



Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

**Time Horizon and Cooperation
in Continuous Time**

Maria Bigoni
Marco Casari
Andrzej Skrzypacz
Giancarlo Spagnolo

Quaderni - Working Paper DSE N° 796



TIME HORIZON AND COOPERATION IN CONTINUOUS TIME*

MARIA BIGONI[†] MARCO CASARI[‡] ANDRZEJ SKRZYPACZ[§]
GIANCARLO SPAGNOLO[¶]

November 16, 2011

Abstract

When subjects interact in continuous time, their ability to cooperate may dramatically increase. In an experiment, we study the impact of different time horizons on cooperation in (quasi) continuous time prisoner's dilemmas. We find that cooperation levels are similar or higher when the horizon is deterministic rather than stochastic. Moreover, a deterministic duration generates different aggregate patterns and individual strategies than a stochastic one. For instance, under a deterministic horizon subjects show high initial cooperation and a strong end-of-period reversal to defection. Moreover, they do not learn to apply backward induction but to postpone defection closer to the end.

JEL Classification Numbers: C72, C73, C91, C92, D74

Keywords: Folk Theorem, Prisoner's dilemma, backward induction, termination rule, infinite horizon

*We thank Pedro Dal Bó, Gary Charness, David Cooper, Dan Friedman, Drew Fudenberg, Rosemarie Nagel, and Al Roth for insightful comments. This paper also benefited from comments received by participants at the ESA meeting in Luxembourg and at the Nordic Workshop in Experimental and Behavioral Economics, and seminar participants at the University of Mannheim, Ewebe in Bertinoro, and the University of Bologna. The usual disclaimer applies. We gratefully acknowledge the financial support from Konkurrensverket and from the ERC grant Strangers 241196.

[†]Department of Economics, University of Bologna E-mail: maria.bigoni@unibo.it

[‡]Department of Economics, University of Bologna. E-mail: marco.casari@unibo.it

[§]Graduate School of Business, Stanford University. E-mail: andy@gsb.stanford.edu

[¶]SITE, Tor Vergata, EIEF & CEPR. E-mail: giancarlo.spagnolo@hhs.se.

1 Introduction

Understanding the determinants of cooperation in social dilemmas is crucial to all social sciences. In many field situations, actors can change actions frequently and asynchronously, a somewhat different situation from that familiar from discretely repeated games. Examples include firms posting prices on Internet or via a centralized and transparent marketplace (as airlines), workers choosing effort in a plant, nearby restaurants choosing menus, and spouses sharing everyday chores. This paper reports results from laboratory experiments on Prisoner's Dilemma games played in (almost) continuous time. We study how cooperation levels change in interactions with different termination rules and of different lengths. This study considers situations with deterministic versus stochastic time horizon (i.e., deterministic versus stochastic termination) and of long versus short expected length (60 and 20 seconds).

The theory of repeated games in discrete time emphasizes the tradeoff between immediate benefits of deviation and future punishments that can be applied only with a delay. This tradeoff is important in many situations involving repeated interactions, for example when bidding in sealed-bid auctions, when traders can sign secret contracts or when producers decide capacity levels that others will discover only at a later time. However, in many other situations the delay in punishment is negligible and may have a little effect on incentives.

When interactions are frequent and players can react quickly, the trade off typical of discretely repeated games may be of second order relevance. In such a case, the assumption of discrete time interaction is not a simple matter of convenience in modelling but may have far reaching implications. Economic theory has provided some answers to the question whether players should behave differently in these different environments, but for finite horizon games it offers conflicting predictions. Running experiments in continuous time allows to empirically investigate what is different in continuous time, and which theories of continuous time games fit best.

Most experimental studies of social dilemmas compare situations where the game is repeated with relatively low frequency. A recent experiment by Friedman and Oprea (forth.) has shown that when actions in a Pris-

oner's Dilemma can be changed at very high frequency, so to approximate a continuous time game, very high levels of cooperation are sustained even when the time horizon is deterministic.

This striking observation led directly to our research question. Would a different termination rule, a stochastic time horizon, generate different results in an (almost) continuous time framework? Friedman and Oprea's results and the leading theories of repeated games do not clearly indicate whether imposing a deterministic or stochastic time horizon would make a difference for continuous time games: theoretical models of cooperation in continuous time do not give an unanimous answer to this question. Providing evidence that the difference is immaterial would blur the line between games of deterministic vs. stochastic duration. On the other hand, if relevant differences emerge, they would be a useful guide for further research on continuous time games.

We report four main results. First, Friedman and Oprea's strikingly high rates of cooperation emerged also in our experiment with deterministic horizon, which provides a robustness check for their findings.

Second, cooperation rates were similar or higher with deterministic horizon than with stochastic horizon in interactions of identical expected duration. In our long duration treatments (60 seconds), cooperation rates are statistically indistinguishable between stochastic and deterministic horizon, while in the short duration treatments (20 seconds) cooperation rates are significantly *higher* with deterministic horizon than with stochastic horizon. These results in continuous time mark a qualitative difference from the findings in Dal Bó (2005) for discretely repeated games, where the author reports the opposite effect of time horizon on cooperation rates.

Third, we find that the within-period pattern of cooperation differed in the stochastic and deterministic horizon treatments. With deterministic horizon, the initial level of cooperation was significantly higher than with stochastic duration, while the final level of cooperation was significantly lower (i.e. lines cross). We are not aware of any theoretical model that would predict a higher initial cooperation with a deterministic horizon. However, since some of the models of behavior in continuous time games yield a folk theorem, it is possible that different time horizon induce subjects to coordinate on different equilibria, and the deterministic con-

tinuous time environment makes it particularly easy to coordinate on the Pareto-efficient one.

Fourth, with deterministic duration, end-of-period effects did not unravel cooperation. In particular, we find that as subjects gained experience the end-of-period effect became less pronounced. This contrasts with discretely repeated games experiments where with finite horizon the end-of-period effect is typically strengthened by experience. It suggests that subjects did not apply backward induction, and instead postponed the end-game-effect closer and closer to the end. Moreover, the end-of-period reversal to defection took place significantly later in the short duration than in the long duration treatments.

As mentioned earlier, these results are not only relevant from a game theoretic or behavioral - experimental point of view. They may explain, for example, why higher prices have been observed in oligopolies when a clear future end-of-the game emerges.¹ And they imply that policies designed for discretely repeated interactions may be ineffective or counterproductive in high frequency environments.²

The next section reviews the related literature; Section 3 discusses the theoretical background; Section 4 describes the experimental design; Section 5 presents our results in detail and Section 6 briefly concludes.

¹Szymanski (1996) noted that the two incumbent shipping companies in the Channel increased prices substantially when the threat of the Eurotunnel taking the best part of their market became real. Assuming a monopolistic market, his model suggested that this happened because of the reduced fear of regulatory intervention given its fixed costs and the fact that the tunnel was expected to soon reduce prices dramatically anyway. However, he admitted he could not explain how this theory could apply to the shipping duopoly that motivated his paper, i.e. why competition among the duopolist did not drive prices down given the Eurotunnel limited the horizon of their interaction.

²For example, Frezal (2006) recently proposed to replace current random industry audits by competition authorities with announced prolonged and intensive audits one industry at a time. By making future collusion impossible for sure for a sufficiently long time period, during the audit, this policy should generate an end-game effect that would make collusion unravel in all markets. Our results suggest that in the electricity auction market and other industries where interaction is highly frequent this policy may actually increase cartel prices.

2 Related Literature

The repeated (or ‘iterated’) Prisoner’s Dilemma with perfect monitoring has probably been the most important set up in which the question ‘what leads people to cooperate’ has been explored experimentally since the early work of Rapoport and Chammah (1965). An important and highly debated issue has been the role played by the time horizon, sometimes called the ‘termination rule’. A large experimental literature has shown that the theoretical prediction that backward induction should apply to finitely repeated games with the features of a Prisoner’s Dilemma often does not hold in the laboratory.³ In field situations, the moment at which a relationship will come to an end is often uncertain. To capture this feature several researchers, starting with Roth and Murnighan (1978) and Murnighan and Roth (1983), have tried to reproduce an indefinite, uncertain horizon in the lab under a stochastic continuation/termination rule for the repeated game. Selten, Mitzkewitz, and Uhlich (1997) argued against the attempt to replicate a potentially infinite horizon in the lab, since no real experiment can have infinite duration and subjects will be aware that the experiment will end in reasonable time, and their beliefs may vary about when exactly. Based on previous experimental evidence (e.g. Selten and Stoecker 1986) they proposed using finitely repeated games, given that the outcome of repeated laboratory games with deterministic and stochastic horizon is similar, apart for the end game effect that only takes place in the last rounds. Dal Bó (2005) offered experimental evidence against this last conclusion. He ran repeated Prisoner’s Dilemma games with two different parameterizations of the stage game payoffs and with deterministic and stochastic horizon with identical but low expected duration. Among other things, he found that cooperation rates in both the first and last rounds of the supergames are significantly lower in treatments with a deterministic horizon than in these with the same expected duration but stochastic horizon. Normann and Wallace (2011) also compared these termination rules (as well as a third, ‘unknown termination’) but in a different set up where the Prisoner’s Dilemma is repeated 22 times before the different termina-

³See e.g. Selten and Stoecker (1986), Andreoni and Miller (1993), Cooper, DeJong, Forsythe, and Ross (1996) and more recently Bereby-Meyer and Roth (2006).

tion rules are introduced, finding no significant differences in cooperation rates.⁴ How far the expected end of the game is seems therefore to play a crucial role for the effects of termination rules in standard repeated games.

By contrast, in Friedman and Oprea (forth.) subjects play a symmetric Prisoner's Dilemma where they could switch actions with latency times on the order of 0.02 seconds for a total period length of exactly 60 seconds, after which the interaction stops with certainty and subjects are rematched in pairs to play another continuous time supergame. Rates of mutual cooperation then reach a median of 90%, and cooperation is typically sustained until the very last seconds of the game when a short but drastic end game effect takes place.⁵ The present study differs from Friedman and Oprea (forth.) because it implemented a series of (almost) continuous time repeated Prisoner's Dilemma supergames both under a deterministic time horizon and under a constant probability of termination generating an identical expected duration of the game. We have treatments with indefinite horizon and look at different expected duration (60 and 20 seconds), but also in other dimensions: we ensure that after each period/match our subjects can never meet again the same opponent (perfect stranger design); and our subjects are asked to choose the starting action rather than having it chosen randomly by the program (and the stage game payoffs are different).⁶

Our work is also related to experimental studies of finitely repeated games played in discrete time at low frequency that, among other things, asked whether subjects learn with experience to apply backwards induction. A consistent finding in this literature, including Selten and Stoecker (1986), Andreoni and Miller (1993), Hauk and Nagel (2001) and Bereby-Meyer and Roth (2006), is that close to the end cooperation rates fall more the more

⁴See also Palfrey and Rosenthal (1994) who compared contributions to a public good in one shot vs. indefinitely repeated games. Engle-Warnick and Slonim (2004) report little differences when comparing a trust game repeated exactly five times vs. repeated with a continuation probability of 0.8.

⁵Charness, Friedman, and Oprea (2011) ran a 4-person public good experiment in continuous time and report a somewhat lower impact of continuous time interaction on cooperation.

⁶An additional difference relative to Friedman and Oprea (forth.) is that in our experiment agents could observe a plot displaying cumulative earnings. In their design, instead, subjects could see the time series of actions both for themselves and the counterpart, and the respective flow payoffs.

subjects gain experience.

3 Theoretical Considerations

The theory of games in continuous time is less developed than its counterpart in discrete time. The topic can be approached from different perspectives; here we sketch three of them that apply to social dilemma games.

A first approach is to treat continuous-time games as the *limit of standard discrete time* games, as each round of interaction is divided into multiple rounds. Hence, agents take more frequent decisions over payoffs that are a fraction of the original ones. When a game is repeated in discrete time, theory predicts that behavior under deterministic vs. stochastic time horizon can be quite different. The standard theory of finitely repeated games in discrete time suggests that cooperation cannot be sustained in equilibrium because of the standard backward induction argument.⁷ In contrast, following the Folk theorems, if future interactions loom sufficiently large, agents can support full cooperation under a stochastic horizon. Hence, under the standard assumptions of rationality and self-regarding preferences, 100 percent cooperation in the initial instant is not sustainable as equilibrium under a deterministic horizon while it can be under a stochastic horizon. This approach predicts that behavior in continuous time games will mirror that in repeated discrete time games.

A second possible approach is to model the games *directly in continuous time*, which entails that deviations can be punished immediately. In continuous time games the backward induction argument breaks down as the real line is not well ordered and a last period cannot be identified even under a deterministic horizon. In other words, in continuous time ‘there is always another period’ in which a deviation can be punished. This setting leads to the prediction that cooperation is an equilibrium regardless of the type of stopping rule or of the length of the interaction.

Finally, the third approach considers discrete-time games with a per-

⁷The induction argument was apparently first made in relation to the finitely repeated Prisoner’s Dilemma by John Nash in private communication reported in Merrill Flood (1952). Models of repeated interaction of stage games with multiple and ranked equilibria that can be used to punish previous defections (Benoit and Krishna, 1985) do not apply to this study.

turbation, which can take several forms. The continuous-time games can be modeled as the *limit of perturbed discrete time* games. This is the approach that prevailed in the literature, and that characterizes the models by Simon and Stinchcombe (1989), Bergin and MacLeod (1993), Radner (1986), Friedman and Oprea (forth.), and Kreps, Milgrom, Roberts, and Wilson (1982).

Simon and Stinchcombe (1989) build a general model of games a finite number of actions and players. When placing some restrictions on strategies, they prove that a prisoner's dilemma played in continuous time admits a unique equilibrium of full cooperation, a prediction that is stronger than the counterpart of Folk Theorems for an infinitely repeated game in discrete time. More in detail, they define the game on a discrete grid in a finite interval and then let the grid interval go to zero, and assume that each strategy admits a uniformly bounded number of moves in the game. Reasoning by backward induction, they obtain that cooperation is typically sustainable in subgame perfect equilibrium and that for the Prisoner's Dilemma full cooperation is the unique equilibrium surviving iterated deletion of weakly dominated strategies. The intuition behind this result is that no player would ever switch from defection to cooperation, when she has only one move left. So if both players can react with a delay that tends to zero, and can switch action at least once in the game, the game will never end in one of the two asymmetric outcomes.⁸ Strictly speaking, this theory suggest that we should not observe sizable end game effects in a continuous time game played under a deterministic horizon.

Bergin and MacLeod (1993) build a related model that includes a degree of inertia in changing actions as interactions are structured in a sequence

⁸The two possible final states are such that either both players have at least one move available and they both cooperate, or they both defect and they have zero or one move left each, or an uneven number of moves. Suppose for example that at some point in time t the two players have the same number n of moves remaining, with $n \geq 2$, and that they both defect. It is optimal for them to switch immediately to cooperation, and to switch back to defection only if they realize that the opponent has not done the same. If instead they both cooperate, and they both have one move available, no one has an incentive to switch to defection, as the gains from defection would be incomparably low with respect to the foregone payoffs from cooperation. Answering the question whether continuous time games have discrete repeated games analogues the authors write "yes, provided that agents payoffs are insensitive to the actions other agents choose near the end of the game" (p. 1200) meaning that continuous time games are intrinsically different than discretely repeated ones precisely with respect to end game effects.

of intervals from t to $t+\epsilon$. They characterize the set of ϵ -subgame-perfect-equilibria and then let ϵ go to zero. This leads to a full Folk Theorem for the continuous time Prisoner's Dilemma that holds for both deterministic and stochastic horizons. The intuition behind these predictions is that if players adopt a trigger strategy that punish a defection after a time interval of size ϵ , the magnitude of the gains of defection also is of the order of ϵ . Thus, as ϵ approaches zero, the incentive to deviate also vanishes. Because of the multiplicity of equilibria, this theory has a weaker predictive power and is consistent with a large number of equilibrium paths observed in the lab, with or without a sizable end game effect.

Radner (1986) puts forward a theory of bounded rationality in discretely repeated games with finite horizon based on ϵ -equilibria (recently adopted and extended by Friedman and Oprea (forth.) to explain their results). He predicts full initial cooperation as long as there is a small probability that the opponent plays a "cooperative" dynamic behavioral strategy. His behavioral restriction is to a class of strategies of the form "cooperate until period k or until the other player defects and defect otherwise," so-called cut-off strategies. He notes that if the players can react swiftly to a defection of the other player, the losses that a player may incur using a cut-off strategy with a very large k are bounded to be very small, while the same strategy allows large gains from prolonged cooperation if the opponent uses a cut-off strategy with a large k . The best response strategy, defect at $k-1$ if the other player waits till k , leads to backward induction and unraveling. Relative to the safe but low non-cooperation payoffs obtained using best reply and the induction argument they trigger, the cooperative strategies become more and more attractive when the number of repetitions grow. This implies that that cooperation can be sustained in deterministic horizon games with many periods or frequent actions if subjects realize that continuing cooperating rather than defecting produces large expected benefits compared to the risk of small losses one is exposed to. This argument applies of course to a stronger extent to continuous time games, as stressed by Friedman and Oprea (forth.), and may be consistent with a small end game effect at the end of the finite horizon. The timing of the end-game effect depends on how far the horizon is, and the reaction time. More specifically, this model predicts that the switch to permanent defection

takes place later, for games with a longer duration and for shorter reaction times.

This last approach also includes the 'gang of four' paper for discrete-time games under a deterministic horizon (Kreps, Milgrom, Roberts, and Wilson, 1982), which can be brought to the limit and extended to continuous time games, without changing the predictions.⁹

None of these theories, however, offer testable predictions on differences in initial, average or median cooperation rates between deterministic and stochastic horizons in continuous time games with the same expected duration. For this reason it is useful to look also at predictions about patterns of behavior within a period. Kreps, Milgrom, Roberts, and Wilson (1982) and Radner (1986) predict an initial high cooperation and a sudden fall in cooperation as the end of the game approaches, while Simon and Stinchcombe (1989) does not predict an end-of-period effect. Moreover, Kreps, Milgrom, Roberts, and Wilson (1982) predict that the duration of the end game effects should be independent of the length of a game and depend on the players' beliefs about the opponent's type, which may change with experience. By contrast, the model by Radner (1986) and its extension by Friedman and Oprea (forth.) predicts that the faster the reaction time, the later appears an end-game effect. On the other hand, they offer no prediction about the impact of experience or length of a game on the end-game effect. Table 1 summarizes the theoretical predictions of the above models.

In brief, the experiment aims at studying three issues for continuous-time games: i) which of the above approaches better predict the differences in overall cooperation levels between a deterministic vs. a stochastic time horizon. ii) which of the above approaches better predict the patterns of behavior within a period. iii) which, if any, patterns observed in the data are inconsistent with all of the above approaches.

4 Experimental Design

The experiment has a two-by-two factorial design. The two treatment variables are the expected duration of each period and the termination

⁹See Mailath and Samuelson (2006) for an excellent survey of models of incomplete information and reputation formation in repeated games.

	Bergin & McLeod (1993)	Radner (1986) Friedman & Oprea (forth.)	Simon & Stinchcombe (1989)	Kreps et al. (1982)	Discrete-time games
Stochastic horizon					
Full initial cooperation is an equilibrium	Yes	Yes	Yes	Yes	Yes
Deterministic horizon					
Full initial cooperation is an equilibrium	Yes	Yes	Yes	Yes	No
End-game effect with det. horizon					
Predicted	No	Yes	No	Yes	–
Pattern emerging with learning	–	Decreasing	–	Unclear, or no effects	Increasing

Table 1: Main theoretical predictions of different models.

rule. Table 2 summarizes the characteristics of each treatment.

In all treatments, subjects played a series of (quasi) continuous time prisoners' dilemmas.¹⁰ Each session comprised a non-overlapping group of 24 subjects, who interacted in pairs for 23 periods. Pairs were formed so that each subject met all the others once and only once in a session (perfect strangers).¹¹

In all treatments, the stage game was as follows. Each subject had to select an initial action for the period between Cooperate (green) and Defect (orange). When all subjects were done, the period began. Within a period, subjects could switch action up to six or seven times per second. More precisely there was a tick every 16/100th of a second, which gave the participants the feeling of continuous time. The PCs had touch screens,

¹⁰As the instructions explained, the experiment was in quasi continuous time: "Within a period, both you and the other will be able to change action as many time as you wish. The time flows in very rapid ticks (of 16th hundredth of a second); in practice there are between six and seven ticks every second, so that if you wish you can change action six or seven times per second." For shortness, from now on we will talk about continuous time experiment.

¹¹In the Short-Deterministic session run on February 2, 2011, due to a technical problem in period 23 subjects met again their opponents of period 1. All reported results hold even if period 23 in that session is dropped.

	Termination rule	
	Deterministic	Stochastic
Short (20 secs.)	N=48 Period endowment: 15 pts. Conversion rate: 50 pts.=1 € - January 24, 2011 - February 4, 2011	N=48 Period endowment: 15 pts. Conversion rate: 50 pts.=1 € Average realized duration: 22.6" - February 2, 2011 - February 4, 2011
Long (60 secs.)	N=48 Period endowment: 50 pts. Conversion rate: 150 pts.=1 € - October 21, 2010 - October 28, 2010	N=48 Period endowment: 50 pts. Conversion rate: 150 pts.=1 € Average realized duration: 68.3" - October 22, 2010 - October 28, 2010


Table 2: Treatments and sessions

hence a switch of action could not be heard by others as subjects simply touched the screen with a finger.

Earnings for all possible combinations of actions were visible on the screen at all time (Figure 1). The payoff matrix showed earnings in tokens per second. The subject's current action was always highlighted in yellow in the payoff matrix. Moreover, every subject could observe her cumulative earnings on a continuously updated graph (Figure 1). Subjects' earnings in every period included an initial endowment (see Table 2), and could stay constant, increase, or decrease over time depending on the choices of the pair. The graph showed these patterns of earnings as a flat, increasing, or decreasing line, respectively. A steeper line indicated a faster accumulation or depletion. The line color was green or orange depending on the subject's own action. Hence, from the graph subjects could unambiguously infer the action taken in any moment by their opponent. The progression of the earnings line marked the timing of the period for the subjects. They could observe at every instant the speed of the game, which ran at the same pace for all subjects in the session. For the Deterministic treatments subjects could always check the time remaining before the end of a period, by looking at the graph on the screen.

In the Long-Deterministic treatment, a period always lasted 60 seconds. In the Long-Stochastic treatment, a period lasted in expectation 60 seconds. Similarly for the short treatments, where the expected duration

Figure 1: Screen-shot of the stage game for the Long treatments



Notes: VERDE = green, ARANCIO = orange, l'azione dell'altro = your opponent's action, guadagno = earnings.

was 20 seconds. In the stochastic treatments, the exact duration was selected at random period by period. As explained in the instructions for the Long(Short)-Stochastic treatment, the period duration depended on a random draw. “Imagine a box with 10,000 (1000) balls, of which 9,973 (992) are black and 27 (8) are white. It is as if a ball is drawn after every tick. If the ball is white, the period ends. If the ball is black, the period continues and the ball is put back into the box. At the next tick, another ball is drawn at random. You have to imagine very fast draws, i.e. one every tick of 16 hundredth of a second. As a consequence of this procedure, we have estimated that periods will last on average 60 (20) seconds. There may be periods that are short and periods that are long.” In case a period lasted beyond 60 seconds, the time-line in the graph automatically shifted forward.

Stage game payoffs are such that cooperation should be easily achieved (at least in the stochastic ending treatments). In continuous time cooperation is always supportable because the instantaneous discount factor is 1: then a grim trigger strategy should in theory always support cooperative play as an equilibrium no matter the arrival rate of the end of the game. But even if agents perceived the game to be played discretely, e.g. because of minimal human reaction time, cooperation should be easily sustained with our parameterization. For example, if subjects react with 1 second delay and treat it as a time interval length of 1 second, then, given our stage game payoffs (see Figure 1), cooperation can be sustained with infinite horizon for discount factors higher than $1/2$, which implies an expected duration of 2 seconds. If the time interval length is 0.25 of a second, then it would be enough to have an expected duration of 0.5 of a second, and so on. Hence the 20 seconds is quite far from the theoretical bound.

Instructions were distributed and then read aloud. Subjects had the opportunity to ask questions, which were answered in private, and then went through three practice periods with a robot opponent that was programmed to switch action in the middle of the period. After the practice period, subjects had to guess the actions taken by the robot, and then completed an on-line quiz to verify their complete understanding of the rules of the game. The experiment started as soon as all subjects answered

correctly to all the four control questions.¹² The session ended with a questionnaire.

Subjects were 192 students at the University of Bologna (primarily), who were randomly assigned to one of the four sessions using an on-line recruitment software (Greiner, 2004). The experiment was run in the Bologna Laboratory for Experiments in Social Sciences using z-Tree (Fischbacher, 2007). Subjects seated at visually isolated computer terminals and could not communicate. A session lasted on average 2 hours for the Long treatments and 1 hour and 10 minutes for the Short ones. Subjects earned on average 16.72 Euros, and 14.81 Euros, respectively, which include a show-up fee of 3 Euros.

5 Results

The presentation of results is organized in three parts. We first provide a comparison of the aggregate cooperation rates across the four treatments (Result 1). We then report about patterns of cooperation within a period (Results 2 and 3). Finally, we describe and comment the effects of experience on cooperation rates (Result 4).

5.1 Aggregate cooperation rates across treatments

Our Long-Deterministic treatment replicates and extends the results reported in Friedman and Oprea (forth.) for different payoff levels. They report only the median cooperation rate after period 12, which ranges from 81% to 93%. Our median cooperation rate is 84.0% over all periods, and exhibits an increasing trend with experience. If we consider only periods after the twelfth, the median rate of cooperation in our data is 91.2%. We provide a robustness check on Friedman and Oprea's results by increasing the number of subjects per session, which allows us to adopt an absolute stranger matching protocol. This design feature reduces repeated game effects within a session.

¹²In the three practice periods, 71% of the subjects always made correct guesses about the sequence of actions taken by the robots. In answering the xx control questions about the instructions, 50% of the subjects made at most one mistake.

The novelty of this study stems from the comparison across all our four treatments.

Result 1 *When period duration is deterministic, cooperation rates are similar or higher than in the case of stochastic duration.*

Support for Result 1 comes from Tables 3 and 4. The unit of observation is the cooperation rate which is defined as the fraction of time R_{ip} a subject i spends cooperating within period p . Given that these observations are not independent, Table 4 compares results across treatments through a panel regression with random effects at the subject level and standard errors robust for clustering at the session level.¹³

Duration	Termination rule	
	Deterministic	Stochastic
Long	65.5 (84.0) √***	~ 66.9 (84.8) √***
Short	63.3 (79.2)	>*** 52.3 (47.0)

Notes: Median cooperation rates are reported in parentheses. The unit of observation is a subject per period. The mean cooperation rate of a session is the average across all 23 periods and all 24 subjects. There are two sessions per treatment, thus N=1104.

Table 3: Cooperation rates

Result 1 holds both for short and for long duration treatments. In the long duration treatments, cooperation rates are statistically indistinguishable between stochastic and deterministic duration (p-value > 0.1, see Table 4). The absolute difference between the two treatments is just 1.4 points in terms of means, and 0.8 points in terms median. By contrast, in the short duration treatments cooperation rates are significantly higher with deterministic duration than with stochastic duration (p-value < 0.001, see Table 4). The absolute difference in cooperation between the

¹³We obtain qualitatively and quantitatively similar results with linear regressions and standard errors robust for clustering at the subject and pair level. We perform the same robustness check for all regressions in the main text and in the Appendix. Results in Tables 4, A.1, and A.2 are also robust to collapsing all the data by subjects and then running the test clustering by session.

two treatments is 11.0 points in terms of means, and 32.2 points in terms of median.

In addition we report that shortening expected period duration from 60 to 20 seconds can have a dramatic impact on cooperation rates. In the stochastic treatments cooperation rate drops by 14.6 points in terms of mean, and by 37.8 in terms of median, and the difference is highly significant (p-value < 0.001, Table 4). To our surprise, the difference between the long and the short deterministic treatments is much smaller, though significant (p-value < 0.01, Table 4). The cooperation rate drops by 2.2 points in terms of mean, and by 4.8 in terms of median.¹⁴

Dependent variable: cooperation rate		
	Coefficient	(s.e.)
Short-Deterministic	-6.082***	(2.000)
Long-Stochastic	1.998	(1.449)
Short-Stochastic	-17.755***	(2.872)
Constant	62.600***	(11.371)
Controls for individual characteristics	Yes	
N	4416	
R-squared overall	0.047	
R-squared between	0.223	
R-squared within	0.000	

Notes: Panel regression with random effects at the subjects' level and standard errors robust for clustering at the session level. The unit of obs. is the fraction of time a subject spends cooperating within a period. Default treatment: Long-Deterministic. The difference between coefficients for the Short-Stochastic and Short-Deterministic treatment is significant at any standard significance level (p-value < 0.001).

Table 4: Panel regression on cooperation rates

These results complement the findings reported in Dal Bó (2005) for games in discrete time. He reports that cooperation is higher with stochastic than with deterministic duration.¹⁵

¹⁴Regression results in Table 4 show that the estimated difference between the Short-Deterministic and the Long-Deterministic treatments is actually slightly larger if we control for individual characteristics.

¹⁵In a repeated game with a much shorter expected duration (the expected number of action choices is 125 in our short treatments, 375 in our long treatments, while it ranges between 2 and 4 in his treatments), Dal Bó finds that, “for every round, [...] the percentage of cooperation in infinitely repeated games [...] is greater than in finitely repeated games of the same expected length [...], with p-values of less than 0.01.” More

5.2 Patterns of cooperation within a period

To deepen the analysis of the difference across treatments, we now turn to the patterns of cooperation within a period. In continuous time, there can be rich dynamics within each period, as the same cooperation rate R_{ip} can result from many different sequences of actions. It turns out that the pattern of cooperation within periods presents some strong regularities, as reported in Results 2 and 3.

Result 2 *With deterministic duration, cooperation does not unravel due to end-of-period effects. End-of-periods effects exist but they arise later and later as subjects gain experience.*

Result 2 suggests that subjects do not apply backward induction, and learn to postpone more and more the end-game-effect.

Support for Result 2 comes from Figure 2 and Table 5. A subject can change action every 0.16 seconds. Figure 2 presents the time profile of the mean share of cooperators, taken across periods and sessions. Our unit of observation is the share of cooperators S_{tp} at a given time interval t of 0.16 seconds, within a period p .

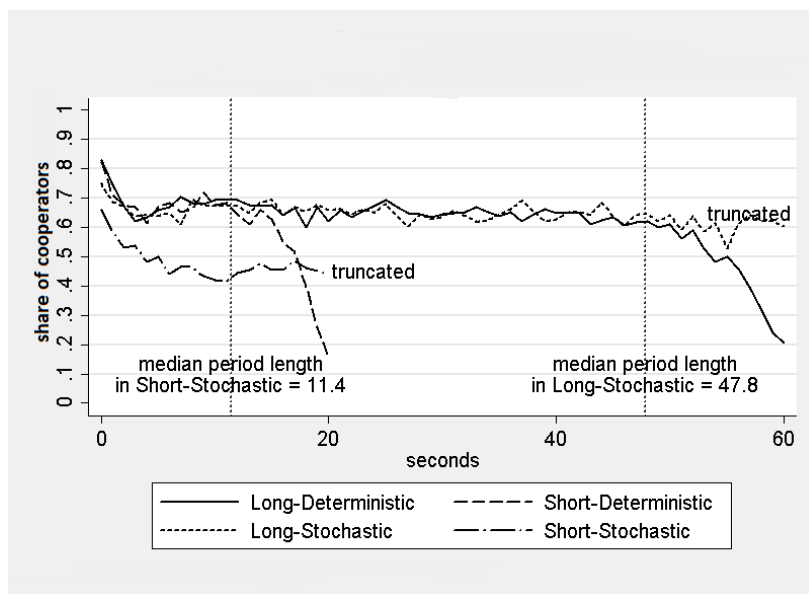
As seen in this figure, subjects facing periods with deterministic duration exhibit a clear end-of-period effect: the share of cooperators suddenly drops a few seconds before the end of the period.

There are of course many ways to quantitatively measure the timing of such switch from cooperation to defection. We describe below one way to measure it that takes as reference all pairs that at some point during a period reached simultaneous cooperation, CC. Out of those pairs, we consider in the calculation only those that switched to defection before the end of the period, i.e. CD, DC, or DD, which were the lion's share of the observations.¹⁶ On average, the end-of-period effects kicked in 3.4 seconds before the end of the period. Table 5 reports in more details the

specifically, when the expected duration is 2 (4) periods, the average cooperation rate is 28.3% (35.2%) with stochastic ending and 12.5% (24.8%) with deterministic ending.

¹⁶468 subjects-period out of a total of 552 in the Long treatment and 460/552 in the Short treatment cooperated simultaneously at least once in a period. Of these, some (54/468 and 51/460, respectively) kept on cooperating until the end of the period, while in others (414/468 and 409/468, respectively) at least one of the subjects in the pair switched to permanent defection.

Figure 2: Time profile of the share of cooperators



Notes: The graph includes the first 60 seconds for Long treatments and the first 20 seconds for Short treatments. The unit of observation is the share of cooperators at a give time interval of 0.16 seconds within a period. All periods and all subjects are included.

Treatment	Periods			
	1-8	9-16	17-23	Overall
Long-Deterministic	5.8	4.1	3.4	4.4
	<i>N=119</i>	<i>N=134</i>	<i>N=134</i>	<i>N=387</i>
Short-Deterministic	3.0	2.4	2.0	2.4
	<i>N=107</i>	<i>N=152</i>	<i>N=150</i>	<i>N=409</i>

Notes: the table reports the number of seconds before the end of the period when a pair in CC permanently switches to defection, i.e. either CD, DC, or DD. In the Long-Deterministic treatment we drop observations in which the end game effect kicks in more than 20 seconds before the end of the period (27 out of 414 of the observations). This was done to preserve comparability with the Short-Deterministic treatment.

Table 5: Timing of the end-of-period effect

mean number of seconds from the end of the period, when this switch to permanent defection took place.

Table 5 shows that, with experience, the end-of-period effect kicks in later in time (1 to 2.4 seconds later). This effect of experience is significant both in the Long-Deterministic (p-value < 0.001) and in the Short-Deterministic treatment (p-value < 0.05).¹⁷ In addition, in the Short-Deterministic treatment the end-of-period effect kicks in significantly later than in the Long-Deterministic treatment (Table 5, p-value < 0.05).¹⁸ Friedman and Oprea (forth.) also report an end-of-period effect. They find that “cooperation level falls below 75 percent only when 5 seconds remain and below 50 percent only when 1 second remains.”

Result 3 *The share of cooperators displays a different time profile in the stochastic vs. the deterministic treatment. With stochastic duration, the initial share of cooperators is lower, while the final share is higher than*

¹⁷P-values obtained from a linear regression with standard errors robust for clustering at the session level. The dependent variable is the timing of the end game effect (in seconds), and among the independent variables we include a dummy taking value 1 for the treatment with short duration, the variable “Period” and their interaction. The regression is not reported here. Regression’s results are available from the authors upon request.

¹⁸P-value obtained from a panel regression with random effects at the subjects’ level and standard errors robust for clustering at the session level. The dependent variable is the timing of the end game effect (in seconds), and the only independent variable is a dummy taking value 1 for the treatment with short duration. The regression is not reported here. Regression’s results are available from the authors upon request.

with deterministic duration (i.e. lines cross).

Support for Result 3 comes from Figure 2 and Table 6. Figure 2 traces the time profile of the share of cooperators second by second and is self-explanatory. Table 6 summarizes frequencies of actions at the initial and final instant of a period and displays the stark contrast between the stochastic and deterministic treatments that is highlighted in Result 3. With stochastic ending, in the relative majority of pairs the two subjects cooperated both in the initial and in the final instant of a period (cells in bold in Table 6 (a) and (d)). On the contrary, with deterministic ending, in the relative majority of pairs both subjects cooperated in the initial instant and defected in the final instant of a period (figures in bold in Table 6 (b) and (c)).

In general, more subjects chose cooperation as their initial action in the Long-Deterministic than in the Long-Stochastic treatment (82.5% vs. 75.1%, Table 6 (a) and (b)). On the contrary, less subjects chose cooperation as their final action in the Long-Deterministic than in the Long-Stochastic treatment (18.1% vs. 64.9%, Table 6 (a) and (b)).

A logit regression on final cooperation shows that this difference is significant at the 1% level (Table A.2 in the Appendix). A similar pattern characterizes results from the short treatments. Initial cooperation is 82.6% vs. 65.9%, and final cooperation is 15.8% vs. 46.8%, respectively (Table 6, (c) and (d)). Logit regressions on initial and on final cooperation show that both differences are significant at the 1% level (Tables A2 and A3 in the Appendix).¹⁹

Result 3 suggests that different termination rules induce the adoption of different individual strategies. This interpretation is supported by the data presented in Table 7, which reports the fraction of subjects whose behavior follows one out of five simple patterns: (i) always defect, (ii) always cooperate, (iii) start defecting then switch to permanent cooperation, (iv) start cooperating and switch to permanent defection when the opponent cooperates (leader), (v) start cooperating then switch to permanent defection when the opponent is defecting (follower). These patterns describe between

¹⁹Support for Result 3 becomes stronger as subjects gain experience. Figure 4 and Figure A.1 in Appendix illustrate initial and final cooperation rates across periods.

Long-Deterministic (a)						Long-Stochastic (b)					
	opponent's initial-final actions						opponent's initial-final actions				
actions	D-D	D-C	C-D	C-C	Total	actions	D-D	D-C	C-D	C-C	Total
D-D	4.3	0.1	10.8	0.5	15.7	D-D	4.3	0.6	8	1.4	14.4
D-C		0.2	1.2	0.4	1.8	D-C		1.4	0.8	7.6	10.5
C-D			47.6	6.6	66.2	C-D			10.1	1.7	20.7
C-C				8.9	16.3	C-C				43.7	54.4
Total					100	Total					100

Short-Deterministic (c)						Short-Stochastic (d)					
	opponent's initial-final actions						opponent's initial-final actions				
actions	D-D	D-C	C-D	C-C	Total	actions	D-D	D-C	C-D	C-C	Total
D-D	3.8	0.4	10.2	0.6	15	D-D	9.8	1.1	13	1.4	25.3
D-C		0.5	0.7	0.7	2.4	D-C		3.1	1.2	3.4	8.8
C-D			53.4	4.8	69.2	C-D			11.4	2.4	27.9
C-C				7.2	13.4	C-C				30.8	38
Total					100	Total					100

Notes: The first letter denotes the initial action, the second letter the action taken when the period ends. For instance D-C denotes subjects who chose defect as their initial action and chose cooperate when the period ended. The fraction of subjects who initially cooperated are the sum of the figures in the “total” row for columns C-D and C-C. The matrices are symmetric and for easier reading the lower triangle has been left blank. The largest figure in each matrix is in bold.

Table 6: Initial and final actions

	Treatment			
	Long Det.	Long Stoch.	Short Det.	Short Stoch.
(i) always D	5.2	4.4	2.8	13.9
(ii) always C	13.0	33.6	8.8	32.7
(iii) start D then C	1.1	2.4	0.5	4.3
(iv) start C then D (leader)	13.6	2.7	18.8	8.2
(v) start C then D (follower)	21.2	4.0	24.4	9.3
(vi) multiple switches	46.0	52.9	44.8	31.5
Total	100.0	100.0	100.0	100.0

Notes: percentage points. The unit of observation is a subject per period, N=1104.

Table 7: Individual patterns of choices.

47.1% and 68.5% of subjects, depending on the treatment. Subjects following (i) and (ii) never switch actions, while subjects following (iii), (iv), and (v) switch action exactly one time. Table 7 also reports the residual category (vi) where a subject engaged in multiple switches between C and D, or vice versa.

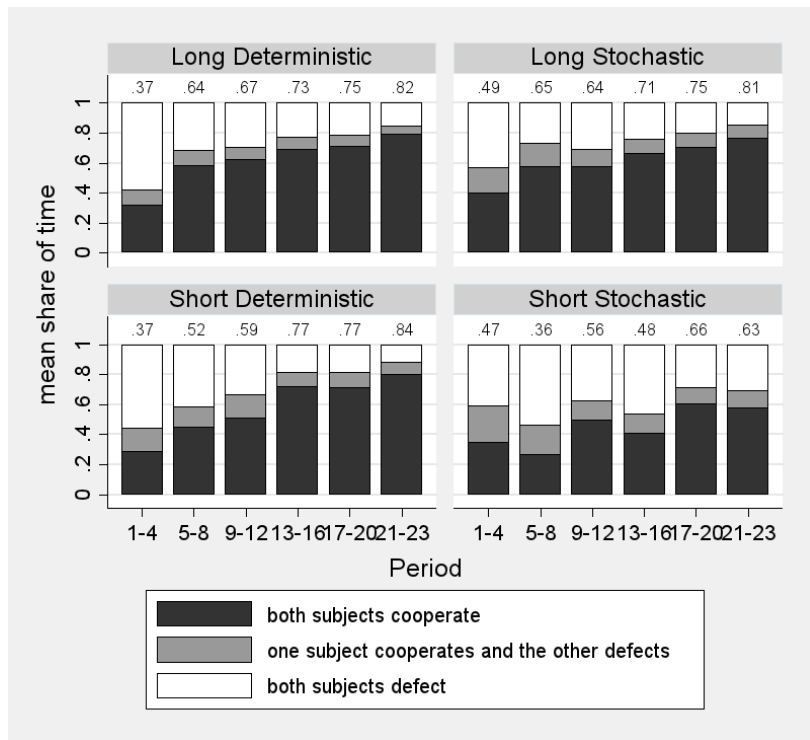
Table 7 shows that subjects use qualitatively different strategies in the stochastic versus deterministic duration treatments. This evidence confirms the intuition gained from Table 5 when one traces a subject's choices in every instant within a period and can hence detect switches and volatility in behavior. With stochastic duration about three times as many subjects follows "always cooperate" than with deterministic duration (33.1% vs. 10.9%, column (ii) averaged across treatments). With stochastic duration, about one third as many subjects follow "start cooperating then switch to permanent defection" than with deterministic duration (12.1% vs. 39.0%, columns (iv) and (v)). Two comments are in order. On the one hand, this evidence can explain why the lines cross especially in the long treatment: a larger fraction of subjects starts cooperating in the deterministic treatment and then switch to defection. On the other hand, with periods of short duration, a high fraction of subjects always defected in the stochastic treatment. This evidence can explain why the gap in cooperation rates between the short treatments is larger than between the long duration treatments.

5.3 Effects of experience on cooperation

In this subsection, we look at how the level of cooperation evolves across periods, as subjects gain experience.

Result 4 *Cooperation rates increase with experience in all treatments.*

Figure 3: Cooperation and experience



Notes: Mean share of time spent in each of the three outcomes. One observation per couple, per period.

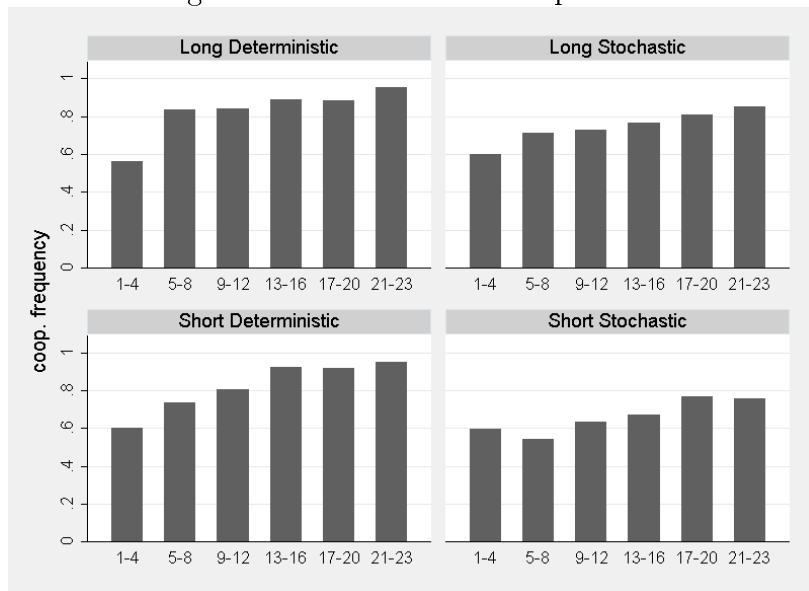
Support for Result 4 comes from Figure 3. On top of each bar, we report the average of the cooperation rate R_{ip} of subject i in period p , across blocks of four periods, and across subjects. In all treatments there is an upward trend, although this trend is weaker with stochastic duration, especially in the Short-Stochastic treatment. More detailed evidence comes from a panel regression reported in the Appendix (Table A.3). Our result for the Long-Deterministic treatment is consistent with Friedman and Oprea (forth.),

who also find that cooperation rates rapidly increase in the first 16 periods and settle afterwards.

In addition, Figure 3 also reports the average share of time in which both subjects in a pair simultaneously cooperated (CC), simultaneously defected (DD), or chose different actions (CD). In all treatments, the share CD decreases with experience (Figure 3).

Dal Bó and Fréchette (forthcoming) suggest that the level of initial cooperation is an additional, important measure of cooperation in treatments with stochastic duration. The reason being that that generally the cooperation rate changes with period duration and periods usually have different durations. Our Result 3 is confirmed when looking at the initial rate of cooperation (Figure 4). The upward trend in the level of initial cooperation is significant in all treatments (Table A.4 in the Appendix). In the Short-Stochastic treatment this upward trend emerges only in the second half of the session (Table A.4), and the overall rate of initial cooperation is significantly lower than in the other treatments (p-value 0.01, see Table A.1).

Figure 4: Rates of initial cooperation



Our results on the impact of experience on cooperation levels are consistent with the findings of Dal Bó and Fréchette (forthcoming). When playing repeated games, the amount of experience is a critical determinant

of outcomes and it takes more than 10 repetitions to settle on a stable level of cooperation.

6 Conclusions

Through an experiment in (quasi) continuous time, we have studied prisoner's dilemma games under deterministic and stochastic horizons. By comparing the data with the theoretical predictions, one can draw several conclusions.

With long periods (60-second expected), we report high levels of cooperation that were similar under a deterministic vs. a stochastic time horizon. This result is compatible with theoretical approaches that model continuous-time games directly or as the limit of discrete-time games with some perturbation.

Instead, with short periods (20-seconds expected) a deterministic horizon led to strictly higher rates of cooperation than a stochastic horizon. This finding is novel and contrasts with existing experimental results about social dilemmas in discrete time where either it is harder to sustain cooperation with a deterministic horizon than a stochastic horizon, or rates are similar. We are not aware of any theoretical approaches to continuous-time games that accounts for this finding. Another result is that with stochastic duration overall rates of cooperation were lower in short than in long periods.

For additional insights one can look at patterns of behavior *within* a period. The time horizon significant impact on the strategies employed and the dynamics of cooperation. With deterministic duration there was a dramatic end-game-effect. Subjects employed cut-off strategies such as "Cooperate until time T and then defect forever." Theoretical models such as those in Radner (1986) and Kreps, Milgrom, Roberts, and Wilson (1982) are compatible with an end-of-period effect. This reversal from cooperation to defection, however, took place later in time the more subjects gained experience within a session. This suggests that subjects are learning to cooperate more close to the end rather than less, as is instead typically observed for discretely repeated games. Moreover, the reversal took place later in short periods than in long ones. None of the theoretical approaches

reviewed here can account for this impact of period length on the timing of the end-game effect. If confirmed by other experiments these findings will be another interesting paradox for the theory to explain.

More generally, under a deterministic horizon cooperation rates within a period were initially higher than under a stochastic horizon, and they were lower toward the end (i.e. lines cross). In the experiment, a deterministic horizon seemed to facilitate an initial coordination on cooperation.

From a theoretical viewpoint these results suggest that behavior in continuous time games is not simply the limit of *standard* discretely repeated games. The canonical theory of repeated games in discrete time focuses on how deviations create a tradeoff between current benefits and future punishments. Our results are consistent with the continuous games theories of Simon and Stinchcombe (1989) and Bergin and MacLeod (1993), where they suggest that this tradeoff is of second order importance when players can react quickly so that the time horizon is not a crucial determinant of cooperation. They support even more strongly theories that predict high cooperation rates followed by end horizon effects, like Kreps, Milgrom, Roberts, and Wilson (1982), Radner (1986) and Friedman and Oprea (forth.). The theoretical approaches that model the game *directly* in continuous time do not find support in our data because they predict no differences across treatments in the ability of subjects to coordinate on cooperation and do not predict an end-of-game effect.

Taken together, these findings may have important implications for a variety of field applications. People facing social dilemmas in which they can react swiftly, as in many productive, labor, sporting, and military activities can easily overcome the challenge of achieving mutual cooperation irrespective of the deterministic or stochastic horizon of the interaction even for short duration activities. In those situations a deterministic horizon is not an impediment to cooperation and may even facilitate it. On collusion practices, our results may explain why higher prices have been observed in oligopolies when the date of the last interaction is made public. This suggests that competition policies designed for discretely repeated interactions may be counter-productive in continuous time.

To draw implications from the experimental results, however, one should keep in mind that these activities must share some well-defined features:

they should involve a continuous time effort by both participants as when carrying together a heavy object or jointly rowing in a boat and participants must perfectly observe the action or effort taken by the opponent. Further work is needed to understand the domain of application of these results, for instance with respect to shorter period lengths or other details. In particular, the introduction of imperfect monitoring of the opponent's action may limit, or remove all-together, the possibility to sustain a cooperative outcome when actions are chosen frequently (as in the theoretical results in Sannikov and Skrzypacz (2007)).

A Additional tables

Dependent variable: initial cooperation (0/1)		
	Marginal effect	(s.e.)
Short-Deterministic	-0.040	(0.044)
Long-Stochastic	-0.025	(0.043)
Short-Stochastic	-0.184***	(0.044)
Controls for individual characteristics	Yes	
N	4416	
Log-likelihood	-1889.180	

Notes: Marginal effects from a logit regression with random effects at the subjects' level and standard errors robust for clustering at the session level. The unit of obs. is the decision of a subject to initiate a period by cooperating (1) or defecting (0). The difference between coefficients for the Short-Stochastic and Short-Deterministic treatment is significant at the 5% significance level (p-value = 0.011).

Table A.1: Impact of the treatment on initial cooperation

Dependent variable: final cooperation (0/1)		
	Marginal effect	(s.e.)
Short-Deterministic	-0.061	(0.043)
Long-Stochastic	0.406***	(0.033)
Short-Stochastic	0.226***	(0.040)
Controls for individual characteristics	Yes	
N	4416	
Log-likelihood	-2324.573	

Notes: Marginal effects from a logit regression with random effects at the subjects' level and standard errors robust for clustering at the session level. The unit of obs. is the decision of a subject of cooperating (1) or defecting (0) in the last instant of a period. The difference between coefficients for the Short-Stochastic and Short-Deterministic treatment is significant at any standard significance level (p-value < 0.001).

Table A.2: Impact of the treatment on final cooperation

Dep. variable: cooperation rate	Treatment			
	Long-Det.	Short-Det.	Long-Stoc.	Short-Stoc.
periods 5-8	26.672*** (1.112)	15.354*** (0.594)	17.902*** (2.714)	-8.314*** (2.893)
periods 9-12	29.622*** (0.452)	22.521*** (2.244)	15.150*** (3.262)	12.812*** (1.568)
periods 13-16	36.357*** (8.431)	40.313*** (0.206)	23.394*** (3.625)	2.648 (5.371)
periods 17-20	37.953*** (6.720)	39.996*** (6.986)	26.279*** (3.109)	21.715*** (6.875)
periods 21-23	45.294*** (8.695)	47.674*** (4.538)	34.568*** (2.567)	19.323*** (2.728)
period's duration			-0.060 (0.039)	-0.217 (0.190)
duration of the previous period			0.069*** (0.010)	0.142*** (0.015)
Constant	35.977* (19.348)	45.981*** (7.034)	17.332 (15.294)	3.626 (11.833)
Controls for individual characteristics				
N	Yes 1104	Yes 1104	Yes 1104	Yes 1104
R-squared overall	0.217	0.292	0.217	0.157
R-squared between	0.329	0.496	0.487	0.355
R-squared within	0.191	0.265	0.127	0.107

Notes: Panel regression with random effects at the subjects' level. Standard errors robust for clustering at the session level. The unit of obs. is a subject in a period.


Table A.3: Panel regression on cooperation and experience

Dep. variable: initial cooperation	Treatment			
	Long-Det.	Short-Det.	Long-Stoc. Short-Stoc.	
periods 5-8	0.179*** (0.029)	0.078*** (0.027)	0.087** (0.035)	-0.068 (0.042)
periods 9-12	0.189*** (0.030)	0.134*** (0.029)	0.088*** (0.034)	0.028 (0.043)
periods 13-16	0.242*** (0.035)	0.262*** (0.037)	0.137*** (0.036)	0.069 (0.043)
periods 17-20	0.236*** (0.034)	0.254*** (0.036)	0.176*** (0.038)	0.187*** (0.046)
periods 21-23	0.355*** (0.055)	0.321*** (0.049)	0.267*** (0.048)	0.177*** (0.050)
duration of the previous period			0.000* (0.000)	0.000 (0.000)
Controls for individual characteristics	Yes	Yes	Yes	Yes
N	1104	1104	1104	1104
Log-likelihood	-348.876	-393.446	-380.489	-529.683

Notes: Marginal effects from a logit regression with random effects at the subjects' level. Standard errors robust for clustering at the session level. The unit of obs. is the decision of a subject to initiate a period by cooperating (1) or defecting (0).

Table A.4: Impact of experience on initial cooperation by treatment

Figure A.1: cooperation rates at the end of the period.



B Instructions

[Instructions for the Long-Stochastic treatment, translated from Italian. the parts that are different in the Long-Deterministic treatment are reported in *italics*.]

Welcome! This is a study about how people take economic decisions. This study is funded by the University of Bologna and other institutions. If you pay attention, the instructions will help you to make your decisions and earn a reasonable amount of money. The earnings will be calculated in points and then converted into euros.

For every 150 points you will receive 1 euro.

In addition, you will receive 3 euros for participation. Your earnings will be paid in cash at the end of today's session.

We ask that you turn off your phone now and do not communicate in any way with the people present in the room until the end of the study. If you have any questions, please raise your hand and we will assist you in private.

This study comprises **23 periods**. In each period you will be paired with another person selected at random from those present in the room.

In every period you will be able to **repeatedly choose between a "GREEN" action and an "ORANGE" action**. Also the person matched with you will be able to repeatedly choose between "green" and "orange" actions. As a consequence, there are four possible combinations: GREEN-green, ORANGE-orange, GREEN-orange, and ORANGE-green. For each combination of actions there is a corresponding cell in Figure B.1 below.

In each cell you can see the gains or losses during the period according to your action and the action of the other. Your action will determine the table row, while the action of the person matched with you will determine the table column.

The earnings described in Figure B.1 above represent earnings per second.

LA TUA AZIONE	l'azione dell'altro	
	verde	arancio
VERDE	guadagni al secondo TU: 1 l'altro: 1	guadagni al secondo TU: -2 l'altro: 2
ARANCIO	guadagni al secondo TU: 2 l'altro: -2	guadagni al secondo TU: 0 l'altro: 0

Figure B.1: Earnings table

For instance, suppose you choose "GREEN" and holds that choice over time: if the other chooses "green" and holds his choice in time, you earn 1 point per second and the other earns 1 point per second; if the other chooses "orange" and holds it, you lose 2 points per second and the other earns 2 points per second. And so on.

In each period, earnings depend on how much time you spend in each cell of Figure B.1. The more time you spend in a cell, the more your average earnings will approximate what is indicated in the cell. For instance, if you spend half of the period in the GREEN-green cell where you earn 1 and half of periods in the ORANGE-orange cell where you earn 0, your earnings will be 0.5 points per second. Are there any questions about how to read the table?

Who is the other person matched with me?

It could be anyone in this room. Your identity and hers will be kept confidential. Also payments will be made in private. There will be 23 periods. At the beginning of each period pairs will be changed. People will be recombined so that **you will never meet the same person twice**.

What should I do? In every period you choose an initial action and then you can decide every instant whether to keep or change that action. The person matched with you can do the same. During a period, both you and the other will be able to change action as many times as you like. Time

flows through very fast ticks (16-hundredths of a second each), in practice there are between six and seven ticks per second, so if you want you can change the action six or seven times per second.




Figure B.2: Earnings table

Earnings

During the period you will receive information in real time on your earnings. In the screen pictured in Figure B.2 above, your cumulated earnings will appear in a graph as a line that will form at every tick of 16-hundredths of a second. In each period you will have an **initial endowment of 50 points** as cumulated earnings. **If during the period, your earnings are zero, then the line will be flat. In case of losses, then the line will be declining. In the case of positive earnings, then the line is increasing.** For instance, if you earn 1 point per second there will be an increasing line that is parallel to the graph grid. If you earn 2 points per second, the line will be increasing, but steeper. Looking at the earnings graph will gives you information on the current action of the other person matched with you. Are there any questions?

To understand how to read the screen, **we will do a trial period**, without consequences on your earnings. For simplicity, the trial period will last 60 seconds and the other will be played by a robot. The robot will start with an action and then, halfway through the period, will change action. Now

please look at the screen and follow the exact guidelines you are given. To start with, choose the 'initial action. Press the screen with your finger on the button that you will be told to choose ("GREEN" or "ORANGE"). Also the robot will choose its initial action ("green" or "orange"). **Please everyone choose "GREEN" now** as the initial action. The selected action will be highlighted in yellow on the table. The period will begin when everybody has chosen their initial action and pressed "OK". From this moment on, the time will begin to run. Then you will see that the graph line is green like your action. Now, please press your finger on the button "OK" to confirm. Does anyone need help? After 10 seconds please everyone press the button "ORANGE." You will see that your action has changed because in the table the line highlighted in yellow will change and that indicates your current action. Moreover the graph line will now be orange in color. After 30 seconds please everyone press again the button "GREEN." Now we ask you to guess what actions did the robot choose. Are there any questions?

We will do two more trial periods, without consequences on your earnings. For simplicity, the trial period will last 60 seconds and the other will be played by a robot. The robot will start with an action and then, halfway through the period, will change action. Now look at your screen. Choose the initial action that you prefer. When everyone has completed, you'll see the time running. You are free to change the action at any time. At the end of the period, we will ask you to guess what actions did the robot choose.

Now we will do the last trial period. Go ahead and choose the action you want. Are there any questions?

For simplicity in the trial periods and the other was a robot and the duration of 60 seconds. However, in the coming periods, the other will be a person in this room *while the duration of each period will be variable and determined randomly. Each period will stop without notice and for everybody at the same moment, and the period duration could vary from less than a second to several minutes.*

How is a period duration established?

The period may stop at every tick of 16 hundredths of a second. This event depends on the result of a random draw. Imagine a box with 10,000 balls, of which 9,973 black and 27 white. It is as if, after every tick, a ball was drawn. If the ball drawn is white, the period ends. If the ball is black, the period continues and the ball is placed back into the box. At the next tick, a new ball is drawn at random. You have to imagine very rapid draws, that is one every tick of 16 hundredths of a second. We calculated that as a result of this, the periods will have an average duration of 60 seconds. There may be some short periods and some long periods. Are there any questions about this?

[DETERMINISTIC: The length of each period will be 60 seconds.]

Very well, then we can start.

References

- ANDREONI, J., AND J. H. MILLER (1993): “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence,” *The Economic Journal*, 103(418), 570–585.
- BENOIT, J.-P., AND V. KRISHNA (1985): “Finitely Repeated Games,” *Econometrica*, 53(4), 905–922.
- BEREBY-MEYER, Y., AND A. E. ROTH (2006): “The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation,” *The American Economic Review*, 96(4), 1029–1042.
- BERGIN, J., AND W. B. MACLEOD (1993): “Continuous Time Repeated Games,” *International Economic Review*, 34(1), 21–37.
- CHARNESS, G., D. FRIEDMAN, AND R. OPREA (2011): “Continuous Time and Communication in a Public-goods Experiment,” mimeo.
- COOPER, R., D. V. DEJONG, R. FORSYTHE, AND T. W. ROSS (1996): “Cooperation without Reputation: Experimental Evidence from Prisoner’s Dilemma Games,” *Games and Economic Behavior*, 12(2), 187 – 218.
- DAL BÓ, P. (2005): “Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games,” *American Economic Review*, 95(5), 1591–1604.
- DAL BÓ, P., AND G. FRÉCHETTE (forthcoming): “The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence,” *American Economic Review*.
- ENGLE-WARNICK, J., AND R. L. SLONIM (2004): “The Evolution of Strategies in a Repeated Trust Game,” *Journal of Economic Behavior & Organization*, 55(4), 553 – 573, Trust and Trustworthiness.
- FISCHBACHER, U. (2007): “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 10(2), 171–178.

- FRIEDMAN, D., AND R. OPREA (forthcoming): “A Continuous Dilemma,” *The American Economic Review*.
- GREINER, B. (2004): “The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics,” University of Cologne, Working Paper Series in Economics 10.
- HAUK, E., AND R. NAGEL (2001): “Choice of Partners in Multiple Two-Person Prisoner’s Dilemma Games: An Experimental Study,” *The Journal of Conflict Resolution*, 45(6), 770–793.
- KREPS, D. M., P. MILGROM, J. ROBERTS, AND R. WILSON (1982): “Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma,” *Journal of Economic Theory*, 27(2), 245–252.
- MAILATH, G. J., AND L. SAMUELSON (2006): *Repeated Games and Reputations: Long-run Relationships*. Oxford University Press.
- MURNIGHAN, J. K., AND A. E. ROTH (1983): “Expecting Continued Play in Prisoner’s Dilemma Games: A Test of Several Models,” *The Journal of Conflict Resolution*, 27(2), 279–300.
- NORMANN, H.-T., AND B. WALLACE (2011): “The Impact of the Termination Rule on Cooperation in a Prisoner’s Dilemma Experiment,” DICE Discussion Papers, 19.
- PALFREY, T. R., AND H. ROSENTHAL (1994): “Repeated Play, Cooperation and Coordination: An Experimental Study,” *The Review of Economic Studies*, 61(3), 545–565.
- RADNER, R. (1986): “Can Bounded Rationality Resolve at the Prisoner’s Dilemma?,” in *Contributions to Mathematical Economics*, ed. by A. Mas-Colell, and W. Hildenbrand, pp. 387–399. North-Holland, Amsterdam.
- RAPOPORT, A., AND A. CHAMMAH (1965): *Prisoner’s Dilemma: a Study in Conflict and Cooperation*. Ann Arbor, University of Michigan Press.
- ROTH, A. E., AND J. K. MURNIGHAN (1978): “Equilibrium Behavior and Repeated Play of the Prisoner’s Dilemma,” *Journal of Mathematical Psychology*, 17(2), 189 – 198.

- SANNIKOV, Y., AND A. SKRZYPACZ (2007): “Impossibility of Collusion under Imperfect Monitoring with Flexible Production,” *American Economic Review*, 97(5), 1794–1823.
- SELTEN, R., M. MITZKEWITZ, AND G. R. UHLICH (1997): “Duopoly Strategies Programmed by Experienced Players,” *Econometrica*, 65(3), 517–556.
- SELTEN, R., AND R. STOECKER (1986): “End behavior in Sequences of Finite Prisoner’s Dilemma Supergames: A Learning Theory Approach,” *Journal of Economic Behavior & Organization*, 7(1), 47 – 70.
- SIMON, L. K., AND M. B. STINCHCOMBE (1989): “Extensive Form Games in Continuous Time: Pure Strategies,” *Econometrica*, 57(5), 1171–1214.
- SZYMANSKI, S. (1996): “Making Hay While the Sun Shines,” *The Journal of Industrial Economics*, 44(1), 1–16.



Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

Strada Maggiore 45
40125 Bologna - Italy
Tel. +39 051 2092604
Fax +39 051 2092664
<http://www.dse.unibo.it>