Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence *

Matthias Blonski[†] Peter Ockenfels[‡] Giancarlo Spagnolo[§]

version August 12, 2009

Abstract

We propose an axiomatic approach for equilibrium selection in the discounted, infinitely repeated symmetric Prisoner's Dilemma. Our axioms characterize a unique selection criterion S^* predicting whether players cooperate or not. S^* depends on all model primitives — in particular on the "sucker's payoff" of a cooperating player meeting a defecting one. Selection criterion S^* results in a critical discount factor δ^* strictly larger than $\underline{\delta}$, the classic criterion often used in applications. In an experimental test we find a striking predictive power of our proposed criterion S^* regarding cooperation. This gets even more pronounced if we restrict attention to equilibrium outcome paths — i.e. if we filter out disequilibrium behavior. When the two critical discount factors move in different directions, observed cooperation follows predictions based on δ^* falsifying those based on $\underline{\delta}$.

JEL Classification: C72, C73, C92
Keywords: cooperation, repeated Prisoner's Dilemma, equilibrium selection, axiomatic approach, experiments, strategic risk, risk dominance, sucker's payoff, collusion, coordination

---

[*] For valuable comments, we are grateful to various seminar participants in Karlsruhe, Jena, at the ESA conference in Rome, Games 2008 and "Sozialwissenschaftlicher Ausschuss des Vereins für Socialpolitik". All errors remain our responsibility. We thank the E-finance Lab at the J.W.Goethe University in Frankfurt for funding the experiment.

[†] Department of Economics, J.W.Goethe University, Frankfurt am Main, 60054 Frankfurt a.M., Germany; e-Mail: *blonski@wiwi.uni-frankfurt.de*

[‡] Department of Finance, J.W.Goethe University, Frankfurt am Main, 60054 Frankfurt a.M., Germany; e-Mail: *peter@ockenfels.com*

[§] University of Rome 'Tor Vergata', SITE - Stockholm School of Economics, and CEPR. e-Mail: *Giancarlo.Spagnolo@uniroma2.it, Giancarlo.Spagnolo@hhs.se.*

1 Introduction

Under which conditions do people cooperate? When should we expect them to successfully coordinate on some cooperative equilibrium or remain stuck in the non-cooperative one? The question of equilibrium selection in infinitely repeated games with complete information is one of the great unsolved puzzles facing economic theory, but it is also crucial to sociology, political science and biology.

This article is a step towards answering this fundamental question. By restricting attention to the symmetric Prisoner's Dilemma we propose an axiomatic approach to equilibrium selection. This approach formulates a minimal set of simple and intuitive conditions on the model primitives any sensible selection theory should satisfy. It yields a unique, simple criterion predicting whether players cooperate or not. Then we present results from a laboratory experiment testing the proposed criterion and showing that it has an unusually strong predictive power, particularly if compared to the criterion most commonly used in applications.

Where do we stand? The theory of infinitely repeated games has already offered many important answers to our opening question. Limiting attention to complete information, it has shown that cooperation is easier when gains from cooperation are relatively large, when short-run gains from 'cheating' are relatively small, when the severeness of expected punishment for cheating is larger, and when people are more patient or interactions more frequent.[1] Recent experimental work appears to support most of these answers.[2]

These observations can be summarized as follows: cooperation is more likely under conditions where players' non-deviation constraints are easy to satisfy. In a symmetric setting, such conditions are identical for all players and their tightness can be quantified by $\delta \geq \underline{\delta}$, where $\underline{\delta}$ is the minimum discount factor for which cooperation is sustainable in equilibrium. Applied theory has very much relied on changes in $\underline{\delta}$ when trying to design institutions or identify real world situations that are more or less conducive to cooperation, both in macroeconomic (e.g. Kocherlakota (1996), or Ligon et al. (2002)) and in microeconomic applications (e.g. Gilo et al. (2006), Athey et al. (2006)).[3] This traditional approach considers cooperation more likely when $\underline{\delta}$ falls and less likely when it goes up. Although the applied literature rarely reflects about theoretical foundations for this interpretation of $\underline{\delta}$, it can be justified by cooperation being Pareto-dominant relative to non-cooperation and therefore being selected when feasible.[4] From here, therefore, we simply call $\underline{\delta}$ the *standard criterion*.

---

[1] Friedman (1971) was the first to formally highlight many of these points. See e.g. Mailath and Samuelson (2006) and Fudenberg and Tirole (1991, Ch. 5) for excellent surveys.

[2] See e.g. Dal Bo (2005).

[3] To name just a few other classic applications, see Tirole (1988, ch. 6.3.2.1), Bernheim and Whinston (1990), and Motta (2004, ch. 4.2.5).

[4] The most widely accepted criterion is efficiency or Pareto-dominance since it is motivated by a normative perspective and it is the relevant criterion to describe the boundary of the equilibrium payoff

Our theoretical contribution There are many possibilities to motivate selection criteria. For example one could formulate an evolutionary model, introduce various kinds of robustness regarding information and mistakes, define the basin of attraction and stability in dynamic models, perform simulations as Axelrod (1984) and so forth.[5] Although we are sympathetic to most of these methods we favor even more a selection theory that is independent of the modeler's taste with respect to theories. Our way to do this is to formulate and motivate three minimal axioms any sensible theoretic selection model should satisfy. Further, we propose and motivate two additional axioms that are sufficient to end up in a sharp prediction regarding equilibrium selection in the discounted repeated Prisoner's Dilemma. This prediction from here is called the *alternative selection criterion*. Both sets of axioms — the minimal set and the full set of axioms — result in testable predictions later denoted as hypotheses.

The first two axioms are standard in the literature on equilibrium selection. The second of them regards the discount rate and is thereby more specific to repeated games. The crucial novel axiom 3 reflects the intuitive idea that cooperation gets more and more risky if the sucker's payoff — earned by cooperating when the opponent defects — gets smaller and smaller. Once it converges to minus infinity any cooperation attempt gets deadly dangerous, hence real world players would never cooperate. Conversely, this kind of *strategic risk* for cooperation vanishes if the sucker's payoff approaches the defection equilibrium payoff, in which case the standard non-deviation constraints remain the only concern. The remaining two axioms — resulting in the critical lower bound on discount rates $\delta^* > \underline{\delta}$ below which cooperation is unsustainable — can be seen as a natural, continuous and monotonous extension of the two principles behind axiom 3. In contrast to the standard criterion, this alternative criterion depends on all payoff-parameters, including the sucker's payoff.

While various related criteria have already been discussed or tested in the early literature[6] we are not aware of theoretical foundations that single out this or another criterion as the relevant cooperation predictor.[7]

---

space. The risk-dominance criterion introduced by Harsanyi and Selten (1988) is not well defined for infinite games. In finite games it is based on the so called bicentric prior and the tracing procedure. In games bigger than 2x2-games Harsanyi and Selten's concept has not been applied often since it is mathematically involved. See Blonski and Spagnolo (2004) for a more detailed discussion.

[5] Compared to Axelrod (1984) we are far less ambitious in the sense that we restrict our attention to the question if people cooperate rather than how cooperation comes about.

[6] See for example Rapoport and Chammah (1965) for the finitely repeated Prisoner's Dilemma or Murnighan and Roth (1978, 1983) for the infinitely repeated Prisoner's Dilemma if a player plays against a given strategy.

[7] An exception is Blonski and Spagnolo (2004) who build on Harsanyi and Selten's (1988) concept of risk dominance and link it to infinitely repeated games. Their theory identifies this same criterion and thereby implicitly connects our axiomatic approach also to Harsanyi and Selten's risk dominance criterion.

Our experimental contribution We test our criterion against the standard criterion with a laboratory experiment that simulates infinitely repeated games with random continuation and matching rules and many experimental subjects. Our experimental setup is closest to Dal Bo (2005), though we design it for testing our theory, changing parameters so that $\underline{\delta}$ and δ^* may change in different directions when comparing couples of treatments.[8]

We find that in all cases in which $\underline{\delta}$ and δ^* change in different directions our criterion based on changes of δ^* predicts correctly, while the standard criterion based on changes on $\underline{\delta}$ fails. This result is very robust since all equilibrium selection theories that satisfy our first 3 axioms yield this same comparative statics prediction. The observed differences in cooperation frequencies are very large and significant at any confidence level. The standard criterion only maintains some predictive power in situations in which δ^* remains constant, hence as a 'residual' of our criterion.

Our second hypothesis derived from our alternative equilibrium selection criterion predicts under which conditions players cooperate in equilibrium. We compare cases where subjects' discount factor – i.e. continuation probability – increases from $\delta < \delta^*$ to $\delta > \delta^*$. We verify in our experiments whether this raises cooperation frequencies more than when it moves from $\delta < \underline{\delta}$ to $\delta > \underline{\delta}$. We find robust support also for this additional hypothesis.This latter observation is also consistent with recent independent experimental evidence in Dal Bó and Frechette (2008).[9]

We first test our hypotheses against the full data set. In a second step — to further improve the match between theoretical hypotheses and experimental test — we ask the question to which extent the alternative selection criterion predicts the mode of behavior among those observations that are consistent with equilibrium behavior. To do this we identify all the observed paths of play that are possible outcome paths of some equilibrium and those that cannot be the outcome paths of any equilibrium. The preciseness of the predictions of the alternative selection criterion further increases when we use this filtered data set. We were surprised to see that more than 85% of our observed outcome paths — i.e. experimental subjects' actual behavior — are indeed equilibrium outcome paths of the repeated game which explains the predictive power of our criterion for the whole original data set. Note that our hypotheses do not explain to which extent players behavior is actually consistent with equilibrium behavior. Among those players that did not play equilibrium outcome paths were for example players who cooperated in parameter constellations where cooperation is not an equilibrium. This non-equilibrium

---

[8] We are not aware of any other experimental studies that allow for this comparison. The closest is Dal Bo (2005) who compares one pair of parameter constellations where $\underline{\delta}$ changes while δ^* remains constant. We discuss Dal Bo (2005) in more detail in section 8.

[9] Our experiments were run independently and simultaneously in (2006) and happen to be complementary. Their focus is mainly on learning effects and players' strategies, while our is on equilibrium selection tools and strategic risk. Unfortunately in their experimental treatments payoffs change so that $\underline{\delta}$ and δ^* are constant or change in the same direction and their players meet again the same opponents with positive probability when they are re-matched. Hence, our main result cannot be tested by their data. See again section 8 for a more detailed comparison.

behavior turns out to fade rapidly towards non-cooperation.

To sum up, our axiomatic approach predicts under which conditions players pick a cooperative equilibrium in the repeated Prisoner's Dilemma. Our experimental evidence supports this novel theoretical prediction on a level of preciseness and robustness rarely seen in experiments testing game theory.

Many experimental studies have been undertaken before to investigate the determinants of cooperation and conflict between real world subjects in Prisoner's Dilemma situations. Restricting focus to experiments on infinitely repeated games with complete information, which start with the pioneering work of Murnighan and Roth (1978, 1983) and include the recent work of Dal Bó (2005) and Duffy and Ochs (2009) among many others (see Section 8), these studies usually find that real players rarely cooperate for $\delta < \underline{\delta}$, i.e. if indefinite cooperation is not sustainable in equilibrium. This evidence suggests a reassuring degree of rationality in the sense that equilibrium non-deviation conditions indeed are robust necessary cooperation conditions for most real world players. However, there are always many parameter constellations for which cooperation is supported as an equilibrium but in which nevertheless real players rarely cooperate (see e.g. the conclusions in Dal Bó 2005). In the last section of the paper we reassess the previous experimental evidence in the light of our selection theory, showing that much of their unexplained variation can be captured by this theory. Given the methodological differences of some of the earlier studies we were surprised to see how much of unexplained variation can be captured by this alternative selection theory.

The rest of the paper is organized as follows. The following Section 2 presents a simple example designed to introduce the basic idea behind the crucial axiom 3. In section 3 we formulate and discuss simple and intuitive axioms guiding players' behavior if there are many equilibria that any theoretical approach should arguably satisfy. In Section 4 we show how the three more general axioms determine a general selection criterion, and together with two additional axioms the more specific criterion δ^*. Based on these general and specific criteria we formulate in section 5 two theoretical predictions to be tested against the standard criterion. Section 6 describes our experimental design, while section 7 presents the experimental results. Section 8 surveys previous experimental investigations within this setting, compare their evidence with ours where possible, and briefly concludes.

2 Numerical example

Compare the following two PD stage games Γ_1 and Γ_2 given by

Γ_1	C	D
C	2 2	4 0.999
D	0.999 4	1 1

with

Γ_2	c	d
c	2 2	3 -1000
d	-1000 3	1 1

and the corresponding discounted infinitely repeated games denoted by $\Gamma_1(\delta)$ and $\Gamma_2(\delta)$. cooperation is not supportable as equilibrium behavior in $\Gamma_1(\delta)$ if the corresponding payoff $\frac{2}{1-\delta}$ is smaller than the payoff $4-\frac{\delta}{1-\delta}$ from a single deviation followed by indefinite defection of both players. This yields a lower bound $\delta \geq \underline{\delta}(\Gamma_1) = \frac{4-2}{4-1} = \frac{2}{3}$ on discount factors in $\Gamma_1(\delta)$. Correspondingly, this lower bound for $\Gamma_2(\delta)$ is $\underline{\delta}(\Gamma_2) = \frac{1}{2}$.

Those players for whom Pareto-efficiency is the relevant criterion should cooperate in $\Gamma_2(\delta)$ but not in $\Gamma_1(\delta)$ if $\delta \in \left[\frac{1}{2}, \frac{2}{3}\right)$. Put differently, the size of the interval $[\underline{\delta}, 1]$ supporting cooperation as the most efficient equilibrium is larger in $\Gamma_2(\delta)$. However, in this example we intuitively expect, that even for very patient players it is far more difficult to cooperate in $\Gamma_2(\delta)$ than in $\Gamma_1(\delta)$. For example, our intuitive prediction in this example is that real players with a discount rate of, say, $\delta = .99$ would almost certainly be able to build up cooperation in $\Gamma_1(.99)$ whereas they most likely would never dare to cooperate in $\Gamma_2(.99)$ whereas the efficiency criterion would predict cooperation in both cases.

Why does the efficiency criterion with its interpretation of $\underline{\delta}$ fail to match this intuitive prediction?

This example suggests that there may be a simple reason for this failure. For the discount rate $\delta = .99$, for example, the long run incentives to cooperate instead of defecting forever are identical for both games $\Gamma_1(\delta)$ and $\Gamma_2(\delta)$ and are given by $\delta\left(\frac{2}{1-\delta} - \frac{1}{1-\delta}\right) = 99$. This is what is at stake in the future in both games — sometimes called "the shadow of the future". We expect, however, that players' short run incentives depend on two effects, (1) the immediate gains from defecting reflected in the standard non-deviation constraint, but also (2) the loss faced by a player who plans to cooperate but meets an opponent who defects. This latter disincentive compares the one-shot payoff of a cooperating player suffering defection by his opponent — from here often called the "sucker's payoff" — with the one-shot payoff from playing defect instead. The first effect corresponds to the left hand side of the standard non-deviation condition of a player considering a cooperative equilibrium, reflected in $\underline{\delta}$. In our example it is given by $4-2=2$ in $\Gamma_1(\delta)$ and $3-2=1$ in $\Gamma_2(\delta)$. This disincentive to cooperate is indeed larger in $\Gamma_1(\delta)$. The second effect, however, is $1 - .999 = .001$, hence small in $\Gamma_1(\delta)$ but overwhelming $1+1000 = 1001$ in $\Gamma_2(\delta)$. Adding up both effects and comparing them with the long run incentive to cooperate yields a total incentive $99 - 2.001 > 0$ in favor of cooperating in $\Gamma_1(\delta)$ and

$99 - 1002 \ll 0$ in favor of defecting in $\Gamma_2(\delta)$.

Clearly, this example is specified such that the whole emphasis is on the second effect which has been neglected by the literature so far. The trade off between these two effects is similar to that faced by players in one shot coordination games where the choice is between a sure and a risky action. Previous experimental studies have shown that efficient but risky equilibria are often not chosen if gains are small or coordination requirements high (see van Huyck et al. (1990) and Heinemann et al (2009)).

In the remainder of this article we show that a criterion that takes into account both incentives outperforms the standard criterion $\underline{\delta}$ in predicting cooperation in repeated interactions with the strategic features of a repeated Prisoner's Dilemma.

3 Axiomatic Approach

Model primitives. Consider the symmetric Prisoner's Dilemma stage game Γ given by

Γ	C	D
C	c c	b a
D	a b	d d

characterized by payoff parameters a, b, c, d with $b > c > d > a$ and $2c > b + a$.[10] Call $\Gamma(a, b, c, d, \delta)$ the respective infinitely repeated game with common discount factor δ. The primitives of the model are given by the parameter set

$$Q = \{(a, b, c, d, \delta) | b > c > d > a, 2c > b + a, 0 < \delta < 1\} \subset \mathbb{R}^5.$$

In this section $\subset$ always means strict subset and we use $\subseteq$ for weak subset.

To embed our model into equilibrium selection theory and the language of Harsanyi and Selten we introduce the following notation. Call any equilibrium of $\Gamma(\delta)$ a *D-equilibrium* if its outcome path is $((D, D), (D, D), \ldots)$ — i.e. contains only defective actions. Conversely, any other equilibrium with at least one cooperative action on its outcome path is called *C-equilibrium.* The set E of all equilibria is then a disjoint union $E = E_C \cup E_D$ of D-equilibria and C-equilibria. We define an *equilibrium selection function* $G : Q \to E$ to be *cooperation-equivalent* to another equilibrium selection function $G' : Q \to E$ iff they predict C-equilibria — and thereby D-equilibria — for the same subset of the parameter space

$$G(a, b, c, d, \delta) \in E_C \Leftrightarrow G'(a, b, c, d, \delta) \in E_C$$

---

[10] Later, we shall normalize parameters as $c \equiv 1$ and $d \equiv 0$ such that the relevant incentives only depend on 2 remaining parameters $a < 0 < 1 < b$. However, since this imposes an assumption on preferences we postpone this till we formulate our axioms.

or $G^{-1}(E_C) = G'^{-1}(E_C) \subseteq Q$. In this article we study the question whether cooperation is possible or not. In terms of equilibrium selection functions this means that we are interested in *cooperation-equivalence classes* of equilibrium selection functions $G : Q \to E$. Any such class is uniquely determined by that subset of parameter values $G^{-1}(E_C) \subseteq Q$ for which this class of equilibrium selection functions predicts that players cooperate — i.e. do not defect forever. Hence, from here a cooperation-equivalence class of equilibrium selection functions or shorter a *selection criterion for cooperation* is simply a subset $S \subseteq Q$ of parameter values for which any equilibrium selection function in the according equivalence class predicts that players play a C-equilibrium. Note that any selection criterion for cooperation S must be contained in the subset of parameter values for which there exist not only D-equilibria but also C-equilibria defined by

$$\left\{(a, b, c, d, \delta) \in Q \,\middle|\, \delta \geq \underline{\delta} := \frac{b-c}{b-d}\right\}.$$

The interesting question is therefore for which part of this set players indeed do cooperate, i.e. select C-equilibria. We consider first the two extreme cases for cooperation criteria. The most "cooperation friendly" criterion is denoted by S_C and is defined by always selecting a C-equilibrium when there is one, i.e. $S_C := \{(a, b, c, d, \delta) \in Q \,|\, \delta \geq \underline{\delta}\}$. Conversely, in the least cooperation friendly criterion denoted by S_D players always defect, hence $S_D = \emptyset$. Any selection criterion for cooperation S must be "between" S_D and S_C, i.e. $S_D \subseteq S \subseteq S_C$.

Our axiomatic approach addresses now the following question. Which properties should a sensible selection criterion for cooperation S satisfy? Later we will also be interested in formulating robust implications of these properties that can be tested in the lab and can verify or falsify our properties.

Cooperation Axioms. Let S be a selection criterion for cooperation with $S_D \subseteq S \subseteq S_C$.

Axiom 1 (Positive Linear Payoff Transformation Invariance) *Let $\tau : \mathbb{R} \to \mathbb{R}$ be a positive linear payoff transformation with $\tau(x) = \alpha x + \beta$ where $\alpha > 0$. Then*

$$(a, b, c, d, \delta) \in S \Leftrightarrow (\tau(a), \tau(b), \tau(c), \tau(d), \delta) \in S.$$

Axiom 1 is well known in equilibrium selection theory. It corresponds to Harsanyi and Selten's *invariance with respect to isomorphisms.*[11] The interpretation of this axiom is that players' payoffs are cardinal — for example represent their von Neumann-Morgenstern utility functions — and abstract away from all framing effects like the choice of the 0-level or any scaling. Note that axiom 1 reduces the parameter set by two degrees

---

[11] Harsanyi and Selten's *isomorphisms* also allow for permutation of players' names which do not matter in symmetric games as ours.

of freedom. Since players do not care about the 0-level and scaling one can always specify $d = 0, c = 1$ since

$$(a, b, c, d, \delta) \in S \Leftrightarrow \left(-\frac{d-a}{c-d}, \frac{b-d}{c-d}, 1, 0, \delta\right) \in S.$$

Since from here we restrict attention to selection criteria S that satisfy axiom 1 we can simplify the notation to

$$S \subseteq Q = \{(a, b, \delta) | b > 1 > 0 > a, b + a < 2, 0 < \delta < 1\} \subset \mathbb{R}^3.$$

Axiom 2 ($\delta-$Monotonicity) *For any payoff parameters* $(a, b) \in \{(a, b) | b > 1 > 0 > a, b + a < 2\}$ *there exists a critical* $\delta_S(a, b)$ *such that*

$$\delta \geq \delta_S(a, b) \Leftrightarrow (a, b, \delta) \in S.$$

While axiom 2 reflects the same logic as Harsanyi and Selten's payoff monotonicity it does not turn up in their setup.[12] Note that axiom 2 implies the weaker condition

$$\delta > \delta', (a, b, \delta') \in S \Rightarrow (a, b, \delta) \in S$$

but not vice versa. The only difference is that the latter condition does not specify a tie-breaking rule, i.e. whether players cooperate once $\delta = \delta_S(a, b)$. We regard this as a purely technical matter. Beyond the latter $\delta-$monotonicity condition axiom 2 simply picks one of two simple tie-breaking rules, namely, always cooperative behavior for $\delta = \delta_S(a, b)$.

The following axiom is related to strategic risk. It formulates the least restrictive condition representing the intuition provided in our introductory example.

Axiom 3 (Boundary Conditions)

i. *Lower Boundary Condition: If the sucker's payoff gets extremely low S selects D-equilibria, i.e. players will defect. Formally,*

$$\forall (a, b, \delta) \in S\ \exists a' < a \text{ such that } (a', b, \delta) \notin S \Leftrightarrow \delta_S(a', b) > \delta.$$

ii. *Upper Boundary Condition: If there are C-equilibria then they are selected by S if the sucker's payoff is high enough. Formally,*

$$(a, b, \delta) \in S_C \Rightarrow \exists a' \in (a, 0) \text{ such that } (a', b, \delta) \in S \Leftrightarrow \delta_S(a', b) \leq \delta.$$

---

[12] The reason is that Harsanyi and Selten provide an axiomatic characterization only for what they call *risk dominance between strong equilibrium points* $2 \times 2-$*games*. Neither does their more general tracing procedure definition of risk dominance cover discounted repeated games since they are infinite.

Some intuition for the boundary condition axioms 3.i and 3.ii was already discussed in our introductory example in section 2. Axiom 3.i requires that players refrain from cooperative actions if *strategic risk* gets overwhelming while conversely 3.ii makes sure that for payoff parameters where C-equilibria exist players indeed cooperate as strategic risk converges to 0. It is instructive to note that axioms 3.i and 3.ii imply that equilibrium selection with respect to cooperation is "strictly between" the two extreme cooperation criteria S_D and S_C. Formally this is

$$S_D \subset S \subset S_C$$

with strict subsets or $\frac{b-1}{b} < \delta_S(a,b) < 1$. In particular, this means that S_C does not satisfy axiom 3.i while S_D does not satisfy axiom 3.ii.

Axiom 4 (Incentive Independence) *The three incentives*

1. *long run cooperation incentive* $\frac{\delta}{1-\delta}$*,*
2. *C-equilibrium non-deviation incentive* $b-1$*, and*
3. *the strategic risk incentive* $-a$

are independent*. Formally, there exists an additively separable function*

$$\gamma\left(\frac{\delta}{1-\delta}, b-1, -a\right) = \frac{\delta}{1-\delta} - \gamma_1(b-1) - \gamma_2(-a)$$

with

$$(a,b,\delta) \in S \Leftrightarrow \gamma\left(\frac{\delta}{1-\delta}, b-1, -a\right) \geq 0.$$

Axiom 4 is more specific and calls for a more detailed discussion. This axiom reflects a more structured view of a repeated game in the sense that there is a trade-off between long-run and short-run incentives.[13] Further, the *incentive independence axiom* 4 is only a small step beyond the linear payoff transformation invariance axiom 1 together with the $\delta-$Monotonicity axiom 2. Why? Axiom 1 already imposes that players only care about payoff differentials $b-c, c-d$ and $d-a$ rather than absolute payoffs a,b,c,d while axiom 2 together with axiom 1 implies that there is a critical discount rate $\delta_S(a,b,c,d)$ above which S selects players to cooperate and below which players defect and that it is a function of the payoff differentials $b-c, c-d$ and $d-a$. Hence, the additional requirement of the incentive independence axiom 4 beyond axioms 1 and 2 is that payoff differential changes — or in our language incentive changes — do not affect each other. In other words, the relative weight between the two short run incentives is not affected by variations of the long run incentive or the relative weight between the long run incentive

---

[13] For example Mailath and Samuelson (2006) build their motivation for repeated games exactly on this perspective, the differences between short run and long run incentives.

and any of the short run incentives is not affected by variations of the the remaining short run incentive. Neither should any two of these 3 incentives reinforce or weaken each other. Finally, it is interesting to note that both extreme cooperation criteria S_C and S_D do satisfy the incentive independence axiom 4.

Axiom 5 (Equal Weight) *The two short run incentives, the C-equilibrium non-deviation incentive* $b-1$ *and the strategic risk incentive* $-a$ *carry* equal weight. *Formally, for a function*

$$\gamma\left(\frac{\delta}{1-\delta}, b-1, -a\right) = \frac{\delta}{1-\delta} - \gamma_1(b-1) - \gamma_2(-a)$$

with

$$(a, b, \delta) \in S \Leftrightarrow \gamma\left(\frac{\delta}{1-\delta}, b-1, -a\right) \geq 0$$

, it holds that

$$\gamma(x, y, z) = \gamma(x, z, y).$$

Axiom 5 builds on the more structured view of the incentive independence axiom 4 and compares only the two short run incentives. In principle there could be any relative weighting of these two payoff differentials. Which of the two short run incentives is more relevant depends on the action chosen by the other player. Without prior knowledge beyond the primitives of the model any other than the equal weighting rule would impose an exogenous asymmetry and according arbitrariness to the problem in question — i.e. the equilibrium selection problem. Therefore, by a similar logic as Harsanyi and Selten's motivation of their *bicentric prior* we also invoke here the *Laplace principle of insufficient reason* stating that without any further knowledge beyond the primitives of the model players should weight the anticipated action of the other player and thereby the relevant short term incentives symmetrically.

Remember from our introduction that S_C can be regarded as the *standard criterion* in the sense that most of the applied literature implicitly or explicitly uses it. Therefore at this point it is interesting to observe that S_C violates the boundary conditions axiom 3 and the equal weighting axiom 5 but satisfies axioms 1, 2 and 4.

4 Theoretical Results

Our first result builds only on the more fundamental axioms. By imposing very little structure it is a very robust result. It generalizes our introductory example by comparing any selection criterion for co-operation $\tilde{S}$ that satisfies axioms 1, 2, and 3 with the classical selection criterion for co-operation S_C and shows that the boundary condition axiom 3 leads to opposing comparative static properties of $\tilde{S}$ and S_C.

Proposition 1 *Let* $\tilde{S}$ *be a selection criterion for co-operation in the discounted infinitely repeated Prisoner's Dilemma that satisfies axioms 1, 2, and 3. Then there exist pairs of*

Prisoner's Dilemma stage games specified by payoff parameters $(a,b),(a',b') \in \{(a,b)|b > 1 > 0 > a, b+a < 2\}$ *such that the comparative statics of selection criterion for co-operation* $\tilde{S}$ *and of the classic criterion* S_C *point in opposite directions. Formally,* $\delta_{\tilde{S}}(a,b) > \delta_{\tilde{S}}(a',b')$*, but* $\delta_{S_C}(a,b) < \delta_{S_C}(a',b')$*.* □

PROOF Pick payoff parameters $1 < b < b'$ such that $\frac{b-1}{b} < \frac{b'-1}{b'}$. This implies the second condition $\delta_{S_C}(a,b) < \delta_{S_C}(a',b')$. Now there are two further degrees of freedom given by the sucker's payoff parameters a, a'. Axiom 3.i makes sure that $\forall\delta$ there exists $a < 0$ such that $(a,b,\delta) \notin \tilde{S}$. Pick such parameters a, δ. By definition of $\delta_{\tilde{S}}(a,b)$ this implies $\delta_{\tilde{S}}(a,b) > \delta$. Conversely axiom 3.ii guarantees that for any δ, in particular for the same δ as before there exists $a' < 0$ such that $(a',b',\delta) \in \tilde{S}$ or $\delta_{\tilde{S}}(a,b) < \delta$. Together this implies the first condition $\delta_{\tilde{S}}(a,b) > \delta_{\tilde{S}}(a',b')$. ■

Our second result characterizes a unique criterion S^* that satisfies all axioms 1 through 5.

Proposition 2 *There is a unique selection criterion for co-operation* S^* *that satisfies axioms 1 through 5 with*

$$\delta_{S^*}(a,b) = \delta^* := \frac{b-a-1}{b-a}.$$

PROOF Incentive independence axiom 4 yields

$$\gamma \geq 0 \Leftrightarrow \delta \geq \frac{\gamma_1(b-1)+\gamma_2(-a)}{1+\gamma_1(b-1)+\gamma_2(-a)}.$$

Axiom 3.ii implies for $a \to 0$

$$\frac{\gamma_1(b-1)+\gamma_2(0)}{1+\gamma_1(b-1)+\gamma_2(0)} = \frac{b-1}{b}$$

or $\gamma_1(b-1)+\gamma_2(0) = b-1$. Further, axiom 3.i yields for $-a \to \infty$

$$\frac{\gamma_1(b-1)+\gamma_2(-a)}{1+\gamma_1(b-1)+\gamma_2(-a)} \to 1$$

from the left since $\frac{\gamma_1(b-1)+\gamma_2(-a)}{1+\gamma_1(b-1)+\gamma_2(-a)} \in [0,1)$. This implies that $\gamma_1(b-1)+\gamma_2(-a) \to \infty$ for $-a \to \infty$. Finally, axiom 5 together with $\gamma(b-1)+\gamma_2(0) = b-1$ implies $\gamma_1(b-1)+\gamma_2(0) = \gamma_1(0)+\gamma_2(b-1) = b-1$ which implies $\gamma_1(x) = \gamma_2(x) = x$. These functional forms of γ_1, γ_2 yield

$$\delta \geq \frac{\gamma_1(b-1)+\gamma_2(-a)}{1+\gamma_1(b-1)+\gamma_2(-a)} = \frac{b-1-a}{1+b-1-a} = \delta^*.$$

5 Theoretical Predictions

In an experiment on repeated games we can only directly observe outcome paths. Since not all outcome paths are equilibrium outcome paths this provides an elegant tool to disentangle the two issues, i.e. the equilibrium selection problem from the equilibrium versus disequilibrium behavior, at least to some extent. Since an equilibrium selection criterion does not predict anything on players that do not play an equilibrium the most accurate way to relate our experimental data to the equilibrium selection theory is to separate observed outcome paths into equilibrium outcome paths and non-equilibrium outcome paths.[14] The predictive power of any equilibrium selection criterion can then only regard the sub-sample of equilibrium outcome paths.

We expect another bias in our experimental data that stems from the fact that discounting in infinitely repeated discounted games is implemented as a termination probability. Hence, what we observe are *finite* stochastic termination realizations of *infinite* outcome paths. Now, consider an outcome path starting with (D, D) that terminates after the first period. According to our definition of C-equilibria and D-equilibria it is impossible to say whether this outcome path results from a D-equilibrium or a C-equilibrium since there are C-equilibria of “more hesitant cooperators” starting with (D, D) and switching to cooperation later. This means that an unknown proportion of finite paths only containing defective actions result from C-equilibria rather than D-equilibria. Conversely any observed cooperative action rules out that a D-equilibrium was played. Together this implies that the proportion of cooperative behavior relative to defective behavior is underestimated by experimental data. In other words, for any selection criterion S our theory predicts that among all observed equilibrium outcome paths we should expect only defective outcome paths for $\delta < \delta_S(a, b)$ and cooperative and defective outcome paths for $\delta > \delta_S(a, b)$.

As real world observers we still expect “unexplained noise” in our observed data. Subjects may differ in the way they evaluate monetary payoffs, in how they perceive the experimental design, in history and beliefs, and in other unobserved details. From the viewpoint of an experimentator this is much less of a problem since much of this unobserved noise can be filtered out by comparing pairs of experimental treatments with different payoff parameters while keeping all unobserved details and the experimental design constant. In combination with our alternative criteria $\tilde{S}$ this leads us to the first of our main testable hypotheses. It rests on proposition 1 predicting that for a selection criterion for cooperation $\tilde{S}$ that satisfies axioms 1, 2, and 3 there exist pairs of Prisoner’s Dilemma stage games with appropriately chosen payoff parameters such that the comparative statics of selection criterion for cooperation $\tilde{S}$ and of the classic criterion S_C point in opposite directions. S^* is a particular such criterion and the appropriately chosen payoff parameters for pairs of Prisoner’s Dilemma stage games depend on the

---

[14] Since we only observe outcome paths rather than strategies an equilibrium outcome path might still result from non-equilibrium strategies. We disregard, however, this latter impreciseness as neither the players nor the experimentator can observe actual strategies.

selection criterion that is to be tested. Nevertheless, besides testing the comparative properties of selection criteria satisfying axioms 1, 2, and 3 we also want to test the precise numerical prediction at which parameter threshold δ^* cooperation frequencies should change contained in our criterion S^*. Our second hypothesis formulates a prediction that tests for the validity of δ^* by again comparing cooperation frequencies for pairs of parameter constellations. Clearly, we are more cautious in our expectations regarding the validity of the second of the following two hypotheses.

Consider two similarly designed experimental treatments Γ_1 and Γ_2 that only differ in payoff parameters.

Hypothesis (i): If $\delta^*(\Gamma_1) < \delta^*(\Gamma_2)$ but $\underline{\delta}(\Gamma_1) > \underline{\delta}(\Gamma_2)$ thereby predicting opposite changes in cooperation frequencies, more subjects will cooperate in Γ_1. Hence, the change of cooperation frequencies will follow predictions based on δ^*.

Hypothesis (ii): If $\delta < \delta^*(\Gamma_1) < \delta^*(\Gamma_2)$ our second more specific criterion δ^* predicts only little change in cooperation frequencies between the two games while for $\delta^*(\Gamma_1) < \delta < \delta^*(\Gamma_2)$ it predicts a visible rise in cooperation frequency in Γ_1 compared to Γ_2.

6 Experimental Design

Our experiments were conducted in the computer lab of the Economics and Business Department of the University of Frankfurt am Main in May, June and November 2006. They were announced to all students with an email account at the department. Most of the participants were business and economics undergraduates. All sessions were computerized, using a program done with z-tree (Fischbacher, 2007). Students were seated randomly at computer terminals. Instructions were given in written form and were read in public. Eventual questions in turn were answered in private. Before the experiment started all subjects were asked questions on the screen to make sure and to make it common knowledge they were understanding all the important ingredients of the decision model. Only after all subjects passed the test correctly the experiment was started. Throughout the sessions students were not allowed to communicate and could not see each others' screens. After the experiments subjects had to answer a questionnaire and were paid out individually. Continuation probabilities were chosen such that the expected duration of a session was less than 75 minutes and the total payoff of a subject varied between 15 and 25 Euro. Since this is a short time and subjects were paid out after the session we suppose that within the time window of a session subjects do not further discount payoffs.

Simulating an infinitely repeated game. The infinitely repeated PD-game is an idealized model which is impossible to implement literally in an experiment with real subjects, as real subjects are aware to have finite lifes. However, it is well known that the mathematical structure of this game allows for another interpretation as a game with a stochastic break off. More precisely, after every stage (sometimes called round) the game ends with probability $1 - \delta$ and the next round formed by the same stage game continues with probability δ. This interpretation was introduced into experimental research by Roth and Murnighan (1978), and meanwhile it has become mainstream since a large number of studies have followed this route, including this one. In the instructions only the continuation probability was mentioned and it was explained that at every stage the expected number of future stages is given by $\frac{1}{1-\delta}$. Since in every session up to 19 repeated games were played longer and shorter realizations of the same repeated game average out and the event that an entire session lasts too long gets extremely unlikely. None of our subjects ever asked about potential time constraints or the existence of a "last round" during the instruction phase or commented on it during or after the experimental session.

Matching procedure. In all sessions there participated exactly 20 students. We used an absolute stranger design, i.e. no subject played a repeated game more than once with the same opponent. After any repeated game the 10 pairs of subjects were rematched. To improve the credibility of our matching design every subject obtained an alias name. Only the alias name of the actual opponent was displayed on the screen while subjects did not get to know their own alias names. By this information policy we wanted to avoid that students could identify themselves after the experiment was over. Clearly, this matching procedure restricts the number of repeated games that any subject can play within a session up to maximally 19.

Payoffs and treatments. A repeated game is a repetition of stage games. In our experiments we tested the *six* different stage games displayed in table 1.

Game No.	payoff parameter			
	a	b	c	d
1	70	100	90	80
2	0	100	90	80
3	30	130	90	70
4	0	100	90	70
5	0	120	90	50
6	0	140	90	30

Table 1: The stage games we use in the experiments measured in ECU (= "Experimental Currency Units").

Treatments differ not only in the stage game but also in the continuation probability δ. Before every round $t \geqq 2$ the program picked a probability δ' from a uniform distribution over $[0,1]$. The next round started for all 10 matched pairs only for realizations $\delta' \leqq \delta$. After their decisions the players were informed about the respective decision of their current opponent, about their own payoff of the current round and about their own total profit of the actual repeated game.

In every session we tested two treatments and changed from the first treatment to the second after repeated game 11. In the treatments with $\delta > .75$ the change was after the repeated game 8 and we ran only 13 repeated games. Any subject's overall monetary payoff is the sum of all realized stage game payoffs. While subjects were paid for all decisions in all repeated games we did not use data from repeated games 1 to 3 for our statistical analysis. In total we observe outcome paths of 1700 different repeated games. They are independent in the sense that no two different outcome paths originate from the same pair of players and therefore should not be correlated. The first three repeated games (that are in total 300) served the only purpose that subjects not only understand but actually experience the game with its payoffs. So we used only the remaining 1400 repeated games. Different treatments within one session were only varied in the payoff parameters of the stage game not in discount factors. Table 2 gives an overview about the sessions.

Session	δ	1. treatment			2. treatment		
		Stage-game	repeated game	#Observ.	Stage-game	repeated game	#Observ.
1	.75	1	1-11	80	2	12-19	80
2	.5	3	1-11	80	4	12-19	80
3	.75	3	1-11	80	-	-	-
4	.875	3	1-8	50	4	9-13	50
5	.75	4	1-11	80	3	12-19	80
6	.75	2	1-11	80	1	12-19	80
7	.75	5	1-11	80	6	12-19	80
8	.875	4	1-8	50	3	9-13	50
9	.75	6	1-11	80	5	12-19	80
10	.75	3	1-11	80	4	12-19	80

Table 2: The table shows the games that were played in the different sessions. In session 3 we had some technical problems and could only conduct one treatment. Because of the time limit we restricted the number of repeated games to 13 in session 4 and 8. #Observ. is the number of repeated games we use for statistical tests.

7 Results

The subjects' choices are summarized in table 3. It lists the rate of cooperation in rounds 1 to 3 and for the average of all rounds.

Stage-game characteristics				Equilibrium			Round			
No	$\underline{\delta}$	δ^*	δ	$\Delta\underline{\delta}$	$\Delta\delta^*$	class	1	2	3	= All
1	.5	.667	.75	+	+	3	.356	.292	.221	.214
2	.5	.9	.75	+	-	2	.044	.025	.028	.028
3	.667	.8	.5	-	-	1	.219	.110	.070	.135
3	.667	.8	.75	+	-	2	.244	.136	.146	.154
3	.667	.8	.875	+	+	3	.390	.267	.283	.266
4	.333	.8	.5	+	-	2	.156	.013	.033	.082
4	.333	.8	.75	+	-	2	.144	.100	.143	.134
4	.333	.8	.875	+	+	3	.385	.169	.179	.217
5	.429	.667	.75	+	+	3	.559	.400	.300	.370
6	.455	.571	.75	+	+	3	.600	.463	.400	.376

Table 3: Results: Rate of cooperation in our experiments.

A “+” for $\Delta\underline{\delta} = (\delta - \underline{\delta})$ indicates that cooperation can be supported as an equilibrium and accordingly a “+” for $\Delta\delta^* = (\delta - \delta^*)$ means that cooperation is supported by our cooperation criterion $\delta \geq \delta^*$ based on the unique selection criterion S^* induced by our axioms axioms 1 through 5. To distinguish theoretical predictions we say that the equilibrium class is 1 if there are two “–” signs, is 2 for a “+”“ –” combination, and is 3 for two ”+” signs.

Definition of Data Sets Our hypotheses make predictions about the frequencies of *C-equilibria*. Accordingly, our objects of analysis are complete outcome paths. Our data set contains 1400 repeated games. The frequencies of outcome paths that contain a cooperative choice are shown in table 4 in column D1400.

But as we explained in section 5 not all outcome paths result from an equilibrium. To test our *equilibrium selection theory* more accurately we approach the set of equilibrium outcome paths from above and from below such that the remaining difference is minimal. To do this we first remove non-equilibrium paths from the data set. For example, for games where only *always-defect* is an equilibrium any outcome path containing an action C cannot result from an equilibrium. Then, we remove further non-equilibrium outcome paths, i.e. outcome paths that have no possible equilibrium-continuation-path after the the end of our observation. How can we distinguish equilibrium paths from non-equilibrium ones?

We know from the folk theorem literature that the set of feasible equilibrium payoffs (π_1, π_2) is bounded from above by two lines through the points

$$\left(\frac{a}{1-\delta}, \frac{b}{1-\delta}\right), \left(\frac{c}{1-\delta}, \frac{c}{1-\delta}\right), \left(\frac{b}{1-\delta}, \frac{a}{1-\delta}\right)$$

and from below by $\frac{d}{1-\delta}$ for any player. Let Π be the set of feasible and strictly individually rational payoffs

$$\Pi = \left\{ (\pi_1, \pi_2) \in \mathbb{R}^2 \;\middle|\; \begin{array}{l} (c-a)\pi_1 + (b-c)\pi_2 - \frac{(b-a)c}{1-\delta} \leq 0, \\ (b-c)\pi_1 + (c-a)\pi_2 - \frac{(b-a)c}{1-\delta} \leq 0, \\ \pi_1 \geq \frac{d}{1-\delta}, \pi_1 \geq \frac{d}{1-\delta} \end{array} \right\}.$$

In the experiment we observe a path $h(T) = (x_{1t}, x_{2t})_{t=1}^{T}$ and player i earns up to that period

$$\pi_{iT} = \sum_{t=1}^{T} \delta^{t-1} u_i(x_{it}, x_{jt}),$$

where $x_{it} \in \{C, D\}$ is i's action in round t.

And the set of feasible payoffs is now given by:

$$\Pi_T(h(T)) = \pi_T + \delta^{T+1}\Pi.$$

Our data set contains paths for which the *always defect* payoff is not in Π_T.[15] Any such path is not an equilibrium outcome path. Our first filtering rule defines a new data set that excludes these non-equilibrium-paths. The according data set is D1272 $= \{h(T) \in \text{D1400} \mid \Pi_T(h(T)) \cap \Pi \neq \emptyset\}$ with 1272 remaining outcome paths. Table 4 lists the frequencies of C-equilibria in column D1272.

The condition defining D1272 is necessary for a C-path to be a C-equilibrium. However, it is not sufficient. For $\delta < 1$ not all payoff combinations in the set Π are actually supported by a pure equilibrium. The folk theorem states this only for δ very close to 1. To account for this latter inaccuracy we define another data filtering rule. The last column of table 4 contains the frequencies of C-paths where we positively identified a continuation such that the observed path is the beginning of a C-equilibrium outcome path. Our positive identification procedure constructs an equilibrium path consisting of three parts. The first part is the observed path $t = 1, \ldots, T$. The second part is the compensation path and the third part establishes indefinite cooperation. In the compensation periods $t = T+1, \ldots, T+\tau$ the player i with the lower discounted payoff up to period t is allowed to defect $\tilde{x}_{it} = D$ while the other player $j \neq i$ must cooperate $\tilde{x}_{jt} = C$.

---

[15] For example in stage game No. 4 with $\delta = .75$ the feasible payoff set Π_3 after the path $h(3) = (DC, CD, CD)$ does not contain the *always defect* payoff $(280, 280)$.

The compensation part has duration $\tau = 0, 1, 2, ...$ periods. Such an outcome path is an equilibrium outcome path if τ satisfies

$$u_i(h(T), (\tilde{x}_{it}, \tilde{x}_{jt})_{t=T+1}^{T+\tau}, (C, C)^{\infty}) \geq \frac{d}{1-\delta}$$

and players always defect off-equilibrium.

We identified for 1196 paths (out of 1272) such a continuation. The test of our two hypotheses is based on D1272 but all results are robust to D1196 and D1400.

Stage-game characteristics				Equilibrium			C-Paths		
No	$\underline{\delta}$	δ^*	δ	$\Delta\underline{\delta}$	$\Delta\delta^*$	class	D1400	D1272	D1196
1	.5	.667	.75	+	+	3	.650	.650	.602
2	.5	.9	.75	+	-	2	.100	.007	.007
3	.667	.8	.5	-	-	1	.400	.000	.000
3	.667	.8	.75	+	-	2	.504	.504	.444
3	.667	.8	.875	+	+	3	.810	.839	.833
4	.333	.8	.5	+	-	2	.313	.018	.018
4	.333	.8	.75	+	-	2	.269	.164	.041
4	.333	.8	.875	+	+	3	.730	.658	.658
5	.429	.667	.75	+	+	3	.806	.803	.790
6	.455	.571	.75	+	+	3	.825	.825	.824

Table 4: Results: Rate of C-equilibria in our experiments.

Hypothesis (i): To test our hypothesis (i) we compare pairs of treatments across which $\underline{\delta}$ and δ^* and thereby $\Delta\underline{\delta}$ and $\Delta\delta^*$ change in opposite directions, as they did in the introductory example of section 2. In particular, we analyze how the frequency of C-equilibria changes in these cases. As mentioned in section 5 we expect that this direct qualitative comparison is robust with respect to subjects' distribution of individual preferences over monetary payoffs that is assumed to be more or less stable across treatments.[16]

Let us first look at pairs of repeated games with the same continuation probability $\delta = .75$. Compare game 2 in the second row of Table 4 — denoted by Γ_{22} — with the

---

[16] For another robustness check see also next section.

second game 3 in the fourth line of Table 3 — which we denote Γ_{34}.

Γ_{22}	c	d
c	90 90	100 0
d	0 100	80 80

and

Γ_{34}	c	d
c	90 90	130 30
d	30 130	70 70

In Γ_{22} we have $\underline{\delta} = 0.5$ while in Γ_{34} we have $\underline{\delta} = 0.667$. Here, the standard criterion $\underline{\delta}$ predicts that cooperation should be easier to sustain, and should therefore be observed more frequently in Γ_{22} than in Γ_{34}. Looking at changes in our alternative criterion δ^*, however, the prediction is the opposite. Cooperation should be easier, hence, be observed more frequently, in Γ_{34} where $\delta^* = 0.8$ than in Γ_{22} where $\delta^* = 0.9$. The experimental results in Table 4 show that cooperation is at least seventeen times more frequent in Γ_{34} than in Γ_{22}. This observation confirms our hypothesis (i) and is consistent with predictions based on changes in δ^*. However, it falsifies predictions based on $\underline{\delta}$.[17]

Similarly, let Γ_{11} denote game 1 at the first row of Table 4, and Γ_{47} denote game 4 at the seventh row of Table 4.

Γ_{11}	c	d
c	90 90	100 70
d	70 100	80 80

and

Γ_{47}	c	d
c	90 90	100 0
d	0 100	70 70

Here $\underline{\delta}$ predicts less cooperation in Γ_{11}, where $\underline{\delta} = \frac{1}{2}$ relative to Γ_{47} where $\underline{\delta}$ falls to $\frac{1}{3}$. The opposite predicts δ^*, as it grows from $\frac{2}{3}$ in Γ_{11} to $\frac{4}{5}$ in Γ_{47}. Again, our experimental results show that cooperation is clearly more frequent in Γ_{11} than in Γ_{47}, as predicted by δ^* and again in contrast with predictions based on $\underline{\delta}$.

Finally, let Γ_{59} denote game 5 at row nine of Table 4, and Γ_{60} denote game 6 at row ten of Table 4; and compare these two games and Γ_{47} discussed before.

Γ_{59}	c	d
c	90 90	120 0
d	0 120	50 50

and

Γ_{60}	c	d
c	90 90	140 0
d	0 140	30 30

---

[17] This and the following similar statements are supported by Wilcoxon rank sum tests, where in all cases we found p-value $\ll 0.0001$. Statistical computations are done by R (R Development Core Team (2006)).

Passing from Γ_{47} to Γ_{59} and then to Γ_{60} we observe $\underline{\delta}$ increasing from 0.164 to 0.803 and then to 0.825, predicting a monotone decrease in the rate of cooperation. On the other hand, δ^* decreases from 0.8 to 0.667 and then to 0.571, predicting the opposite, a monotone increase in cooperation. The experimental results show that indeed the frequency of cooperation increases monotonically moving from Γ_{47} to Γ_{59} and then to Γ_{60}, again confirming predictions based on δ^* and rejecting those based on $\underline{\delta}$.

We summarize these comparisons as follows.

Hypothesis (i) Result: When $\underline{\delta}$ and δ^* change in opposite directions, the frequency of cooperation changes as predicted by changes in δ^*, and contradicting predictions based on $\underline{\delta}$.

This result provides unambiguous support for our hypothesis (i). In section 8 we show that in all previous experimental studies we are aware of $\underline{\delta}$ and δ^* never change simultaneously in opposite directions across treatments. We believe that our experiments are novel in the sense that they are the first that can differentiate so clearly with respect to the two competing criterions.

Hypothesis (ii): If our hypothesis (ii) is also correct, we should find significantly more cooperation in games within equilibrium class 3 compared to any other parameter constellation. A first look at table 4 shows that the frequency of C-equilibria differs for different games. By definition of our data set the frequency of C-equilibria is 0 for games with $\delta < \underline{\delta}$ since we have removed observations where outcome paths are not supported by equilibrium behavior. But even in the the class of games where $\underline{\delta} < \delta < \delta^*$ frequency remains low.

Hypothesis (ii) Result: The overall frequency of C-equilibria in class 2 is 25%. In class 3 the frequency is 75%. A Wilcoxon rank sum test shows with very high significance that there is a difference.

A compelling graphical representation of hypothesis (ii) is provided by figure 1. It shows a logistic estimation[18] of C-equilibria frequencies dependent on the difference $(\delta - \delta^*)$ for D1272.

The estimated parameters of the logistic function are given in table 5.

For D1272 only the p-value of b is highly significant $p < 10^{-15}$. The p-value of a is 0.165. In this sense we can conclude that the value where the logistic function has its turning point is at $\delta = \delta^*$. This establishes strong evidence in favor of our hypothesis (ii). For the other data sets the turning point is also very close to δ^*.

---

[18] The logistic function is defined by $1 - \frac{1}{1+\exp(a+bx)}$.

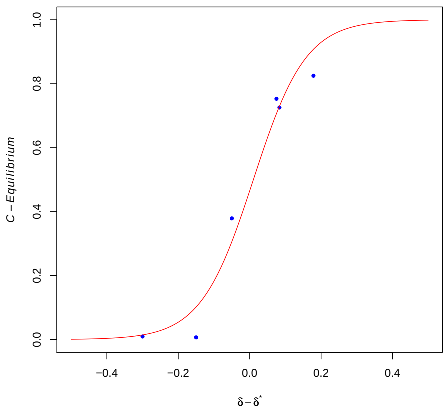

Figure 1: Frequencies of C-equilibria depending on $(\delta - \delta^*)$.

data	a	b
D1440	0.259 ***	6.560 ***
D1272	-0.096	12.887 ***
D1196	-0.316 ***	13.918 ***

Signif. codes: $0 \leq$ *** $< 0.001 \leq$ ** $< 0.01 \leq$ * $< 0.05 \leq$. < 0.1

Table 5: Logistic estimates for the frequencies of C-path.

8 Results of Other Experimental Studies

As mentioned in the introduction there are several previous experimental studies of cooperation in the infinitely repeated Prisoner's Dilemma reported in the literature, and one study that was undertaken simultaneously and independently. Table 6 offers a synthetic overview of these results. The table has the same structure as table 3 displaying our own experimental results. In some of the earlier studies the data are missing or reported only in an aggregated form, so we were not able to complete the full table. The experiments are too diverse to allow serious conclusions based on our main hypothesis. Nevertheless, if we look at the cooperation frequencies, the mean values of cooperation are 0.14 if $\delta < \underline{\delta}$, 0.26 if $\underline{\delta} < \delta < \delta^*$ and 0.43 if $\delta^* < \delta$. This is not inconsistent with our hypothesis in the sense that we expect "higher frequencies of cooperation" for the third equilibrium class, and even moderately supports it.

Anyway, it is interesting to look more closely at the studies from our perspective. Roth and Murnighan (1978) is the first study that analyzed equilibrium behavior in an uncertain horizon repeated PD game. They found that the frequencies of cooperation increase in δ. However, their values for the two cases $\delta < \delta^*$ are high compared to our observations. The reason may be that in their experiments subjects played against particular robot strategies rather than against another. Some subjects may have believed, for example, that they play against the well known strategy *Tit-For-Tat* which would

Game				Equilibrium			Round			
Study	$\underline{\delta}$	δ^*	δ	$\Delta\underline{\delta}$	$\Delta\delta^*$	class	1	2	3	All
DF	.72	.812	.5	-	-	1	.098			.098
DF	.4	.605	.5	+	-	2	.187			.180
DF	.08	.395	.5	+	+	3	.390			.353
DF	.72	.812	.75	+	-	2	.256			.203
DF	.4	.605	.75	+	+	3	.611			.587
DF	.08	.395	.75	+	+	3	.851			.764
DO14	.5	.667	.9	+	+	3	.482			.549
DO6	.5	.667	.9	+	+	3	.625			.627
D	.538	.667	.5	-	-	1				.032
D	.538	.667	.75	+	+	3				.207
D	.455	.667	.5	+	-	2				.188
D	.455	.667	.75	+	+	3				.256
FH	.407	.5	.333	-	-	1				.233
FH	.407	.5	.667	+	+	3				.360
FH	.407	.5	.833	+	+	3				.389
MR				-	-	1	.173			
MR				+	-	2	.409			
MR				+	+	3	.320			
RM	.333	.5	.105	-	-	1	.19			
RM	.333	.5	.5	+	+	3	.298			
RM	.333	.5	.895	+	+	3	.364			

Table 6: Rate of cooperation in some experiments reported in the literature, DF = Dal Bó and Frèchette (2008), DO14 = Duffy and Ochs (2006), 14 subjects; DO6 = Duffy and Ochs (2009), 6 subjects; D = Dal Bó (2005), game 1; FH = Feinberg and Husted (1993); MR = Murnighan and Roth (1982); RM = Roth and Murnighan (1978).

encourage more of them to try to build up cooperation compared to other less forgiving types of equilibria. This may also be the reason behind the high frequencies in Murnighan and Roth (1982), where subjects play against the experimenter. In the table we listed only aggregated data. The experiment includes 12 games with different payoff matrices, i.e. different values for $\underline{\delta}$, δ^* and 3 different continuation probabilities. In our terminology they observe 36 different treatments of which 14 are part of equilibrium class 1 ($\delta < \underline{\delta}$) and 16 are part of equilibrium class 3 ($\delta^* < \delta$). Only the remaining 6 belong to the middle equilibrium class 2 and offer some clue regarding our hypothesis.

The work of Feinberg and Husted (1993) frames a PD game as a duopoly game. In their study the probability δ is composed of two parts. Besides the continuation probability they consider a discount factor reducing payoffs in successive rounds. Their three treatments differ only in this latter factor (equal to 1, 0.8, 0.4). Consequently, in table 6 we listed the product of both. According to our main hypothesis we believe that the explanation for their comparatively high cooperation frequencies for equilibrium class 1 is buried in the unobserved details of this particular experimental design. Nevertheless, FH's observed cooperation rates rise markedly in equilibrium class 3 which in their case supports the traditional prediction based on $\underline{\delta}$ as much as ours.

Dal Bò (2005) used a 2 by 3 by 2 design. This means, he studies two infinite repeated games with three continuation probabilities (0, 0.5, 0.75) and compares this with repeated games of fixed duration (1, 2, 4). His fixed durations correspond to the expected lengths of the according infinite games. We did not list the results of the treatment with fixed duration nor those with the continuation probability 0. Dal Bò observes a significant difference in the levels of cooperative behavior between these two types of games. In line with game theoretical wisdom the cooperation frequencies in treatments with fixed duration are significantly lower. More interesting for our context is the second result stating that cooperation frequencies increase with δ. In particular, the difference in the level of cooperation between equilibrium class 2 (0.188) and equilibrium class 3 (0.207 and 0.256) is positive though not as clearly as in our experiment.

Duffy and Ochs (2006) study experimentally the hypothesis of Kandori (1992) that cooperation may emerge in a group of subjects randomly selected to play a PD-game. They analyze treatments with 14 and others with 6 subjects and compare the levels of cooperation in treatments in which subjects are randomly rematched after each round with fixed matching treatments. Kandori's main hypothesis is not supported by their experiment, in the sense that substantial cooperation supported by the threat of contagion does not emerge. The level of cooperation is much higher in the fixed paired treatments as in the randomly selected ones. In our table we only listed the results of the fixed matching treatments, which are closer to our topic. The treatments they used in their experiment fit in our equilibrium class 3. They found relatively high levels of cooperation, much in line with our Figure 1.

As already mentioned, at the same time as we ran our experiments Dal Bò and Fréchette (2008) ran experiments on cooperation and learning that are related to what we did, as they also partly take into account strategic risk. Their results are much in accordance with our results, as they provide independent additional support in favor of Blonski and Spagnolo (2004), which as we mentioned comes up with the same predictions as these tested here through a theoretical derivation of a risk dominance indicator for repeated games. Again, they observe slightly higher frequencies of cooperation than we do, but this could probably be explained, at least in part, by differences in their experimental design. For example, in our design a subject could never meet the same subject again, as in Dal Bò (2005). In the design of Dal Bò and Fréchette there is a positive probability to meet again the same subject. In a pool of 12 to 20 subjects every subject plays 23 to 77 repeated games. On average any subject meets any other subject 3.3 times. If a subject expects to meet the opponent over and over at later instants the assumed continuation probabilities may not correspond to the perceived ones. Also, their pool of subjects appears to contain less economics and business students than ours. Unfortunately, in their treatments $\underline{\delta}$ and δ^* were always chosen to change in the same direction, so that direct comparisons of the kind we did in Section 4.1 are not possible based on their experimental data.

There are some other experimental studies on infinite repeated games modifying the standard PD-game. For example, van Huyck, Wildenthal and Battalio (2002) reported an experiment on repeated dominance solvable games. In one of their four treatments similar to a PD-game of equilibrium class 3 they found after a time of learning a pointedly high level of cooperation. Less related are experiments by Aoyagi and Fréchette (2003) who show that in infinitely repeated prisoner's dilemma games with imperfect public monitoring the level of cooperation increases with the quality of the public signal.

To sum up, we conclude that studies that ignore strategic risk or the role of the "sucker's payoff", by only looking at changes in the incentive compatibility conditions — summarized by $\underline{\delta}$ — to predict changes in agents' ability and willingness to cooperate or collude when the environment changes, may yield incorrect or misleading results. The available experimental evidence to which we add here indicates that our δ^* clearly fares much better as a tool for predicting changes in the frequency of cooperation among real subjects when the relevant institutions change.

References

Athey, S., K. Bagwell, and C. Sanchirico (2004), "Collusion and price rigidity". *Review of Economic Studies*, 71, pp 317–49.

Axelrod, R. (1984), "The Evolution of Cooperation". *Basic Books*

Bernheim, B.D. and M.D. Whinston (1990), "Multimarket contact and collusive behavior". *Rand Journal of Economics*, 21, pp 1–26.

Blonski, M. and G. Spagnolo (2004), "Prisoners' other Dilemma". J.W.Goethe University of Frankfurt, working paper.

Dal Bó, P. (2005), "Co-operation under the shadow of the future: Experimental evidence from infinitely repeated games". *American Economic Review*, 95, pp 1591–1604.

Dal Bó, P. and G. Frèchette (2008), "The Evolution of co-operation in infinitely repeated games: Experimental evidence". Brown and New York Universities, working paper.

Duffy, J. and J. Ochs (2009), "Co-operative Behavior and the Frequency of Social Interaction". *Games and Economic Behavior*, 66(2), pp 782-812.

Feinberg, R. and T. Husted (1993), "Experimental Test of Discount-Rate Effects on Collusive Behaviour in Duopoly Markets". *The Journal of Industrial Economics*, 16, pp 153–160.

Fischbacher, U. (2007), "z-Tree: Zurich Toolbox for Ready-made Economic Experiments". *Experimental Economics*, 10(2), pp 171–178.

Friedman, J. (1971), "A Non-co-operative Equilibrium for Supergames". *Review of Economic Studies*, 38, pp 1–12.

Fudenberg, D. and J. Tirole (1991), "Game Theory". Cambridge, Mass., M.I.T. Press.

Gilo, D., Moshe Y. and Y. Spiegel (2006), "Partial cross ownership and tacit collusion". *Rand Journal of Economics*, 37(1), pp 81–99.

Harsanyi, J. and R. Selten (1988), "A General Theory of Equilibrium Selection in Games". Boston, Mass., M.I.T. Press.

Harsanyi, J. (1995), "A New Theory of Equilibrium Selection for Games with Complete Information". *Games and Economic Behavior*, 8, pp 91–122.

Heinemann, F., R. Nagel and P. Ockenfels (2009), "Measuring Strategic Uncertainty in Coordination Games". *Review of Economic Studies*, 76, pp 181–221.

van Huyck, J.B., R.C. Battalio and R.O. Beil (1990), "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure". *American Economic Review*, 80, pp 234-248.

van Huyck, J.B., J.M. Wildenthal and R.C. Battalio (2002), "Tacit Co-operation, Strategic Uncertainty, and Coordination Failure: Evidence From Repeated Dominance Solvable Games". *Games and Economic Behavior*, 38, pp 156–175.

Kandori, M. (1992), "Social Norms and Community Enforcement". *Review of Economic Studies*, 59, pp 63–80.

Kocherlakota, N. (1996), "Implications of Efficient Risk Sharing without Commitment". *Review of Economic Studies*, 63, pp 595–609.

Ligon, E., P. Thomas and T. Worrall (2002), “Informal Insurance Agreements with Limited Commitment: Theory and Evidence from Village Economies”. *Review of Economic Studies*, 69, pp 209–244.

Mailath, G. and L. Samuelson (2006), “Repeated Games and Reputations: Long-run Relationships”. Oxford University Press.

Motta, M. (2004), “Competition Policy, Theory and Practice”. Cambridge University Press.

Murnighan, K. and A. Roth (1983), “Expecting Continued Play in Prisoner’s Dilemma Games: A Test of Several Models”. *Journal of Conflict Resolution*, 27, pp 279–300.

Rapoport, A. and A.M. Chammah (1965), “Prisoner’s Dilemma”. University of Michigan Press

R Development Core Team (2006), “R: A language and environment for statistical computing”. R Foundation for Statistical Computing”. Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org.

Roth, A. and K. Murnighan (1978), “Equilibrium Behaviour and Repeated Play of the Prisoner’s Dilemma”. *Mathematical Psychology*, 17, pp 1123–1142.

Tirole, J. (1988), “The Theory of Industrial Organization”. Cambridge, Mass., M.I.T. Press.