

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

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

version October 19, 2010

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^* is also useful as a tool for applied comparative statics exercises as it results in a critical discount factor δ^* strictly larger than $\underline{\delta}$, the standard criterion often used in applications. In an experimental test we find a strong predictive power of our proposed criterion S^* that gets even more pronounced if we restrict attention to equilibrium outcome paths — i.e. if we filter out disequilibrium behavior. For parameter changes where the standard and our criterion predict differently, changes in observed cooperation follow predictions based on δ^* .

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

*We thank the E-finance Lab at the J.W.Goethe University in Frankfurt for funding the experiment. For valuable comments, we are grateful to Werner Güth, Edi Karni, Wolfgang Leininger, Eric Maskin, John Nash, Larry Samuelson, Reinhard Selten, and many seminar participants in Karlsruhe, Jena, at the ESA conference in Rome, Games 2008 at Northwestern University, “Sozialwissenschaftlicher Ausschuss des Vereins für Socialpolitik”, and the Brazilian Workshop of the Game Theory Society in Sao Paulo 2010.

[†]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? In this article we study this question in the context of the infinitely repeated Prisoner's Dilemma. Progress in this topic is important for game theory itself, but is also critical for the numerous applications of infinitely repeated games in economics, sociology, political science, biology and other disciplines.

We propose an axiomatic approach that formulates (i) a minimal set of simple and intuitive conditions on the model primitives a sensible selection theory should satisfy and (ii) a more comprehensive set of conditions resulting in a unique cooperation criterion. While the parsimonious formulation (i) has implications for the qualitative question whether cooperation should increase or decrease when parameters change the more specific model (ii) based on a longer list of axioms comes up with a quantified parameter-frontier above which it predicts cooperation. An axiomatic approach can only convince with simple and intuitive axioms that build only on model primitives. The first part of the paper builds intuition for the novel axioms and derives their theoretical implications. The second part presents results from a laboratory experiment testing the implications from both the parsimonious (i) and comprehensive (ii) versions of the theory.

Where do we stand? The main emphasis of the theory of infinitely repeated games has been on folk theorems characterizing the payoff space if player's patience approaches infinity. In particular it states that in this case there exist many equilibria including cooperative ones and non-cooperative ones. There is broad consensus in the profession that equilibrium multiplicity causes a big lack of predictive power in this theory.¹ In the perfect monitoring context equilibrium non-deviation constraints for cooperative equilibria are easier to be satisfied when gains from cooperation are larger, when short-run gains from 'cheating' are smaller, when the severeness of punishment for cheating is larger, and when players are more patient or interact more frequently.

In a symmetric setting, the non-deviation conditions are identical for all players and their tightness is often quantified by the minimum discount factor $\underline{\delta}$ for which cooperation is sustainable in equilibrium. Recent experimental evidence has measured cooperation frequencies against parameter variations within this inequality and appears to support the view that easier to satisfy non-deviation constraints lead to more cooperation.² Applied theory has built on changes in $\underline{\delta}$ when trying to design institutions or identify real world situations that are more or less conducive to cooperation, for example 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 All these applications implicitly interpret cooperation to be more likely when $\underline{\delta}$ falls and less likely when it goes up. A theoretical foundation for this interpretation of $\underline{\delta}$ is Pareto-dominance as an equilibrium selection

¹Friedman (1971) was among the first to formulate these observations formally. See e.g. Mailath and Samuelson (2006) and Fudenberg and Tirole (1991, Ch. 5) for excellent surveys.

²See e.g. Dal Bo (2005).

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

criterion.⁴ From here, therefore, we simply call $\underline{\delta}$ the *standard criterion*.⁵

Our theoretical contribution Equilibrium selection in the infinitely repeated Prisoners' Dilemma can be organized into increasing layers of detail. The coarsest question is *whether* players cooperate at all. If players do cooperate one could ask how frequently they cooperate. This raises more particular questions how players can learn and maintain to cooperate. This in turn leads to the next level of detail how cooperation is supported — i.e. how strategies react to off-equilibrium misbehavior. One of the puzzles in repeated game theory is that for each of these questions there exist not just more than one but a whole universe of possible and consistent answers. Intuitively, it is in the best interest of players to build up cooperation as fast as possible and then cooperate as much as possible if they can cooperate at all. This intuition emphasizes the first and most basic question whether players cooperate at all as the most crucial and consequential part of the problem. By identifying any equilibrium with cooperative actions on its outcome path as *cooperation* we concentrate here just on this basic question.

There are various possibilities to motivate selection criteria. For example one could formulate an evolutionary model, introduce different 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.⁶ Although we are sympathetic to all these methods we favor even more a selection theory that is independent of the modeler's taste with respect to theory. Our way to do this is first to formulate and motivate three minimal axioms any sensible theoretic selection model should satisfy. Second, we propose and motivate two additional axioms that are sufficient to end up in a unique prediction regarding whether or not to cooperate. 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 axiom is most standard in the literature on equilibrium selection. It corresponds to Harsanyi and Selten's *invariance with respect to isomorphisms*. The second axiom requires that once players cooperate increases in players' discount factor should not

⁴Efficiency or Pareto-dominance is the most widely accepted criterion since it is motivated normatively and it is the relevant criterion to describe the boundary of the equilibrium payoff 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 on this.

⁵It should be pointed out that there exist many reasonable circumstances where the comparative statics of the standard criterion and our criterion coincide qualitatively. This will become clear after seeing our numerical example in section 2. The whole emphasis of this project is on those theoretical cases where the predictions differ. Since this article pursues the theoretical question when players select cooperative behavior here we do not attempt to identify those applications where our objection matters most and derive its implications. This, however, is a promising avenue of future research.

⁶Compared 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.

destroy cooperation. It implies that any selection criterion satisfying this axiom can be formulated in terms of critical discount factor depending on the stage game payoff parameters above which players cooperate and below which they don't cooperate. This makes sure that any two selection criteria satisfying axiom 2 are easily comparable through the two corresponding critical discount factors. This observation simplifies equilibrium selection theory for repeated games a great deal and will be used frequently in this article.

The crucial novel idea is axiom 3. It 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 avoid this risk by never cooperating. Conversely, this kind of *strategic risk* for cooperation gets smaller if the sucker's payoff gets larger and eventually approaches the defection equilibrium payoff, in which case the standard non-deviation constraints remain the only concern. In the parsimonious setting axiom 3 is the critical axiom that is violated by the standard selection criterion.⁷

Axioms 1 through 3 determine the qualitative comparative static properties of the selection criteria satisfying them. They do not quantify, however, the critical benchmark where cooperation breaks down. This is what the remaining two axioms *incentive independence* and *equal weight* accomplish. Together with the axioms 1 to 3 they induce a unique selection criterion and by axiom 2 a critical lower bound on discount factors δ^* below which cooperation breaks down. The latter two axioms will be motivated in detail in sections 2 and 3. They can be seen as one way to continuously and monotonously extend the two principles behind axiom 3. In contrast to the standard criterion, the alternative criterion depends on all payoff-parameters, including the sucker's payoff.

While various related criteria have been discussed or tested in the earlier literature⁸ we are not aware of theoretical foundations that single out one criterion over another as the relevant cooperation predictor.⁹

Our experimental contribution We test our criterion against the standard criterion with a laboratory experiment that simulates infinitely repeated games with random

⁷We have analyzed several possibilities to generalize the crucial axiom 3 with respect to stage games. Given that there is a broad consensus in the profession that the Prisoner's Dilemma is a relevant starting point to analyze cooperation and at our current stage of understanding none of the possible generalizations is obviously superior to others we decided to stay within the Prisoner's Dilemma world in this article.

⁸See 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.

⁹In contrast, in one shot games several theory based experimental studies have shown that strategic risk has explanatory power. Efficient but risky equilibria are often not chosen if gains are small or coordination requirements high (see for example van Huyck et al. (1990) and Heinemann et al (2009)). An exception for the repeated Prisoner's Dilemma is Blonski and Spagnolo (2004) who build on Harsanyi and Selten's (1988) concept of risk dominance and link this to infinitely repeated games. They identify the same criterion as this paper by a different method and thereby indirectly connect this axiomatic approach to Harsanyi and Selten's risk dominance concept.

continuation and matching rules and many experimental subjects. Our experimental setup tests our theory, changing parameters so that $\underline{\delta}$ and δ^* may change in different directions when comparing couples of treatments.¹⁰

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. This result is very robust since any equilibrium selection theory that satisfies only 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 maintains some predictive power only in situations in which δ^* remains constant, hence as a 'residual' of the alternative criterion.

Our second hypothesis derived from our alternative equilibrium selection criterion positively 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 hypothesis. This latter observation is also consistent with recent independent experimental evidence in Dal Bó and Frechette (2008).¹¹

We first test our hypotheses against the full data set. Then, in a second step — to improve the connection between theory and experiment — we investigate the question to which extent the alternative selection criterion predicts the mode of behavior only 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 quality 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 80% of our observed outcome paths — i.e. experimental subjects' actual behavior — are indeed equilibrium outcome paths of the repeated game which conversely 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 behavior turns out to fade rapidly towards non-cooperation.

To sum up, our axiomatic approach predicts under which conditions players cooperate in the repeated Prisoner's Dilemma. Our experimental evidence supports this novel

¹⁰We 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 7.

¹¹Our experiments were run independently and simultaneously in (2006) and happen to be complementary. Their focus is on learning effects and players' strategies, while ours is on equilibrium selection and strategic risk. In Dal Bó and Frechette's experimental treatments payoffs always change such that $\underline{\delta}$ and δ^* are constant or change in the same direction. Further, they do not attempt to rule out group learning effects as their players meet again the same opponents with positive probability when they are re-matched. See again section 7 for a more detailed comparison.

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, in all studies there were 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 and show that not all of their unexplained variation can be captured by our theory. Nevertheless, given the methodological differences of some of the earlier studies we were still 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 numerical example and motivates the critical axioms 3 and 5. In section 3 we formulate and further discuss all axioms. In Section 4 we show how the three first more general axioms determine a general selection criterion, and together with two additional axioms the more specific criterion S^* . Based on these general and specific criteria we formulate two hypotheses to be tested against the standard criterion. Section 5 describes the experimental design, while section 6 presents our experimental results. Section 7 surveys previous experimental investigations within this setting and compares their evidence with ours where possible.

2 Intuition for Critical Axioms

Consider the following two Prisoner’s Dilemma stage games Γ_1 and Γ_2 given by

Γ_1	C	D	with	Γ_2	c	d
C	2	3		c	2	2.5
	2	0.9		d	2	-99
D	0.9	1			-99	1
	3	1		2.5	1	

We compare 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 $3 + \frac{\delta}{1-\delta}$ from a single deviation followed

by indefinite defection of both players. This yields a lower bound $\delta \geq \underline{\delta}(\Gamma_1) = \frac{3-2}{3-1} = \frac{1}{2}$ on discount factors in $\Gamma_1(\delta)$. Correspondingly, this lower bound for $\Gamma_2(\delta)$ is $\underline{\delta}(\Gamma_2) = \frac{1}{3}$.

Those players for whom Pareto-efficiency is the relevant criterion should cooperate in $\Gamma_2(\delta)$ but not in $\Gamma_1(\delta)$ if $\delta \in [\frac{1}{3}, \frac{1}{2})$. The range of discount factors for which cooperation can be supported as an equilibrium is larger in $\Gamma_2(\delta)$. In absence of another criterion the applied literature building on the standard criterion $\underline{\delta}$ and its interpretation has concluded that an environment or institution as in $\Gamma_2(\delta)$ is more conducive to cooperation. However, in this example we intuitively expect, that even for quite patient players it is far more difficult to cooperate in $\Gamma_2(\delta)$ than in $\Gamma_1(\delta)$. For example, real players with a discount factor of, say, $\delta = .9$ may be able to build up cooperation in $\Gamma_1(.9)$ whereas they most likely would not dare to cooperate in $\Gamma_2(.9)$.

For the discount factor $\delta = .9$, the long run incentives to cooperate instead of defecting forever are identical for both infinitely repeated games $\Gamma_1(.9)$ and $\Gamma_2(.9)$ and are at most $\delta \left(\frac{2}{1-\delta} - \frac{1}{1-\delta} \right) = 9$. This is an upper bound on what is at stake in the future in both games — sometimes called “the shadow of the future”.

In any cooperative equilibrium a player in some period considers to play a cooperative rather than a defective action. This player faces two possibilities. Either the opponent cooperates as well in this period denoted as case (c) or defects denoted as case (d). In a Prisoner’s Dilemma *both* cases yield immediate gains in favor of choosing the defective action. The short run gains for case (c) are $3 - 2 = 2$ in $\Gamma_1(.9)$ and $2.5 - 2 = .5$ in $\Gamma_2(.9)$, hence small in both examples compared to the “shadow of the future” given by 9. However, for case (d) the short run gains from defecting are $1 - .9 = .1$, hence small in $\Gamma_1(.9)$ but $1 + 99 = 100$, i.e. overwhelming in $\Gamma_2(.9)$. Obviously, the two examples differ most in the sucker’s payoff .9 versus 99.

Our critical axioms 3 and 5 build on the idea that in a game with cooperative and non-cooperative equilibria players cannot know whether case (c) or case (d) is relevant. Since equilibrium non-deviation constraints such as those defining $\underline{\delta}$ only depend on case (c) they do not take into account this sucker’s payoff. Axiom 3, called *boundary conditions*, intuitively formulates equilibrium selection for the two most obvious and extreme cases where incentive (d) is overwhelming — i.e. converges to infinity, and where it is negligible — i.e. converges to 0. Moreover, the intuition for *equal weighting* axiom 5 is that both cases (c) and (d) bear equal weight in comparing short run and long run incentives as long as both can be supported by equilibrium strategies. All other axioms are less critical in the sense that they are also satisfied by the classical criterion. We postpone the further discussion of the axioms to the following section once the notation is set up.

Both axioms 3 and 5 relate to “strategic risk” which is not taken into account by equilibrium non-deviation constraints for a cooperative equilibrium. Harsanyi and Selten (1988) formulate axioms for equilibrium selection and define strategic risk in 2×2 -coordination games with two equilibria. They emphasize the desirability of an axiomatic approach for more general settings.¹² In this article we follow this route and propose

¹²Harsanyi and Selten’s more general definition for strategic risk, in contrast, is not based on axioms. It imposes more structure, i.e. the bicentric prior and the tracing procedure and is not defined for infinite

such a theory for the infinitely repeated Prisoner’s Dilemma and thereby go a first step towards an equilibrium selection theory for repeated games.¹³

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 $2c > b + a$.¹⁴ 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

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

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

Call any equilibrium of $\Gamma(\delta)$ a *D-equilibrium* if its outcome path is $((D, D), (D, D), \dots)$ — i.e. contains only defective actions. Conversely, any 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.

A *selection criterion for cooperation* is defined as a subset $S \subseteq Q$ of parameter values for which 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

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

The interesting question is therefore for which part of this set players indeed cooperate, i.e. select C-equilibria. We consider first the two extreme cases for cooperation criteria. The most “cooperation friendly” criterion denoted by S_C is defined by always selecting

games including infinitely repeated games.

¹³Similarly as 2×2 - coordination games are prototypical for equilibrium selection in one-shot games we believe that the equilibrium selection problem in the repeated Prisoner’s Dilemma is prototypical for equilibrium selection in more general repeated games. Though we hope for the contrary, it is well possible that similarly as in Harsanyi and Selten’s context there is no natural and obvious way to generalize this axiomatic approach without imposing a lot more structure.

¹⁴It is possible to normalize parameters as $c \equiv 1$ and $d \equiv 0$ such that the relevant incentives only depend on 2 remaining parameters $a < 0 < 1 < b$. See for example Stahl (1991). However, since this imposes an assumption on preferences related to our axioms we postpone here the discussion.

a C-equilibrium when there is one, i.e. $S_C := \{(a, b, c, d, \delta) \in Q \mid \delta \geq \underline{\delta}\}$. 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 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.

3.1 Parsimonious Cooperation Axioms 1 – 3

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

Axiom 1 (Positive Linear Payoff Transformation Invariance) *Let $\tau : \mathbb{R} \rightarrow \mathbb{R}$ 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*.¹⁵ The interpretation of this axiom is that players’ payoffs are cardinal — for example represent their von Neumann-Morgenstern utility functions — and thereby abstract away from all framing effects like the choice of the 0-level or any scaling. Adding an arbitrary constant β implies that for all cooperation selection criteria satisfying axiom 1 only the payoff differentials rather than the absolute payoffs matter. For further reference we call

1. $\frac{\delta}{1-\delta}(c - d)$ the long run incentive to cooperate,
2. $b - c$ the short run incentive to defect if the opponent cooperates and
3. $d - a$ the short run incentive to defect if the opponent defects as well.

It turns out to be useful to normalize the first payoff differential $c - d$ by multiplying it with $\frac{\delta}{1-\delta}$ since this is what we called the shadow of the future in our numerical example. It defines an upper bound on what is lost by picking the defective continuation equilibrium instead of a cooperative one with a symmetric Pareto efficient outcome path.¹⁶ If axiom 1 were the only axiom the parameter set could even shrink by two degrees of freedom. Since

¹⁵Harsanyi and Selten’s *isomorphisms* also allow for permutation of players’ names which do not matter in symmetric games as ours.

¹⁶The symmetric Pareto-efficient cooperative outcome path can always be supported by maximal punishment as grim trigger punishment or asymmetric punishments a la van Damme (1989) if it is an equilibrium outcome path. More generally, a big variety of other off-equilibrium behavior supports cooperative equilibria once $\delta > \underline{\delta}$. By defining this upper bound maximal punishment such as grim trigger plays a particular role in defining the long run incentive.

players do not care about scaling one could specify parameters α, β such that $d = 0, c = 1$ and thereby

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

However, it is more intuitive to continue with all 5 parameters since all three incentives matter for the comprehensive set of axioms and in our experiments we also vary the long run incentive and thereby the payoff difference $c - d$.

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

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

Axiom 2 is called δ -monotonicity since it implies the weaker condition

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

It is a little stronger though, i.e. the latter monotonicity condition does not imply axiom 2. The difference is that the weaker monotonicity condition does not specify a tie-breaking rule, i.e. whether players cooperate once $\delta = \delta_S(a, b, c, d)$. However, we regard this as a purely technical matter. Beyond the latter δ -monotonicity condition axiom 2 simply picks one of two simple tie-breaking rules, namely, always cooperative behavior for $\delta = \delta_S(a, b, c, d)$. The critical discount factors for the two extreme cooperation selection criteria are given by $\delta_{S_C}(a, b, c, d) = \underline{\delta}$ and $\delta_{S_D}(a, b, c, d) = 1$.

As mentioned before the long run incentive defines an upper bound on what is at stake in the future. This upper bound in turn is defined by the harshest possible sub-game perfect punishment payoffs, i.e. grim trigger strategies. Any more forgiving strategy decreases the long run incentive to cooperate by raising the future payoff of being punished. It is natural therefore that if a criterion selects defection with grim trigger strategy it will also do so with any milder type of strategies. In this sense grim trigger strategies naturally play a salient role for the critical discount factor of any selection criterion. For the motivation of the last two axioms 4 and 5 it is very helpful to keep in mind here that for any cooperation selection criterion S that satisfies both axioms 1 and 2 the function $\delta_S(a, b, c, d)$ could be written as $\delta_S(c - d, b - c, d - a)$, i.e. only depends on the payoff differentials $c - d, b - c$ and $d - a$.

The next axiom is related to strategic risk and aims to formulate the least restrictive condition representing the intuition provided in our introductory example in section 2.

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, $a \rightarrow -\infty \Rightarrow \delta_S(a, b, c, d) \rightarrow 1$.*

ii. *Upper Boundary Condition: If there are C-equilibria then they are selected by S if the sucker's payoff is high enough. Formally, $a \rightarrow d \Rightarrow \delta_S(a, b, c, d) \rightarrow \underline{\delta} = \frac{b-c}{b-d}$.*

The basic intuition for the boundary condition axiom 3 was already discussed 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 conditions 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 $\frac{b-c}{b-d} < \delta_S(a, b, c, d) < 1$. In particular, this means that S_C does not satisfy condition 3.i while S_D does not satisfy condition 3.ii.

3.2 Comprehensive Cooperation Axioms 4 – 5

The following *incentive independence axiom* 4 builds on a more structured view of a repeated game as a trade-off between long-run and short-run incentives.¹⁷ As mentioned before it follows the logic of the linear payoff transformation invariance axiom 1 together with the δ -monotonicity axiom 2. In which sense? We have seen that the critical discount factor $\delta_S(c-d, b-c, d-a)$ is a function that depends only on the payoff differentials $b-c$, $c-d$ and $d-a$. Now suppose any of the three incentives changes. The following axiom 4 then requires that the relative weight between the other two incentives remains unaffected. Neither should any two of these 3 incentives reinforce or weaken each other. Mathematically, this means that $(a, b, c, d, \delta) \in S$ is representable by an additively separable function of the three incentives.¹⁸

Axiom 4 (Incentive Independence) *The three incentives, i.e. long-run cooperation incentive $\frac{\delta}{1-\delta}(c-d)$ and the two short run incentives $b-c$ and $d-a$, are independent. Formally, there exists an additively separable function $\sigma\left(\frac{\delta(c-d)}{1-\delta}, b-c, d-a\right)$ of the three incentives such that*

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

It is interesting to note that both extreme cooperation criteria, the standard criterion S_C and also S_D , do satisfy the incentive independence axiom 4.

¹⁷Mailath and Samuelson (2006) in their introductory chapter motivate repeated games by the trade-off between short run and long run incentives.

¹⁸Note that for any given delta the long run incentive $\frac{\delta(c-d)}{1-\delta}$ is proportional to the parameter difference $c-d$ so that the following independence axiom 4 relates the three parameter differences to each other.

The final axiom 5 quantifies the threshold that separates C-equilibria from D-equilibria. It builds on the more structured view of the incentive independence axiom 4 and compares only the two short run incentives $b - c$ and $d - a$. In principle $\delta_S(c - d, b - c, d - a)$ could reflect any relative weighting between these two payoff differentials. Which of the two short run incentives is more relevant actually depends on the action chosen by the other player. We argue that at the beginning of the game without prior knowledge beyond the primitives of the model any other than the equal weighting rule would impose an exogenous asymmetry and accordingly 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 actions of the other player and thereby the relevant short term incentives symmetrically.

Axiom 5 (Equal Weight) *The two short run incentives, $b - c$ and $d - a$ carry equal weight. Formally, consider the two payoff parameter constellations a, b, c, d and a', b', c, d with similar long run incentive $c - d$ but exchanged short run incentives $b - c = d - a'$ and $d - a = b' - c$. This implies*

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

The *standard criterion* S_C violates axiom 5. It can be seen as the case where players compare the long run incentive only with the first of the two short run incentives since $\frac{\delta}{1-\delta}(c - d) \geq b - c \Leftrightarrow \delta \geq \frac{b-c}{b-d} = \underline{\delta}$. In other words, in the standard criterion players put 100% weight on $b - c$ — i.e. the short run incentive to defect if the opponent cooperates — but fully neglect $d - a$ and thereby the influence of the sucker’s payoff. According to our earlier interpretation this violates the principle of insufficient reason in the sense that by comparing long run and short run incentives a player only considers the case where the opponent cooperates once it is an equilibrium without any prior knowledge about the chosen equilibrium based on the primitives of the model. The *standard criterion* S_C thereby violates the boundary conditions axiom 3 and the equal weighting axiom 5 but satisfies axioms 1, 2 and 4.

4 Theoretical Results and Predictions

Theoretical results. Our first general result builds only on the first three axioms. By imposing less structure it is more robust. 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 $(a, b, c, d), (a', b', c', d') \in \{(a, b, c, d) | b > c > d > a, b + a < 2c\}$ 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, $\delta_{\tilde{S}}(a, b, c, d) > \delta_{\tilde{S}}(a', b', c', d')$, but $\delta_{S_C}(a, b, c, d) < \delta_{S_C}(a', b', c', d')$. \square*

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

Our second more specific result quantifies the cooperation-threshold and 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 characterized by*

$$(a, b, c, d, \delta) \in S^* \Leftrightarrow \delta \geq \delta_{S^*}(a, b, c, d) = \delta^* := \frac{b - a - c + d}{b - a}.$$

PROOF Axiom 4 implies that there are functions $\sigma_1(\cdot), \sigma_2(\cdot)$ such that $\sigma(\cdot)$ can be written in the form

$$\sigma\left(\frac{\delta(c-d)}{1-\delta}, b-c, d-a\right) = \frac{\delta(c-d)}{1-\delta} - \sigma_1(b-c) - \sigma_2(d-a)$$

with

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

This implies

$$\sigma(\cdot) \geq 0 \Leftrightarrow \delta \geq \frac{\sigma_1(b-c) + \sigma_2(d-a)}{c-d + \sigma_1(b-c) + \sigma_2(d-a)}.$$

Axiom 3.ii implies for $a \rightarrow d$

$$\frac{\sigma_1(b-c) + \sigma_2(0)}{c-d + \sigma_1(b-c) + \sigma_2(0)} = \frac{b-c}{b-d}$$

or $\sigma_1(b-c) + \sigma_2(0) = b-c$. Further, axiom 3.i yields for $a \rightarrow -\infty$

$$\frac{\sigma_1(b-c) + \sigma_2(d-a)}{c-d + \sigma_1(b-c) + \sigma_2(d-a)} \rightarrow 1$$

from the left since $\frac{\sigma_1(b-c) + \sigma_2(d-a)}{c-d + \sigma_1(b-c) + \sigma_2(d-a)} \in [0, 1)$. This implies that $\sigma_1(b-c) + \sigma_2(d-a) \rightarrow \infty$ for $a \rightarrow -\infty$. Finally, axiom 5 together with $\sigma_1(b-c) + \sigma_2(0) = b-c$ implies

$\sigma_1(b - c) + \sigma_2(0) = \sigma_1(0) + \sigma_2(b - c) = b - c$ which implies $\sigma_1(x) = \sigma_2(x) = x$. These functional forms of σ_1, σ_2 yield

$$\delta \geq \frac{\sigma_1(b - c) + \sigma_2(d - a)}{c - d + \sigma_1(b - c) + \sigma_2(d - a)} = \frac{b - c + d - a}{c - d + b - c + d - a} = \frac{b - a - c + d}{b - a} = \delta^*.$$

Predictions. As real world observers we 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 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 other 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 \tilde{S} and the appropriately chosen payoff parameters for pairs of Prisoner’s Dilemma stage games depend on the selection criterion that is to be tested. Nevertheless, besides testing the comparative statics properties of selection criteria satisfying axioms 1, 2, and 3 we also want to test the quantified 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 since it rests on two additional axioms.

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 .

5 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”), $1 \text{ €} = 270 \text{ ECU}$.

Treatments differ not only in the stage game but also in the continuation probability δ . Before every round $t \geq 2$ 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 $\delta' \leq \delta$. 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.

Learning. As pointed out in the introduction in this paper we are interested in cooperation as an equilibrium selection phenomenon per se. Therefore we want to deemphasize learning, as (i) learning within a repeated game and (ii) learning by playing many repeated games against an increasing sample of partners who gain their own experience from the same pool of partners. Both learning processes are involved in several respects. By observing an evolving outcome path the within-game-learning-process (i) step by step rules out potential equilibria.¹⁹ In this context our equal weight axiom is not innocuous since over time players can base their beliefs on observed play which should affect their weighting. While this limitation decreases the predictive power of the comprehensive axioms part of our theory over time it does not affect the the initial period of any game and our parsimonious axioms predictions. Learning process (ii) by playing many repeated

¹⁹See e.g. Kalai and Lehrer (1993).

games contains experimenting and experiencing the payoffs of the games under various histories against different partners. By doing this as a player you may learn at the same time about your partners' (sample of) dispositions and chosen strategies and your partner's beliefs about the according distributions in the population, about the learning of your partners and so on. One might think that the best way to disregard this complex dynamic process would be to let every player just play one repeated game. The downside of this latter method would be, however, to study only fully unexperienced behavior. While this is interesting in itself, in particular in simpler games it is not what we are after in this very complex game. We want to study the behavior of players with sufficient experience to "understand" at least the termination probability and the basic tradeoffs in the payoff parameters. We believe that for an infinitely repeated game with termination probability a small sample of correctly answered test questions cannot substitute for experience. Hence, in order to test our hypotheses in the lab we must compromise to some extent. Our way of doing this in this project is to implicitly suppose that most of this learning happens in the first couple of games since the incremental gain in information decreases over time.²⁰ We take account of this by not evaluating the first three repeated games (that are in total 300) where players already got paid. We used only the remaining 1400 repeated games. Since variation of the termination probability would be rather abstract we did not alter it within one session. We varied only the payoff parameters of the stage game, and at most once during one session. Table 2 gives an overview about the sessions. The resulting set of observed truncated outcome paths is called D1400.

Equilibrium filtering. In an experiment on repeated games we cannot observe strategies. The only directly observable information are outcome paths. Since not all outcome paths are equilibrium outcome paths this provides an elegant tool to disentangle 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 equilibrium selection theory is to separate observed outcome paths into equilibrium outcome paths and non-equilibrium outcome paths.²¹ The predictive power of any equilibrium selection criterion can then only regard the sub-sample of equilibrium outcome paths.

To precisely identify all equilibrium outcome paths one would have to characterize the set of all infinite equilibrium outcome paths and check for any observed truncated path $h(T) = ((x_{11}, x_{21}), \dots, (x_{1T}, x_{2T}))$ of length T if there is an infinite equilibrium outcome path starting with this observed path. Here we propose a pragmatic and much simpler way of approaching this identification. We proceed in two steps. The first rule formulates

²⁰The experimental evidence in Dal Bó and Frechette (2008) supports this assumption. See their fig 1. They let some players play more than 60 repeated games. It turns out that after a couple of periods cooperation rates tend to be rather stable.

²¹Since 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.

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.

a necessary condition and the second rule a sufficient condition for paths to be equilibrium outcome paths.

Filtering rule 1: Remove all C-paths from the sample that are not individually rational or the discount rate is too low, i.e. do not satisfy

$$\sum_{t=1}^T \delta^{t-1} u_i(x_{1t}, x_{2t}) + \delta^T \pi_i \geq \frac{d}{1-\delta} \text{ for } i = 1, 2 \text{ or}$$

$$\delta \geq \underline{\delta}$$

where the continuation payoffs (π_1, π_2) of the unobserved infinite paths are elements of the equilibrium payoff space which we describe after these rules. The remaining set of paths surviving filtering rule 1 contains 1272 out of 1400 paths and is called D1272.

Filtering rule 2: From the remaining set D1272 remove all paths for which we cannot construct the following equilibrium outcome path. The first part is the observed path $t = 1, \dots, T$. The second part establishes indefinite cooperation for both players. If both players always defect off-equilibrium such an outcome path is a subgame perfect equilibrium outcome path if all non-deviation conditions

$$\sum_{\tau=1}^T \delta^{\tau-1} u_i(x_{1\tau}, x_{2\tau}) + \delta^T \frac{c}{1-\delta} \geq \sum_{\tau=1}^{t-1} \delta^{\tau-1} u_i(x_{1\tau}, x_{2\tau}) + \delta^{t-1} u_i(\tilde{x}_{i,t}, x_{-i,t}) + \delta^t \frac{d}{1-\delta}$$

for both players $i = 1, 2$ and $\forall t \in \{1, \dots, T\}$ are satisfied. Here $\tilde{x}_{i,t} \neq x_{i,t}$ is the deviating action (for example D instead of C) for player i relative to the observed path in period t and players switch to indefinite defection from then. We know that for $t > T$ the conditions are satisfied if $\delta \geq \underline{\delta}$. The remaining set of paths surviving filtering rule 2 contains 1098 out of 1272 paths and is called D1098.

We know from the folk theorem literature that for $\delta \rightarrow 1$ the equilibrium payoff space is bounded by a polygon Π given as

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

The Pareto frontier of Π is formed by two lines through the points $(\frac{a}{1-\delta}, \frac{b}{1-\delta})$, $(\frac{c}{1-\delta}, \frac{c}{1-\delta})$, and $(\frac{b}{1-\delta}, \frac{a}{1-\delta})$. The lower bound $\frac{d}{1-\delta}$ for any player is called individual rationality constraint. The equilibrium payoff space for any given δ must be contained in Π . Hence, a necessary condition for a path to survive filtering rule 1 is that the individual rationality condition in filtering rule 1 is satisfied for both players and at least for one of the three vertices on the Pareto frontier of Π . Otherwise individual rationality is violated for all points of Π and therefore also for the equilibrium payoff space of the repeated game. Besides inequality $\delta \geq \underline{\delta}$ this procedure defines for any $h(T)$ a set of another 6 inequalities to be checked for filtering rule 1.²² Accordingly there is a finite set of $2T$ non-deviation inequalities to be checked for filtering rule 2.

If our experimental results do not depend on whether we test the predictions for D1272 versus D1098 we have strong reasons to believe that they would hold as well for the set of equilibrium outcome paths that must be between the two samples. Moreover, though this paper is about equilibrium selection there is no reason why our axioms should not be empirically valid for non-equilibrium behavior.²³ Therefore we test our predictions as well for D1400.

We expect another bias in our experimental data that stems again 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

²²For example in stage game 4 with $\delta = .75$ the feasible payoff set $\delta^3\Pi$ after the path $h(3) = (DC, CD, CD)$ does not contain the *always defect* payoff (280, 280).

²³We thank Werner Güth and Wolfgang Leininger who both pointed this out independently.

defective behavior is underestimated by our 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 $\delta > \delta_S(a, b)$.

6 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 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.

Our hypotheses make predictions about the frequencies of *C-equilibria*. Our data set contains outcome paths of 1400 repeated games. The frequencies of C-outcome paths — i.e. those that contain a cooperative choice — are shown in table 4 in column D1400.

The frequencies in table 4 are means over all repeated games of a treatment. Our statistical analysis may be inaccurate if we ignore learning and correlation — i.e. if the behavior in one repeated game depends on the behavior in previous repeated games. Note that we have an absolute stranger design, but subjects may still learn to cooperate

Stage-game characteristics			Equilibrium				C-Paths		
No	$\underline{\delta}$	δ^*	δ	$\Delta\underline{\delta}$	$\Delta\delta^*$	class	D1400	D1272	D1098
1	.5	.667	.75	+	+	3	.650	.650	.529
2	.5	.9	.75	+	-	2	.100	.007	.007
3	.667	.8	.5	-	-	1	.400	.000	.000
3	.667	.8	.75	+	-	2	.504	.504	.331
3	.667	.8	.875	+	+	3	.810	.839	.762
4	.333	.8	.5	+	-	2	.313	.018	.018
4	.333	.8	.75	+	-	2	.269	.164	.041
4	.333	.8	.875	+	+	3	.730	.658	.645
5	.429	.667	.75	+	+	3	.806	.803	.767
6	.455	.571	.75	+	+	3	.825	.825	.801

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

or defect based on their accumulated experience or may just display reciprocity to an anonymous subject in a later repeated game. In the appendix we show the evolution of C-equilibria over all repeated games in all our 10 sessions. The figures are based on the data set D1400. In almost all treatments we do not see a clear trend. In those figures where there is a trend as in Session 5 it supports the direction of our predictions. Given these observations we believe we have good reasons to trust our statistical analysis.

Although our theoretical predictions are a little bit sharper for the filtered samples of observations we want to point out that none of our results actually depends on this filtering process. This means that the empirical validity of our axioms for non-equilibrium behavior cannot be rejected by our experimental results.²⁴

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 4 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.²⁵

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

²⁴This confirms Werner Güth's and Wolfgang Leininger's observation mentioned in section 5.

²⁵For 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	100
d	0	80

and

Γ_{34}	c	d
c	90	130
d	30	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 much 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 $\underline{\delta}$.²⁶

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	100
d	70	80

and

Γ_{47}	c	d
c	90	100
d	0	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	120
d	0	50

and

Γ_{60}	c	d
c	90	140
d	0	30

²⁶This 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 0.333 to 0.429 and then to 0.455, 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^* , contradicting predictions based on $\underline{\delta}$. This holds for all data sets D1400, D1272, and D1098.*

This result provides unambiguous support for our hypothesis (i). In section 7 we show that in all previous experimental studies that 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 criteria.

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 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 of C-equilibria frequencies dependent on the difference $(\delta - \delta^*)$ for D1272.²⁷

We can omit the according figures for D1400 and D1098 since the estimated parameters of the logistic function given in table 5 show that all figures look very similar to figure 1.

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^* .

²⁷The 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 ***
D1098	-0.620 ***	14.715 ***

Signif. codes: 0 ≤ *** < 0.001 ≤ ** < 0.01 ≤ * < 0.05 ≤ . < 0.1

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

7 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

Study	Game		Equilibrium				Round			
	$\underline{\delta}$	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 (2009), 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 potentially 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 (2009) 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 than 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 (2009) 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

- Aoyagi, M., G. Fréchette (2009), “Collusion as public monitoring becomes noisy: Experimental evidence”. *Journal of Economic Theory*, 144, pp 1135–65.
- 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.
- Kalai, E. and E.Lehrer (1993), “Rational Learning Leads to Nash Equilibrium”. *Econometrica*, 61(5), pp 63–80.
- 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.
- Stahl, D.O. (1991), “The Graph of Prisoner’s Dilemma Supergame Payoffs as a Function of the Discount Factor”. *Games and Economic Behavior*, 3, pp 368–384.
- Tirole, J. (1988), “The Theory of Industrial Organization”. Cambridge, Mass., M.I.T. Press.
- van Damme, E. (1989), “Renegotiation-proof Equilibria in Repeated Prisoner’s Dilemma”. *Journal of Economic Theory*, 47, pp 206–217.

A Appendix

The following figures show the evolution of C-equilibria in all sessions. The dotted blue lines are the average frequencies of C-equilibria reported in table 4 for D1400. The red lines separate the two different treatments within one session.




Figure A.1: $\delta = 0.75$, Game 1 and 2




Figure A.2: $\delta = 0.5$, Game 3 and 4




Figure A.3: $\delta = 0.75$, Game 3

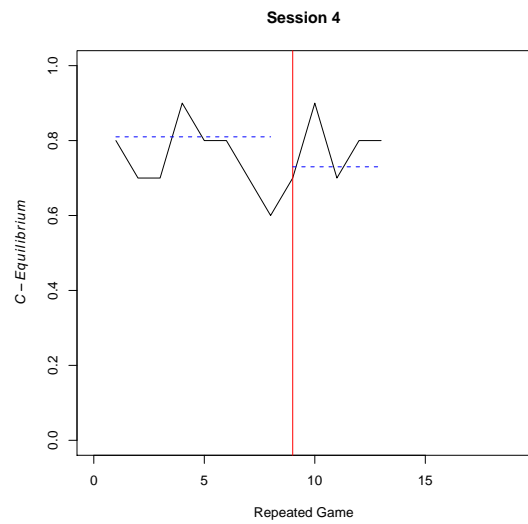


Figure A.4: $\delta = 0.875$, Game 3 and 4




Figure A.5: $\delta = 0.75$, Game 4 and 3




Figure A.6: $\delta = 0.75$, Game 2 and 1




Figure A.7: $\delta = 0.75$, Game 5 and 6




Figure A.8: $\delta = 0.875$, Game 4 and 3




Figure A.9: $\delta = 0.75$, Game 6 and 5




Figure A.10: $\delta = 0.75$, Game 3 and 4