rates in UNCERTAIN and UNCERTAIN2, separated by the decision of the recipient of the small signal in R-5

4.4 Stationary cooperation in long supergames?

In a treatment with random continuation, Bruttel and Kamecke (2007) found that participants do not believe in large repetition numbers. Instead, a continuation is considered less likely, the more periods have been played, triggering declining cooperation rates over time. To analyze whether subjects in our experiment similarly distinguish between supergames with less and more than 20 periods, Figure 4 depicts the cooperation rates over the last 10 rounds of the respective supergames.

In treatment KNOWN there is no significant difference in the cooperation rates of shorter and longer supergames. In the two treatments with uncertain horizon, however, the curves do not show the same pattern in supergames with less and more than 20 periods. In particular, in UNCERTAIN2, participants, having already experienced that the interval of possible durations extends from 10 to 30 periods, cooperate less in the final periods of a supergame lasting more than 20 periods. In supergames with more than 20 periods, cooperation rates begin decreasing in the sixth period before the end and remain about 15% points below that in supergames with less than 20 periods. In treatment UNCERTAIN with inexperienced participants, this difference is less obvious; here the cooperation rates in supergames with less than 20 periods are only slightly lower than in longer supergames, although the drop in cooperation in the fifth last period is more pronounced in supergames with more than 20 periods. This indicates that for $R > 20$, participants are less inclined to believe that the supergame will continue after their signal period. In other words, they seem to reason that their signal equals R and that their opponent, having received a smaller signal, already knows this and will surely defect in their signal period.




Figure 4: Average cooperation rate, depending on the length of a supergame

5 Discussion

To avoid a straightforward finite horizon, we introduce uncertainty about the true duration of a repeated prisoners' dilemma game. Each participant receives a signal about the duration of the supergame, knowing that this signal is only correct with 50% probability. With the remaining 50% probability the opponent received the correct information. This mechanism allows for equilibrium cooperation in the repeated game – except for the last round – even if its duration is limited in advance. We find that this, in contrast to a finite horizon, prevents the endgame effect shifting to earlier periods, not only (game) theoretically but also experimentally and more so in later supergames. However, we also find less final cooperation when the supergames last longer than participants on average expect.

Let us, finally, relate our specific study to the more general context of efficiency enhancing retribution,¹¹ meaning that the anticipation of punishment stabilizes voluntary cooperation. One has to distinguish between equilibrium and non-equilibrium punishment. That non-equilibrium threats can guarantee not only efficiency but also fairness is convincingly and robustly demonstrated by many ultimatum experiments (see Camerer, 2003, for a survey), showing that many responders turn down substantial but seemingly unfair offers¹² which, if anticipated by proposers, prevents such proposals in the first place. Assuming a prisoners' dilemma as a base game for repeated interaction as we do is a way to explore equilibrium punishment since defecting is a strictly dominant strategy in the base game. The idea here is that retribution is an aspect of (subgame perfect) equilibria but not necessarily of (subgame perfect) equilibrium play. Whereas Folk Theorems (see Aumann, 1981, for an older survey, and Neyman, 1999, for a more

¹¹By restricting retribution to the purpose of efficiency enhancement, one rules purely ex-post interpretations of retributive emotions as expressed in the principle “an eye for an eye, ...” Here the hurt is accepted as a fact and the question is only whether its victim, and possibly also society, desires revenge.

¹²Similar evidence is provided by public goods experiments with an additional subsequent punishment phase (e.g., Fehr and Gächter, 2000; Güth et al., 2007, study a more natural punishment institution).

recent theoretical study of “infinite” horizon games) can avoid inefficiency altogether, we explored a situation where this is possible in all but the last round of interaction.¹³ In this sense, we have studied equilibrium retribution to enhance efficiency, a rather familiar topic in game theory but much less explored experimentally.

6 Appendix: Instructions

(Translated from the German)

Thank you for participating in this experiment.

These instructions are identical for all participants. Please read them carefully. Raise your hand if you have any questions regarding the experiment. We will then come directly to your place. Please be quiet during the experiment and do not talk to other participants.

You will participate in the following supergame 20 times. Your payment depends on your decisions and on the decisions of other participants. You will not learn who the other participants are, and they will not learn about your identity.

Each of the 20 supergames will consist of several periods. [KNOWN: *The supergames have different lengths. At the beginning of each supergame you and the other participant will be informed about the duration of the current supergame.*] In each supergame, you play against the same opponent during all periods, but you play against different opponents in different supergames. After having participated in 20 supergames with

¹³This is similar to Folk Theorems for finite horizon repeated games, assuming base games with multiple equilibria (see Benoit and Krishna, 1985)

different opponents, we will add up your profits of all 20 supergames which determine your payoff. Your gains and losses during the experiment are counted in points. The exchange rate is 1,700 points for 1 euro. You will receive your payoff directly after the end of the experiment.

In each period you have to choose one out of two alternative decisions, X and Y. Depending on the decision of your opponent, your payoff is then determined according to the table below. You find your own payoff at the first position of each item and the payoff of your opponent in the second position. At the end of each period, you learn about both your and your opponent's decision as well as your own payoff in points.

	Other's decision	
	X	Y
Your decision	X	80;80 0; 140
	Y	140; 0 10;10

- If you and the other participant both choose X, each of you receives 80 points.
- If you and the other participant both choose Y, each of you receives 10 points.
- If one of you chooses X and the other Y,
 - the one who chooses X gets 0 points
 - and
 - the one who chooses Y gets 140 points.

[UNCERTAIN: *The supergames have different lengths. At the beginning of each supergame you and the other participant will receive a message n about the duration of the current supergame. You only know your own message but not the message of the other*

participant. One of these messages contains the actual number of rounds, whereas the other message contains a number minus five rounds. Suppose, for example, you received the message $N = 73$. In this case, the other participant either received the (wrong) message $n' = 68$, and the actual number of rounds is $R = 73$ - or the other participant received the (correct) message $n' = 78$, and the actual number of rounds is $R = 78$. Since you do not know the message of the other participant, your message $n = 73$ therefore tells you only that the supergame lasts either 73 or 78 rounds.

How are the messages being determined?

Previous to this experiment, the actual number of rounds R was arbitrarily determined for each of the supergames. With 50% probability your message contains the number R , and the other participant receives the message $n' = R - 5$. With the remaining 50% probability you receive the message $n = R - 5$ while the other participant receives the actual number of rounds R . Of course, your message does not give any information whether your number contains R or $R - 5$.

References

- Andreoni, J. and Miller, J.H. (1993). Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence, *The Economics Journal*, vol. 103, pp. 570-585.
- Aumann, R.J. (1981). Survey of Repeated Games, in: Aumann, R.J. (Ed.), *Essays in Game Theory and Mathematical Economics*, pp. 11-42, Bibliographisches Institut BI, Mannheim.
- Benoit, J.-P. and Krishna, V. (1985), Finitely Repeated Games, *Econometrica*, vol. 53, no. 4., pp. 905-922.
- Brosig, J., Riechmann, T. and Weimann, J. (2007). Selfish in the End? An Investigation of Consistency and Stability of Individual Behavior, *Munich Personal RePEc Archive*, vol. 2035.
- Camerer, C. F. (2003). *Behavioral Game Theory. Experiments in Strategic Interaction*, Princeton University Press, Princeton.
- Bruttel, L. and Kamecke, U. (2007). Infinity in the Lab - The Role of the Continuation Rule for Cooperation in Repeated Games, *Working Paper*.
- Dal Bó, P. (2005). Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games, *The American Economic Review*, vol. 95, no. 5, pp. 1591-1604.
- Fehr, E. and S. Gaechter (2000), Cooperation and Punishment in Public Goods Experiments, *American Economic Review*, vol. 90, no. 4., pp. 980-994.
- Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics*, vol. 10, no. 2, pp. 171-178.
- González, L.G., Gueth, W., and Levati, M.V. (2005). When does the game end? Public goods experiments with non-definite and non-commonly known time horizons, *Economics Letters*, vol. 88, no. 2, pp. 221-226.
- Gueth, W., Levati, M.V., Sutter, M., and Heijden, E. van der (2007). Leading by example

with and without exclusion power in voluntary contribution experiments, *Journal of Public Economics*, vol. 91, no. 5-6, pp. 1023-1042.

Kreps, D.M., Milgrom, P., Roberts, J. and Wilson, R. (1982). Rational Cooperation in the Finitely Repeated Prisoner's Dilemma, *Journal of Economic Theory*, vol. 27, pp. 245-252.

Neyman, A. (1999). Cooperation in Repeated Games When the Number of Stages is not Commonly Known, *Econometrica*, vol. 67, no. 1., pp. 45-64.

Normann, H. and Wallace, B. (2005). The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment, Working Paper, Royal Holloway.

Selten, R. (1991). Anticipatory Learning in 2-Person Games, in: Selten, R. (ed.), *Game Equilibrium Models I*, pp. 98-154, Springer, Berlin.

Selten, R., Mitzkewitz, M. and Uhlich, G.R. (1997). Duopoly Strategies Programmed by Experienced Players, *Econometrica*, vol. 65, no. 3, pp. 517-555.

Selten, R. and Stoecker, R. (1986). End Behavior in Sequences of Finite Prisoner's Dilemma Supergames, *Journal of Economic Behavior and Organization*, vol. 7, pp. 47-70.