



University of Trento

Doctoral School of Social Sciences

**Too much of anything is bad for you,  
even information:  
how information can be detrimental to  
cooperation and coordination**

Chiara D'Arcangelo

Submitted in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy in Economics and Management of the University of Trento.

October 2018



*If we had deliberately built, or were consciously shaping, the structure of human action, we would merely have to ask individuals why they had interacted with any particular structure.*

*Whereas, in fact, specialized students, even after generations of effort, find it exceedingly difficult to explain such matters, and cannot agree on what are the causes or what will be the effects of particular events. The curious task of economics is to demonstrate to men how little they really know about what they imagine they can design.*

(Hayek, 1988, pp.76-77)



## Abstract

Repeated games of cooperation share the same equilibrium selection problem as coordination games. In both settings, providing information might help players coordinating on efficient equilibria. Which equilibrium is most likely to be selected strictly depends on the type and the amount of information provided. It is then natural to ask under which conditions providing more information increases efficiency.

The present thesis makes a step in answering this question. It analyzes how the presence of information regarding either the opponent, or the options that are available for choice, might change players' behavior. It focuses on two settings where increasing information might be detrimental for players: a repeated Prisoner's dilemma, and a coordination game.

The first chapter develops a theoretical model in which players have limited information about the opponents' previous moves. When applied to the Trust Game, we show that by increasing the amount of information disclosed to the first player, more exploitative equilibria appear, in which that player obtains a smaller payoff. These equilibria disappear in settings in which the information the first player obtains about the second player's past behavior is limited. This is a case in which providing a player more information may reduce his payoff in equilibrium.

In the second chapter, we test this latter result with a laboratory experiment, and we show that subjects do understand that different behavior might be optimal in different settings. Subjects tend to use a fully cooperative strategy more often when only minimal information is available. Moreover, subjects trying to exploit the opponent succeeded in gaining more than the mutual cooperation payoff only when the information provided to the opponent is sufficiently rich, that is when our model predicts that exploitative outcomes are equilibria.

The last chapter considers the effects of introducing information about the options available for choice in a coordination game. It reports the results from a simulated crowdfunding experiment. We show that the presence of non payoff-relevant information is able to make a project focal. However, when returns from coordination are uncertain, the presence of information is instead detrimental for coordination.



# Contents

<b>List of Figures</b>	<b>ix</b>
<b>List of Tables</b>	<b>xi</b>
<b>Introduction</b>	<b>1</b>
<b>1 You don't fool me: Memory one strategies in repeated games</b>	<b>5</b>
1.1 Introduction . . . . .	5
1.2 Literature Review . . . . .	7
1.2.1 The repeated Prisoner's Dilemma: mutual cooperation . . . . .	8
1.2.2 The repeated Prisoner's Dilemma: Extortion . . . . .	11
1.2.3 Repeated games with memory-one strategies: the equilibria . . . . .	15
1.2.4 Reputation and the Trust Game . . . . .	17
1.3 Methods . . . . .	20
1.3.1 General setting . . . . .	20
1.3.2 Memory-one strategies . . . . .	23
1.3.3 Zero Determinant Strategies . . . . .	27
1.4 Equilibria in the PD . . . . .	34
1.4.1 Equilibria on the Pareto frontier . . . . .	34
1.4.2 Equilibria below the Pareto frontier . . . . .	36
1.4.3 Equilibrium strategies . . . . .	39

1.5	The Trust Game . . . . .	43
1.5.1	The repeated Trust Game . . . . .	43
1.5.2	Zero Determinant and Unconditional strategies . . . . .	45
1.5.3	Equilibrium payoffs . . . . .	47
1.5.4	Equilibrium strategies . . . . .	50
1.6	Short-run vs Long-run: does it pay to extort? . . . . .	54
1.6.1	Short-run and long-run players with memory-one . . . . .	54
1.6.2	Observable mixtures . . . . .	55
1.7	Discussion and Conclusions . . . . .	58
<b>Appendices</b>		<b>59</b>
1.A	Zero Determinant strategies with no discounting (Press and Dyson, 2012) . . . . .	59
1.B	Belief-free equilibria (Ely and Valimaki, 2002) . . . . .	61
1.C	Proofs . . . . .	64
1.D	Reactive strategies with no discounting (Baklanov, 2018) . . . . .	79
<b>2</b>	<b>Be nice, but not always: Extortion in the Lab</b>	<b>81</b>
2.1	Introduction . . . . .	82
2.2	Trust in the Lab . . . . .	84
2.3	Should you be nice? Equilibria in the Trust Game . . . . .	87
2.4	Extortion in the Lab . . . . .	90
2.4.1	Experimental Design . . . . .	90
2.4.2	Predictions . . . . .	92
2.5	The experiment . . . . .	94
2.5.1	Setting . . . . .	94
2.5.2	Results . . . . .	96
2.6	Conclusions . . . . .	106



<b>Appendices</b>	<b>108</b>
2.A More Figures . . . . .	108
2.B Instructions . . . . .	112
<b>3 You go first: Coordination in a simulated Crowdfunding Experiment</b>	<b>117</b>
3.1 Introduction . . . . .	118
3.2 Literature Review . . . . .	121
3.3 Experimental Design . . . . .	124
3.4 Predictions and Hypotheses . . . . .	130
3.4.1 Theoretical Framework . . . . .	130
3.4.2 Herding and Behavioral Predictions . . . . .	133
3.5 Results . . . . .	135
3.5.1 Winning Projects . . . . .	137
3.5.2 Coordination . . . . .	138
3.5.3 Choices over time . . . . .	143
3.5.4 The role of <i>cheap</i> information . . . . .	149
3.5.5 Strategies and choices . . . . .	152
3.6 General discussion and Conclusions . . . . .	153
<b>Appendices</b>	<b>155</b>
3.A Notes on Coordination . . . . .	155
3.B Projects . . . . .	157
3.C Instructions . . . . .	157
<b>Conclusions</b>	<b>167</b>
<b>Bibliography</b>	<b>169</b>



# List of Figures

1.1 ZD strategies - Examples . . . . .	12
1.2 Stackelberg payoff . . . . .	18
1.3 Feasible and rational payoffs - PD and TG . . . . .	22
1.4 TfT, Grim, and WSLS . . . . .	26
1.5 ZD strategies - PD . . . . .	28
1.6 Equalizers, Extortionate, and Generous strategies . . . . .	31
1.7 Equilibria in memory-one strategies - PD . . . . .	34
1.8 Nash equilibria in ZD strategies - PD . . . . .	40
1.9 Set of feasible payoffs - TG . . . . .	44
1.10 ZD strategies - TG . . . . .	45
1.11 Equilibria in memory-one strategies - TG . . . . .	48
1.12 Nash equilibria in ZD strategies - TG . . . . .	51
1.13 Mixed SPNE: Maximum payoff for the second player - TG . . . . .	53
1.14 Equilibria with observable mixtures - TG . . . . .	57
1.B.1 Belief-free equilibria - PD . . . . .	63
2.1 The Trust Game . . . . .	82
2.1 Feasible payoff profiles in the repeated TG . . . . .	88
2.1 Nash equilibria - Memory one vs Observable mixtures . . . . .	93
2.1 Results . . . . .	96

2.2 <i>f</i> -Reward strategies . . . . .	102
2.3 <i>f</i> -Reward strategies: Trust and Payoffs . . . . .	103
2.4 Information - Trust and Reward . . . . .	105
2.A.1 Average Reward I . . . . .	108
2.A.2 Average Reward II . . . . .	108
2.A.3 Always Reward . . . . .	109
2.A.4 Almost Reward . . . . .	109
2.A.5 Mild exploiters . . . . .	110
2.A.6 Exploiters . . . . .	110
2.A.7 Hard exploiters . . . . .	111
2.A.8 Almost Not Reward . . . . .	111
2.B.2 Choice window - First player . . . . .	115
2.B.3 Choice window - Second player . . . . .	115
3.1 Funding over time by Market Session . . . . .	143
3.2 Funding over time by Treatment, First Market Session . . . . .	144
3.3 Funding over time: Italy and Finland . . . . .	145
3.4 Funding over time by Treatment . . . . .	146
3.5 Clicks and Choices . . . . .	149
3.C.1 Example of a project . . . . .	159
3.C.2 Choice Window - Baseline . . . . .	161
3.C.3 Choice Window - Info . . . . .	161
3.C.4 Choice Window - Uncertainty . . . . .	162
3.C.5 Choice Window - Combined . . . . .	162

# List of Tables

1.1 Stage game payoffs in the PD . . . . .	8
1.2 Stage game payoffs in the TG . . . . .	17
1.3 Stage game payoffs . . . . .	21
1.4 Normal form - Example . . . . .	33
1.5 Stage game payoffs in the TG . . . . .	43
2.1 Stage game payoffs in the TG . . . . .	87
2.1 Payoffs in the experiment . . . . .	91
2.1 Actions and Outcomes - Averages . . . . .	98
2.2 Outcomes over time . . . . .	99
2.3 Actions and outcomes - Rewarding types . . . . .	100
2.4 Ratios - $r_1, r_2, r_3$ . . . . .	101
2.5 Strategies - Choices over time . . . . .	104
3.1 Baseline and Treatments . . . . .	125
3.1 Sample demographics . . . . .	136
3.2 Winning projects . . . . .	137
3.3 Coordination measure . . . . .	138
3.4 Determinants of choosing the successful project . . . . .	139
3.5 Coordination measure, Italy . . . . .	140
3.6 Determinants of choosing the successful project, Italy and Finland . . . . .	141

3.7	Coordination measure, Finland . . . . .	141
3.8	First movers . . . . .	146
3.9	Contributions over time - Winning projects . . . . .	147
3.10	Determinants of choosing the successful project: learning . . . . .	148
3.11	Debriefing question: Decisions . . . . .	151
3.12	Debriefing question: Information . . . . .	151
3.13	Debriefing question: Strategies . . . . .	152
3.B.1	Selected projects . . . . .	157

# Introduction

*In an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.*

(Simon, 1971, pp.40)

The presence of information can dramatically change the outcome of a strategic interaction. In settings characterized by multiple equilibria, introducing information might help solving the equilibrium selection problem. In other settings, characterized by inefficient outcomes, introducing information might allow players to reach more efficient outcomes, but, at the same time, it might create an equilibrium selection problem.

Which equilibrium will be selected crucially depends on the *type* and the *amount* of information provided. In order for information to be beneficial, it is necessary that different individuals are able to interpret it in the same way. This process might be difficult in situations characterized by abundance of information, which can be harder to interpret (Simon, 1971). Indeed, there are cases in which the presence of *little* information can be beneficial, while the presence of *too much* information can be detrimental.

This thesis tries to explore both possibilities. The main question it tries to answer is to what extent introducing information is beneficial for players, in the sense of allowing players to reach outcomes that are more favorable to them.

The first part of the thesis considers settings characterized by inefficient equilibria, where information is on players' past behavior. Specifically, it deals with the Prisoner's Dilemma (PD) and the Trust Game (TG), where the only equilibrium is mutual defection if no information about past behavior is available. Introducing information allow players to reach more efficient equilibria, but it also creates an equilibrium selection problem.

The second part of the thesis considers settings characterized by multiple equilibria, where there can be public information exogenously given to the players. Specifically, it deals with a multiplayer coordination game, where different pieces of information are given over the different available options. Introducing information might help players in solving the equilibrium selection problem, but only if it is able to make one option focal for the majority of them.

The main finding of this thesis is that, perhaps counterintuitively, increasing the amount of information available doesn't imply a higher rate of cooperation or coordination. In both cases, there are conditions under which players could be better off by knowing less, meaning that they are able to reach a higher payoff when the disclosed information is *limited*.

Chapter One deals with repeated interactions between players with limited memory.

The first part of the chapter focuses on the PD, with the aim of describing equilibria when players are restricted to strategies that only condition on the previous outcome (so called *memory-one* strategies). Specifically, it considers whether equilibria are possible where one player is able to gain more than the mutual cooperation payoff, that is an equilibrium that is not Pareto-dominated by mutual cooperation (we call it an *extortionate* equilibrium). This question is interesting because, if the answer is negative, then the best that players can do is mutual cooperation. Indeed, this part of the thesis shows that, when players are restricted to strategies that can only condition on what happened in the previous period, the only efficient equilibrium is mutual cooperation. Moreover, in *any* equilibrium, players' payoffs are bounded from above by the mutual cooperation payoff. Thus, no extortionate equilibria exist if players are restricted to memory one strategies. Comparing this result with the case where players can use strategies that condition on more than one period, the implication is that, by restricting player's memory (thus reducing the information available) it is possible to insure players against the risk of falling into equilibria where the opponent *extorts* a higher payoff.

The second part of the chapter applies the previous findings to a TG played between a long-run second player against a sequence of short-run first players. In those situations we say that the



second player is able to build a *reputation*, which is represented by the information available to the first players regarding his behavior in previous interactions. The amount of information given to first players will then determine the type of reputation that the second player can build. This part of the thesis shows that, when the information is only relative the last period, the best reputation the second player can build is fully cooperative. When instead the information is relative to more than one period, extortionate equilibria exist, and the best reputation would imply a less-than fully cooperative strategy for the long-run player.

This result is interesting because it shows that, in situations where information is about a player's previous behavior, having little information could be beneficial. If the model presented in Chapter One is correct, people should be more willing to fully cooperate when only minimal information is available, compared to settings in which richer information is considered.

This prediction is tested with a laboratory experiment in Chapter Two. The experiment involves a repeated TG, and is designed to test whether different amount of information available to the *first* player would trigger different strategies choices by the *second* player. Overall, the results from the experiment are in line with the model's predictions. On the one hand, subjects in the experiment do try to form a fully cooperative reputation more often when less information is available. On the other hand, subjects who tried to build a less-than-fully cooperative reputation succeeded in gaining more than the mutual cooperation payoff only when more information was available. Those results suggest that subjects in the experiment are able to recognize the different settings implied by the different amount of information. They are ready to cooperate more when mutual cooperation is the only efficient equilibrium, but they are also ready to extort more when also extortionate outcomes can be sustained in equilibrium.

The results from Chapter Two may imply that the low level of cooperation observed in the experimental literature on repeated games can be due to subjects trying to coordinate on different equilibria. When the amount of available information on previous behavior is large, many new equilibria appear, and a repeated PD might resemble a coordination game.

In coordination games, introducing public information exogenously given may help players to solve the equilibrium selection problem. The large literature on salience demonstrates that it is easy to foresee players' decisions by means of focal points. While this is true in simple settings in which it is relatively easy to understand which option is focal, it is still an open question whether the presence of richer information would still be able to achieve the same result.

To explore this possibility, Chapter Three considers a coordination experiment designed to simulate a crowdfunding platform, where different *levels* of information are disclosed to the players. In crowdfunding, potential investors have access only to the information that the designers of the projects decide to disclose, and it is necessary that enough people invest on the same project in order to finance it. Information about previous behavior is then extremely valuable, and there might be an incentive to wait and see what the others will do, in order to invest in a project only when it has a higher chance to be fully funded. If everybody waits for the others, nobody will ever chose anything, not solving the coordination problem. The presence of exogenous information in this case may help players to focus on the same project. Indeed, the results from the experiment suggest that information did play a role in driving subjects' decisions. Non payoff-relevant information is effective in making a project focal, and subjects who use the information available are more likely to choose the winning project. However, when returns from the projects are uncertain, the beneficial effect of information disappears, as subjects could reach a lower level of coordination compared to the case where no information was provided. Those results imply that, in the presence of uncertainty, subjects might be better off by not having additional information.

This thesis tries to explore the relation between the availability of information and equilibrium selection. It considers two specific cases, cooperative and coordination games, where more information could be thought to increase efficiency. However, the present results show that there are conditions under which the opposite is true. In those cases, increasing the amount of information disclosed to players has a deleterious effect, in the sense that those players might have been better off by knowing less. This is relevant when considering the design of mechanisms able to foster cooperation or coordination, as it suggests that explicit incentives may be needed, and extreme care should be put when deciding which information to disclose.

# Chapter 1

## You don't fool me: Memory one strategies in repeated games

### 1.1 Introduction

The theory of repeated games has been widely used to explain the emergence of cooperation in games like the Prisoner's Dilemma (PD), in which all Nash equilibria (NE) are Pareto inefficient (see Nowak (2006) for a comprehensive review of this literature). When a game is repeated over time, a large number of new equilibria appear. A long known results in game theory, the so called Folk Theorem, states that when a game is repeated, and players are sufficiently patient, all outcomes that Pareto dominate inefficient NE of the stage game can be sustained as subgame perfect NE<sup>1</sup> (SPNE). This result is reassuring, as it implies that, with a sufficiently long time horizon, players may sustain efficient mutual cooperation also in games, like the PD, in which all one shot equilibria are inefficient. However, the Folk Theorem also implies that, beside cooperative equilibria, there are many others that are less appealing. In some of them, cooperation is never observed. For example, no matter how patient players may be, mutual defection is always an equilibrium. In others, one player cooperates more than the other. For example, there may be equilibria in which one player alternates between cooperation and defection, while the other player cooperates continuously. In these equilibria we say that one of the players "extorts" a larger payoff from the other.

---

<sup>1</sup>The Folk theorem is in fact more general than this. See Mailath and Samuelson (2006) for a textbook presentation of the results in this literature.

Given the wealth of equilibria created by repetition, it is not surprising that a large literature formed on the equilibrium selection problem. For a long time, attention has been focused on the strategies that are able to sustain mutual cooperation (Imhof et al., 2007). The main concern was the type of punishment which is more effective in maintaining cooperation, and the relative stability of different (cooperative and non-cooperative) equilibria. Scholars investigated whether cooperation is easier to sustain by popular strategies like Tit-for-Tat, Pavlov, or Grim, and whether it is easier to move from cooperative to non-cooperative equilibria or *vice versa*.

More recently, in a seminal paper, Press and Dyson (2012) focused on the second class of equilibria, those in which one player cooperates less often than the other, and hence enjoys a larger payoff. They consider a repeated PD where players are constrained to use memory-one strategies, i.e. strategies that only condition on the last outcome of the game, and they singled out a class of strategies, which they dub *extortionate*. A player using an extortionate strategy is able to ensure that his payoff is never below the one of his opponent. In fact, by playing against an extortionate strategy, a player cannot do better than always cooperate, allowing the opponent to reap a larger part of the benefits of cooperation, and to enjoy a payoff that is larger than the mutual cooperation payoff. The relevance of extortionate strategies for the emergence of cooperation in repeated games is at the moment hotly debated (see, for example, Hilbe, Nowak, and Traulsen (2013) and Stewart and Plotkin (2012)).

We contribute to this literature in two ways. First, we characterize the full set of equilibrium strategies and payoffs when players are restricted to memory-one strategies, investigating whether extortionate strategies may be part of any NE. This is important in order to clarify the role that this type of strategies may have in the emergence of cooperation. If no extortionate strategy is part of a NE, then no such strategy can pass more stringent tests like evolutionarily stability. Our main result shows that, as long as players are constrained to use memory-one strategies, in any NE no player can get more than the mutual cooperation payoff. In other words, extortion in repeated games is only possible when at least one player conditions his current choice on more than just the last outcome of the game. The intuition behind this result is straightforward. To get more than the mutual cooperation payoff, a player should randomize between cooperation and defection. However, when the opponent only conditions to the last outcome of the game, no deviation from a mixed strategy can be observed and hence punished, and strategies that cannot punish deviations cannot be NE of the repeated game, unless they are a NE of the stage game.

As we shall see in a while (see below Section 1.2), it is a common observation that extortionate strategies are unlikely to succeed in symmetric contexts like, for example, a PD played by individuals belonging to a single population. In a symmetric context, if each player tries to get more than his opponent, the only possible outcome can be mutual defection. Extortion may have a better chance to play a role in asymmetric situations in which players have different strategic opportunities, or, in evolutionary contexts, when the game is played by two different populations. Our second contribution points in this direction. We study extortion in the Trust Game (TG), which is an asymmetric version of the PD in which one player chooses first whether to cooperate or not and, if he chooses to cooperate, his opponent decides whether to reciprocate or not. We extend the analysis of memory-one strategies to the TG and conclude that, just like in a PD, in no NE one player can get more than her mutual cooperation payoff.

Motivated by this result, we further extend the analysis to the case in which one player (the trustor in the TG) can condition his choice on the *frequency* with which the opponent cooperated in the previous interactions. We do a first step in this direction by studying equilibria in the repeated TG with observable mixtures. That is, the first player observes the probability with which the second player has chosen to cooperate in the past. We prove that with observable mixtures, there is a continuum of extortionate equilibria in which the second player gets more than what he could get by mutual cooperation. Finally, we explore the analogies of our results with the classical literature opened by Fudenberg and Levine (1989), in which a single long-run player interacts with a population of short-run players.

This chapter is organized as follows. Section 1.2 revises the existing literature. Section 1.3 presents the main methods used to study memory-one strategies. Section 1.4 characterizes equilibrium strategies and payoffs in the Prisoner's Dilemma. Section 1.5 does the same for the Trust Game. Section 1.6 considers the case of a long-run player facing a population of short-run players, and Section 1.7 concludes.

## 1.2 Literature Review

**Definitions and notation** Before moving to the review of the literature, we need some definitions. Consider a standard repeated PD where, in each round, players have to choose whether to Cooperate (C) or Defect (D). The stage game payoffs are shown in Table 1.1:  $R$  is the mutual cooperation payoff,  $P$  is the mutual defection payoff,  $T$  is the temptation payoff,

and  $S$  is the sucker payoff, with  $T > R > P > S$  and  $2R > S + T$ . A PD has *equal gains from switching* if  $T - R = P - S$ . In each round, there is a probability  $\delta \leq 1$  of going to the next round. Notice that  $\delta$  is sometimes interpreted as the discount factor of future players' payoffs.

	$C$	$D$
$C$	$R, R$	$S, T$
$D$	$T, S$	$P, P$

Table 1.1: Stage game payoffs in the PD

The history of the game up to period  $t$  is the sequence of action profiles chosen by the two players until period  $t - 1$ , and a strategy for the repeated game has to specify an action for each possible history. A strategy is memory-one if it prescribes the same behavior after any history that has the same outcome in the previous period. In other words, a memory-one strategy only conditions on what happened in the previous period. A general memory-one strategy  $\mathbf{p}$  is thus fully characterized by 5 probabilities:

$$\mathbf{p} = (p_0, p_{CC}, p_{CD}, p_{DC}, p_{DD})$$

where  $p_0$  is the probability to cooperate in the first round (i.e. when no history is available), and  $p_w$  is the probability to cooperate after each of the four possible outcomes  $w$ , that is after each of the four possible action profiles that can be chosen in each period. A memory-one strategy is *reactive* if it only conditions on what the opponent did in the previous round. Formally  $p_{CC} = p_{DC} = p_C$  and  $p_{DD} = p_{CD} = p_D$ . Finally, a reactive strategy is *unconditional* if it specifies the same probability to cooperate after any possible history:  $p_C = p_D = p_0$ .

### 1.2.1 The repeated Prisoner's Dilemma: mutual cooperation

One of the first systematic attempts to find the “best” strategy for the repeated PD was the famous tournament organized by Axelrod in the late 70s (see Axelrod (1980a) and Axelrod (1980b)). He considered a standard repeated PD, and asked participants to submit strategies, in the form of computer programs, to play it<sup>2</sup>. A strategy defined a move for the first round, and a move for each subsequent round, and could condition on all the previous history of the game. From the results of the tournaments, Axelrod and Hamilton (1981) identified 3 properties that

---

<sup>2</sup>In the first tournament, there were 14 different strategies submitted by game theory experts, while in the second there were 62 strategies submitted by experts from various disciplines and non-experts computer hobbyists.

best performing strategies had in common: they were never the first to defect (niceness), they were ready to respond if the opponent defected (provocability), but they were also ready to go back and cooperate after the punishment (forgiveness).

Even if some of the submitted strategies were rather complicated, the winning strategy of both tournaments was the rather simple Tit for Tat (TfT). TfT starts with cooperation and then plays the same action played by the opponent in the previous round. TfT is thus nice, provokable and forgiving. Indeed, TfT is able both to get the mutual cooperation payoff when matched with another nice strategy, and to avoid being exploited by more defecting strategies. Moreover, TfT is a symmetric NE, and, as Axelrod (1984) showed, it is *collectively* stable, that is, no strategy have a selective advantage over it (in a more familiar terminology, TfT is *neutrally stable*). Nonetheless, TfT has some flaws. First, it is a SPNE only in games with equal gain from switching, and only for certain values of the discount factor (Kalai et al., 1988). Second, in the presence of noise, it might enter into cycles of  $(C, D)$  and  $(D, C)$ , implying that it gets the same payoff of a strategy that plays randomly (Selten and Hammerstein, 1984). Moreover, already Axelrod (1980b) noticed that the relative success of TfT in the tournament was largely determined by the particular composition of the competing population.

When considering the full strategy set, although TfT is “collectively” stable as in Axelrod’s definition, it is not evolutionarily stable. Strategies like All Cooperation (AllC), (as well as any other *nice* strategy, that is a strategy that starts by cooperating and keeps cooperating with probability one after mutual cooperation), are able to neutrally invade a population of TfT players. Selten and Hammerstein (1984) proved that in the repeated PD there are no (pure) strategies that are evolutionarily stable. As such, TfT is only one of the many neutrally stable strategies of the iterated PD (Bendor and Swistak, 1995).

More recently, Garcia and Van Veelen (2016) proposed an analysis of the repeated PD using a stability criterion which is intermediate between neutral stability and evolutionarily stability. Van Veelen (2012) calls a strategy  $s$  Resistant Against Indirect Invasion (RAII) if there is no chain of neutral mutants through which  $s$  can be invaded by another strategy that does have a selective advantage. TfT fails the RAI criterion, as it can be neutrally invaded by AllC, which in turn can be invaded by AllD. Garcia and Van Veelen (2016) prove that no strategy passes the RAI test in a large class of games, including the PD.

A different approach to select among equilibria is to relax some of the assumptions of the Folk Theorem, thus restricting players’ strategy sets. The Folk theorem assumes that players have

perfect recall and can condition their choices over the entire history of the game. Notice that this implicitly assumes that players have an implausible amount of memory to keep track of all possible contingencies of the game. This assumption is particularly questionable in evolutionary settings in which the emphasis is on bounded rationality and learning. A natural question is then whether the Folk theorem continues to hold when restrictions are put on player's memory (Compte and Postlewaite, 2015; Piccione, 2002).

A first step in this direction is to assume that memory is bounded to the immediate previous round, i.e. that strategies are memory-one. Building on the results of Axelrod's tournament, Nowak and Sigmund (1988) and Nowak and Sigmund (1989) extensively study how the probabilities that define stochastic reactive strategies change under the replicator dynamics. They prove that, in the presence of noise, the dynamics leads either towards AllD, or towards a state where cooperation is sustained by a strategy that is more forgiving than Tft.

Nowak (1990), found conditions for reactive strategies to be NE in a repeated PD with no discounting. He noticed that, with reactive strategies and no discounting, "in each round the cooperation of the first and the second player is independent" (p.95). This greatly simplifies the analysis, and allows to prove that also in this setting there are no evolutionarily stable strategies, although there are several, cooperative and non-cooperative Nash equilibria. Nowak and Sigmund (1990) extended this analysis to the case of discounting. Discounting complicates the matter and leads the authors to characterize Nash equilibria only for a restricted set of strategies and for the subclass of PD with equal gains from switching.

Nowak and Sigmund (1995), further extended the analysis to general memory-one strategies, considering also an alternating version of the iterated PD. Using simulations, they show that payoffs are usually close to either the mutual cooperation or the mutual defection payoff. Moreover, if the mutual cooperation payoff is sufficiently large, the most frequent strategy is Pavlov.

More recently, also thanks to the advance of computing and simulations programs, several papers tried to further analyze the behavior of memory-one strategies in the iterated PD, with great emphasis on strategies' performance in one or two populations, and little regards to whether strategies could be part of a Nash equilibrium. The main focus is still on the performance of Tft, compared with other well known strategies such as AllC, AllD, Pavlov.

For example, Zagorsky et al. (2013) considered all possible strategies described by one and two state automata (that is, pure memory-one strategies), in the alternating PD with equal gains



from switching. In their setting, players can make mistakes, but they cannot observe their own errors. Thus, it is not surprising that Tft performed poorly: AllD and Grim are the only possible strict Nash equilibria (depending on the error rates and the value of the benefit of cooperation). Nonetheless, convergence to one of those equilibria was always rare. When the payoff from mutual cooperation is sufficiently large, the most likely outcome is a convergence to a mixed equilibrium where the prevalent strategy is Forgiver, that is a strategy that punish defections but attempts to re-establish cooperation even after multiple defections.

Baek et al. (2016) extended the analysis by considering mixed strategies in games with equal gains from switching. They call the gain from switching the cost of cooperation. They show that, for high cost of cooperation, cooperation rates are generally low, but relatively higher for reactive pure strategies. For low cost of cooperation, cooperation rates with memory-one strategies are always higher than with reactive strategies. In this case, mixed strategies increase the rate of cooperation, but only for reactive strategies. Interestingly, for intermediate values of cooperation costs, the dynamic leads to the (somehow strange) strategy which only cooperates after begin exploited, that is the exact opposite of a forgiving strategy.

Thus, even if we have some insights on how memory-one strategies behave, the search for the “best” strategy is still ongoing.

## 1.2.2 The repeated Prisoner’s Dilemma: Extortion

### Zero-determinant strategies: definitions

While the strategies considered in the previous sections are well-known, Press and Dyson (2012) focused on a new class of memory-one strategies for the repeated PD, which they called Zero Determinant (ZD). They noticed that, when playing against each other, two memory-one strategies induce a Markov chain whose invariant distribution can be explicitly calculated. By manipulating the transition matrix, they discover that there are particular strategies that are able to enforce a linear relation between players’ payoffs. A player using such a strategy (i.e. a ZD strategy) is able to fix the slope and the intercept of the line on which his own and the opponent’s payoff will lie in the payoff space, as if he was using an unconditional strategy. Contrary to unconditional strategies, that in the PD all have a negative slope, ZD strategies can have a positive (or null) slope, meaning that the best reply to a ZD strategy can also maximize the payoff of the ZD player.

Figure 1.1 shows some examples of ZD strategies. The kite-shaped area represents the set of feasible payoff profiles for the repeated PD. The second player (on the x-axes) can choose the probabilities of his memory-one strategy so that, regardless of the strategy played by his opponent, their payoff profiles will lie on a straight line. The linear set of payoff profiles associated to a ZD strategy played by the second player is represented by the set of black dots. Each dot in Figure 1.1 represents the payoff profile associated to a fixed ZD strategy of the second player, and one randomly generated memory-one strategy of the first player.

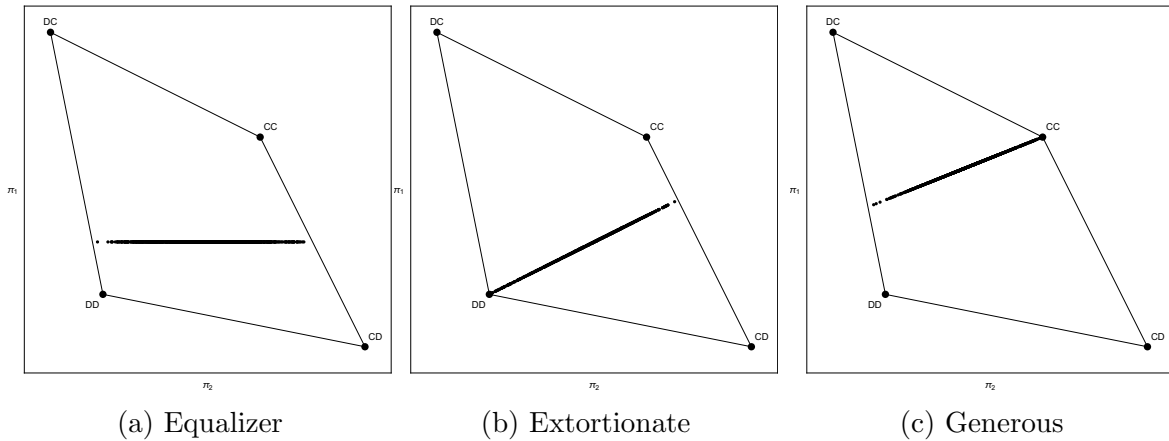


Figure 1.1: Examples of Zero Determinant Strategies for the second player

Among all ZD strategies, Press and Dyson (2012) focused on those which they called *extortionate* (ExS, see Figure 1.1b). By using an ExS, a player can be sure that his payoff is never below the one of the opponent. Moreover, against an ExS opponent, a player cannot do better than always cooperate. The resulting payoff profile is then a point on the Pareto frontier in which the player who uses the ExS reaps a larger share of the benefits of cooperation, and is able to get a payoff that is larger than the mutual cooperation payoff.

Another important class of ZD strategies, that will play a crucial role in our theoretical results, is made by all strategies that are able to fix the opponent's payoff at a certain level, thus forcing the payoff profiles to lie on an horizontal (or vertical, depending on the player) line. The existence of this kind of strategies was already noticed by Boerlijst et al. (1997), who called them *Equalizer* (EqS, see Figure 1.1a). In this case, no matter which strategy the first player chooses, his payoff will only depend upon the EqS strategy chosen by the second player.

EqS are less responsive than ExS to the moves of the opponent. Typically, EqS have low probability to change actions: starting with cooperation (defection), they have an high probability to keep cooperating (defecting). ExS instead will always defect after mutual defection, but

will cooperate with a positive probability (smaller than one) in any other case. ExS are more responsive to the behavior of the opponent, and will cooperate often enough so as to make AllC a best reply, but will defect often enough to ensure for themselves a payoff that can be higher than the mutual cooperation payoff. Moreover, when facing another ExS or AllD, an ExS gets the mutual defection payoff. As such, ExS are “unbeatable” (Hilbe, Röhl, et al., 2014) that is, in direct competition, no strategy exists that gains against an ExS more than the ExS itself.

Stewart and Plotkin (2012) introduced a further type of ZD strategy, which they call *generous* (GeS, see Figure 1.1c). A player using a generous strategy can ensure his payoff to never be above the one of the opponent. Typically, generous strategies cooperate with probability one after mutual cooperation, and, like ExS, they reward cooperation by cooperating more often after a cooperation than after a defection. When facing each other, GeS are able to get the mutual cooperation payoff, but they are susceptible to be exploited by AllD.

Finally, notice that, in discussing their results, Press and Dyson (2012) suggest that a player facing a memory-one opponent cannot benefit from having a longer memory, as for any of his longer memory strategies there is a memory-one strategy able to get the same payoff<sup>3</sup>.

### **The evolution of extortion: the symmetric case**

ExS are never Nash equilibria with themselves, as the best reply to any of them is unconditional cooperation. It is not surprising that they usually perform poorly in evolutionary settings. Already Press and Dyson (2012) admitted that ExS could succeed only against a player who is able to “learn” to play the best reply to any strategy played by her opponent. Otherwise, in standard symmetric settings, more generous strategies perform better than ExS. This intuitively plausible result was first confirmed by Stewart and Plotkin (2012). They re-run Axelrod’s tournament by adding one ExS and one GeS. Within their setting, ExS performed worse than any other strategy, except AllD, while GeS performed better than all the other strategies, including TFT.

The success of GeS and the defeat of ExS was later confirmed by Hilbe, Nowak, and Traulsen (2013). They derive the adaptive dynamics for all ZD strategies, and study whether a finite

---

<sup>3</sup>This conclusion was somehow challenged by Li and Kendall (2014), who show that “longer memory strategies outperform shorter memory strategies statistically in the sense of evolutionary stability” (p.819) and by C. Lee et al. (2015), who constructed a longer memory strategy that is able to invade all possible memory-one strategies.

population would move towards a cooperative or a non-cooperative state, when varying the population size. Even if the population never converges to a state dominated by ExS, they discovered an unexpected role that this type of strategy can play. For large populations ( $N > 10$ ), the presence of ExS helps the population to escape from the AllD state. The reason for this is that ExS can neutrally invade AllD, thus allowing for the emergence of more cooperative strategies. For example, in a population of GeS, AllC can neutrally invade, opening the door to defectors, and ExS are needed to escape from the resulting state of mutual defection<sup>4</sup>.

The important role of ExS as facilitators of the transition from defection to cooperation is robust to several modifications of the general setting. For example, it holds in an the alternating PD (McAvoy and Hauert, 2017), in PD with continuous action spaces (McAvoy and Hauert, 2016), in PD with structured populations (Szolnoki and Perc, 2014a; Szolnoki and Perc, 2014b), in multiplayer games (Hilbe, Wu, et al., 2015) and in the presence of noise (Hao et al., 2015).

Finally, ExS strategies perform better if they are able to recognize themselves and to cooperate with each other. For example, Adami and Hintze (2013) show that ExS can be evolutionarily stable if players can use a “tag based strategy”, that is if they can send a reliable signal of the type of strategy they are using (Garcia, Van Veelen, and Traulsen (2014) explore a similar idea). A similar result would hold if ExS can infer the type of strategy used by their opponent by observing its past play. Notice however that this would require more than one period memory.

### **The evolution of extortion: asymmetric contexts**

From what we said in the previous section, it should be clear that to give extortionate strategies the best chance to succeed one should consider asymmetric contexts in which players have different strategic possibilities, different constraints on memory, or evolve at different rates. The literature has broadly confirmed this intuition. Hilbe, Nowak, and Sigmund (2013), for example, considered a repeated PD in a two populations setting in which the two populations evolve at different rates. In those conditions, in the long-run the slowest population will evolve towards ExS, forcing the fastest adapting population to fully cooperate, and enjoying a higher payoff. Hence, in asymmetric contexts, it pays to be the slower to adapt, an instance of the

---

<sup>4</sup>As a side note, the authors noticed that their results were similar to the ones obtained for reactive strategies by Nowak and Sigmund (1995). This is not a surprise, as they both consider games with equal gains from switching, where every ZD strategy is reactive (Hilbe, Nowak, and Sigmund, 2013). In both cases, there is a region they call *cooperation rewarding* zone, where every strategy with a higher probability to reward cooperation can invade. In the area outside this region the dynamics leads to AllD.

so-called Red King effect (see Bergstrom and Lachmann, 2003).

J. Chen and Zinger (2014) discuss a model in which a memory-one strategy is matched with an adaptive agent who learns to play a best response to any strategy played by the opponent. They show that for every ExS or GeS, every possible adaptation path will give the maximum possible payoff to both players. If matched against an ExS, the adaptive player will obtain a smaller payoff, which confirms the original intuition of Press and Dyson (2012).

ExS might succeed also in settings where individuals in a population have different learning abilities (Chatterjee et al., 2012), or when they have asymmetric incentives, either because they have different stage game payoffs, or because they have different future incentives due to a different time discount factors. Our main concern is with the latter case, when players have asymmetric incentives. We will show that having different stage game payoffs is not sufficient for extortion to emerge in the realm of memory-one strategy. In line with previous results, we shall prove that the less patient player needs to have a longer memory, which is needed to detect the use by the other player of a mixed strategy. If this is the case, the less patient player can punish the opponent if he tries to extort too much. Otherwise, if mixed strategy are not observable, as it is the case with memory-one strategies, no extortion is possible in equilibrium.

### 1.2.3 Repeated games with memory-one strategies: the equilibria

Aumann (1981) suggested that putting restrictions on players' memory would reduce the set of payoff profiles that can be sustained in equilibrium (see also Barlo et al., 2009). In this vein, a literature emerged that studies the existence of equilibria when memory is restricted to the last period. The most general theoretical result was originally proven by Ely and Valimaki (2002) as a preliminary step for the study of games with imperfect monitoring. They construct equilibria that are called "belief-free" (Ely, Hörner, et al., 2005). In these equilibria, each player is indifferent between his actions at each round, no matter which action the opponent is currently playing. In a PD, for example, at each round both players are indifferent between cooperation and defection, and they would remain indifferent even if they learned which action the other player is currently choosing. Their main result is that there is no belief-free equilibrium in which one player obtains more than the mutual cooperation payoff. We will generalize this result and show that in *any* NE and SPNE in memory one strategies (not necessarily belief-free), players' payoffs are never above the mutual cooperation payoff,  $R$ .

In a related study, Dutta and Siconolfi (2010) deal with general games, not necessarily PD. They prove that totally mixed belief-free equilibria in memory-one strategies exist in all games in which payoffs satisfy a condition they dub *reverse dominance*. These are games in which each player strictly prefers that the opponent plays one of his two actions, no matter which action he intends to use. The PD satisfies reverse dominance, as, no matter what a player intends to do, he prefers the other to cooperate. Dutta and Siconolfi (2010) prove that only payoff profiles that are between players' second and third best payoff can be sustained in a totally mixed belief-free equilibrium. For the Prisoner's Dilemma this is the same result proven by Ely and Valimaki (2002).

These theoretical results are clearly relevant for the literature on extortion in memory-one strategies, as they prove that in belief-free equilibria no player can get more than the mutual cooperation payoff. Yet, so far very few attempts have been made to connect the two. An exception is Baklanov (2018), who extensively studies equilibria in reactive strategies in games without discounting. He shows that, when players are restricted to use reactive strategies, and there is no discounting, a mild form of extortion can take place in equilibrium. That is, equilibria are possible in which one player obtains more than the mutual cooperation payoff. Extortion is mild, as the payoff that accrues to the extortionate strategy<sup>5</sup> is bounded away from the Pareto frontier. Indeed, in this setting an ExS player can only obtain slightly more than what he could obtain from mutual cooperation (see Appendix 1.D).

Barlo et al. (2009) proves a results which is crucial for the second part of our analysis. They prove that, in the context of repeated games with observable mixtures, the folk theorem is reestablished with memory-one strategies. In other words, when players can observe the probability with which each pure action is played at each round, any payoff profile that Pareto dominates the mutual defection payoff can be sustained in equilibrium<sup>6</sup>. This is clearly due to the fact that when mixtures are observable, deviations from the mixed strategy can be punished, just like deviation from pure strategies are. We shall investigate the consequences of this result in our analysis of the Trust Game, to which the following section is devoted.

---

<sup>5</sup>Notice that strategies considered by Baklanov (2018) are ExS, in the sense that a player's payoff is never below the one of the opponent, but they are not linear. See Appendix 1.D.

<sup>6</sup>Their result is in fact much more general than this, as it applies to any game, like Cournot and Bertrand duopoly, in which players have a strategy set which is compact.

## 1.2.4 Reputation and the Trust Game

In Section 1.2.2 we noticed that ExS are more likely to succeed in asymmetric settings in which players have different strategy sets or different time horizons. We explore this possibility in Section 1.5, by studying ExS in the asymmetric version of the PD in which one player chooses first, the so called Trust Game (TG). In the TG, the first player chooses whether to Trust (T) or Not Trust (NT) the second player. If he chooses T, then the second player decides whether to Reward (R) or Not Reward (NR) the first player's Trust. The normal form of the game as well as the stage game payoffs are shown in Table 1.2, with  $T > R > P > S$  and  $2R > T + P$ .

	R	NR
T	R,R,	S,T
NT	P,P	P,P

Table 1.2: Stage game payoffs in the TG

We shall consider a repeated version of the TG where the first player is able to observe the choice made by the second player even when he plays NT<sup>7</sup>. This version of the Trust Game is a special case of the *quality choice* game studied by Kreps and Wilson (1982) and Mailath and Samuelson (2006), in which the first player is a buyer who has to choose whether to buy (T) or not (NT) a product, while the second player chooses to deliver a high (R) or low (NR) quality product. In a simplified analysis of the repeated version of this game, it is convenient to assume that the buyer can observe the quality of the good, even if he decides not to buy it. Figure 1.2 represents the set of feasible payoffs for the repeated TG. The shaded area is the set of payoff profiles that can be supported in a NE when players are sufficiently patient. Although there is a superficial similarity between the TG and the PD, the existing literature points to a very important difference between the two games, when their repeated version is considered. This difference will prove to be very important in our discussion of extortion in the TG.

To illustrate how the TG differs from the PD, we shall discuss the standard reputation model due to Kreps and Wilson (1982) in which a single long-run second player interacts with a large number of short-run first players. For ease of exposition, we shall refer to the short-run player as the "customer" and to the long-run player as the "store". The type of situation we have in mind is the one in which a large number of customers are served by the same retail store: each customer only plays once against the store, but can observe the way in which previous

<sup>7</sup>Notice that this is not necessarily true if the genuine sequential version of the TG is considered.

customers were treated. It follows that while the store cares about its reputation with future customers, each customer is only interested in his current payoff. Customers will thus play a myopic best reply to whatever strategy they believe the store will choose in the current round. Notice that if the game being played is a PD, no new equilibrium would emerge in this setting, as the short-run player would play D, no matter what he expects the long-run player to choose. In turn, the long-run player cannot do better than play D as well. In the TG the situation is different, because the customer's best reply depends upon the strategy he believes the store will use. Playing T is a best reply against R, while NT is a best reply to NR. It follows that for a cooperative equilibrium to exist, it suffices that the store persuades the customers that R will be played with a sufficiently large probability to make T a best reply.




Figure 1.2: Feasible payoffs in the TG.

It is instructive to think about this setting in terms of the type of reputation the long-run player would like to have. By playing constantly R, the long-run player can induce short-run players to always play T, and would thus get the mutual cooperation payoff,  $R$ . Apparently, he can do better than this, by forming a reputation for playing R with a sufficiently large probability to induce the short-run players to play T. Simple algebra suffices to show that this probability should be slightly larger than  $q^S = \frac{P-S}{R-S}$ . We call  $q^S$  the long-run player's Stackelberg mixed strategy (black dots in Figure 1.2), and the corresponding payoff his Stackelberg equilibrium payoff. This is the payoff the store could get in a game in which it were able to commit to play a (possibly mixed) strategy. If all short-run players choose T, the resulting payoff profile corresponds to the point  $b_2$  in Figure 1.2.

To get his Stackelberg equilibrium payoff, the long-run player should then be able to form a reputation for always playing  $q^S$ . However, in no NE the long-run player can obtain more than



the mutually cooperative payoff  $R$ . This follows from the observation that, if mixed actions are not observable, the long-run player will always have an incentive to play NR in the current round, implying that mixing cannot be optimal. Figure 1.2 shows the long-run player's Stackelberg strategy (black dots), against fully cooperative short-run players (red dots). In this context, on the x-axis there is the long-run player's average payoff, and on the y-axis there is the average payoff of all the short-run players. The long-run player will then choose a point on the red segment that maximizes his payoff. If enough short-run players choose T (as is the case in Figure 1.2), the long-run player will have a strong incentive to choose NR. This result can be easily extended to any other strategy profile in which the long-run player is required to mix at some rounds (Mailath and Samuelson, 2006, Sec. 2.7.2). The result is somehow counterintuitive, and it is not surprising that different modifications of the same basic setting can deliver very different results.

In two influential papers, Fudenberg and Levine (1989) and Fudenberg and Levine (1992) studied the interaction between short and long-run players in the context of incomplete information games and proved that the long-run player can obtain his Stackelberg payoff even when its corresponding strategy is mixed. A fundamental assumption in this class of models is that the long-run player can be of several "types", which are private information<sup>8</sup>.

Among them, there is the "Stackelberg type", that is a player who plays invariably his (possibly mixed) Stackelberg strategy. Their main result is that a sufficiently patient long-run player can obtain his (possibly mixed) Stackelberg payoff. The intuition is that the "normal type" (i.e. the type who maximizes his expected payoff along the repeated game) mimics the behavior of the "type" who is committed to the Stackelberg strategy. Short-run players' beliefs about types will eventually put a probability sufficiently large on the long-run player being a "Stackelberg type", and so they will play a best response to the strategy he plays. If the "normal type" is sufficiently patient, he will prefer to build a reputation for being a Stackelberg type and he would get the corresponding payoff in the long-run.

---

<sup>8</sup>These are known as "commitment types". In their original contribution, Fudenberg and Levine (1989) followed a different path, as they assumed that short-run uncertainty concerned the long-run player's preferences. These are known as "payoffs types". A payoff type, is a player for whom playing the Stackelberg strategy through the repeated game is a strictly dominant strategy. This creates special problems when the Stackelberg strategy is mixed, as one needs to prove that a mixed strategy is strictly dominant for players whose preferences obey the axioms of decision under uncertainty. The introduction of "commitment types" drastically simplify matters, although this comes at the cost of assuming that there may be "crazy" types who mechanically play always the same strategy. See Mailath and Samuelson (2006) p. 464 for a discussion of this matter.

## 1.3 Methods

### 1.3.1 General setting

In this section we present the main methods used to deal with repeated games when players are restricted to use memory-one strategies. We consider a repeated game with a continuation probability of  $0 < \delta \leq 1$ . The game is played between two individuals, called “first player” and “second player”. In every period, player  $i$ 's (pure) action set is  $A_i = \{C, D\}$ , and  $\Delta(A_i)$  is his mixed action set, with  $a_i \in A_i$  denoting a pure action, and  $\alpha_i \in \Delta(A_i)$  denoting a (possibly) mixed one. We denote by  $\Delta(W) = \prod_i \Delta(A_i)$  the set of possible actions profiles, with elements  $\alpha = (\alpha_1, \alpha_2)$ , where  $\alpha_1$  is the (mixed) action chosen by first player, and  $\alpha_2$  is the (mixed) action chosen by the second player. In every period, the set of possible outcomes is  $W = \{CC, CD, DC, DD\} \subset \Delta(W)$ , with a typical element  $w = (a_1 a_2)$ , where  $a_1$  is the realized action of the first player, and  $a_2$  is the realized action of the second player.

The outcome of period  $t$  is then  $w_t = (a_{1t} a_{2t})$ , and the history available in period  $t$  is the sequence of outcomes from 0 to  $t - 1$ ,  $h_t = \{w_0 \dots w_{t-1}\}$ . We denote with  $H_t$  the set of possible histories at round  $t$  and with  $H$  the set of all possible histories. A strategy  $s_i \in S_i$  for player  $i$ , is then a map from the set of possible histories to the set of (mixed) actions:  $s_i : H \rightarrow \Delta(A_i)$ , where  $S_i$  is the set of strategies. The set of strategy profiles is then  $S = S_1 \times S_2$ , and a typical element  $s$  specifies a strategy for each player:  $s = (s_1, s_2)$ .

We denote by  $s_i(h_t)$  player  $i$ 's continuation strategy after  $h_t$ , and by  $\alpha_i(s_i(h_t))$  the (possibly mixed) action that strategy  $s_i$  prescribes to play after  $h_t$ . We give the following definitions:

**Definition 1.** A strategy  $s_i$  is:

- (i) pure if  $\forall h_t \quad \alpha_i(s_i(h_t)) \in \{0, 1\}$ .
- (ii) mixed if  $\exists h_t$  s.t.  $0 < \alpha_i(s_i(h_t)) < 1$ .
- (iii) totally mixed if  $\forall h_t \quad 0 < \alpha_i(s_i(h_t)) < 1$ .

The stage game payoffs are shown in Table 1.3. We are interested in two specific games: the Prisoner's Dilemma (PD), and the Trust Game (TG). The payoff from mutual cooperation, i.e. if both players choose action  $C$ , is  $R$ , while the payoff from mutual defection, i.e. if they both choose action  $D$ , is  $P < R$ , for both players. If the first player cooperates and the second player defects, the first player gets  $S_1 < P$ , and the second player gets  $T_2 > R$ . If the first player

defects and the second cooperates, then, in the PD, the first player gets  $T_1 > R$ , and the second player gets  $S_2 < P$ . In the TG instead, if the first player defects, the payoff of both players is  $T_1 = S_2 = P$ . To sum up, in the PD payoffs satisfy  $T_i > R > P > S_i$ , while in the TG payoffs satisfy  $T_2 > R > P = S_2 = T_1 > S_1$ . Moreover, we assume that  $2R > T_i + S_i$ , so that, in both games, mutual cooperation is preferred by both players to a fair chance of getting  $T_i$  or  $S_i$ .

	$C$	$D$
$C$	$R, R$	$S_1, T_2$
$D$	$T_1, S_2$	$P, P$

Table 1.3: Stage game payoffs

For each possible outcome  $w \in \{CC, CD, DC, DD\}$ ,  $\pi_i(w)$  is the corresponding payoff for player  $i$ , which belongs to the stage game payoff vector  $\pi_i$ , with  $\pi_1 = \{R, S_1, T_1, P\}$  and  $\pi_2 = \{R, T_2, S_2, P\}$ . We denote with  $\pi_i(w_t)$  the stage game payoff for player  $i$  when the outcome in period  $t$  is  $w$ , and by  $\pi(w_t) = (\pi_1(w_t), \pi_2(w_t))$  the corresponding payoff profile. We will sometimes drop the subscript  $t$  when it is clear from the context. For example,  $\pi_1(CC) = R$  is the payoff for player one corresponding to the outcome  $CC$ , and  $\pi(CC) = (R, R)$  is the associated payoff profile. Following Mailath and Samuelson (2006), we denote the set of stage game payoffs generated by pure action profiles as

$$F_p = \{u \in R^2 : \exists w \in W \quad s.t. \quad \pi(w) = u\}$$

The set of *feasible* payoff profiles is then  $F = co(F_p)$ , that is the convex hull of  $F_p$ , and  $F_b = bo(F)$  is the boundary of  $F$ .

Let  $s = (s_1, s_2)$  be a strategy profile and let  $w_t(s)$  be the outcome at time  $t$  in the history generated by  $s$ . We denote with  $\Pi_i(s) = \Pi_i(s_1, s_2) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \pi_i(w_t(s))$ , the total average payoff of player  $i$  when the strategy profile is  $s$ . Since payoffs are averaged by  $(1 - \delta)$ , we can ensure that the payoff profile  $\Pi(s) \in F$  for all  $s$ .

In the set  $F$ , a payoff profile  $u$  is inefficient if there exists  $u' \neq u$  s.t.  $u'_i > u_i$  for all  $i$ . In this case we say that  $u'$  strictly Pareto dominates  $u$ . We denote by  $\hat{F} \subset F_b$  the set of efficient payoff profiles, that is the set of payoff profiles that are not dominated by any other profile. Finally, a player (pure) *minmax* payoff,  $\underline{u}_i$ , is defined as the maximum payoff player  $i$  can get when he is playing a best reply, and his opponent is playing in order to minimize his payoff. Formally:

$$\underline{u}_i = \min_{a_{-i}} \max_{a_i} \pi_i(a_i a_{-i})$$

In the present setting, the minmax payoff is always  $P$  for both players.

A payoff profile  $u$  is individually rational if  $u_i \geq \underline{u}_i$  for all  $i$ , i.e. if it gives to both players at least their minmax payoff. We denote the set of feasible and rational payoffs as  $F_r = \{u \in F : u_i \geq \underline{u}_i \quad \forall i\}$ . Figure 1.3 shows the set of feasible payoff,  $F$ , for a PD (left) and a TG (right). The light gray area is the set of feasible and rational payoffs,  $F_r$ , i.e. the set of feasible payoffs such that no player receives less than his minmax payoff. The red line is the set of efficient and rational payoffs, that is the intersection between  $F_r$  and  $\hat{F}$ , the Pareto frontier.




Figure 1.3: The set of feasible and rational payoffs for a PD (left) and a TG (right). Payoff values are  $R = 4$ ,  $P = 1$ ,  $S_1 = 0$ ,  $T_2 = 6$ , with  $S_2 = 0$ ,  $T_1 = 6$  for the PD, and  $S_2 = T_1 = P = 1$  for the TG. The Stackelberg payoff profile for the second player is then  $b_2 = (1, 5.5)$ .

In figure 1.3,  $b_2$  ( $b_1$ ) is the profile that gives to the second (first) player his maximum efficient and rational payoff, and it corresponds to the payoff profile:

$$b_2 = \left( P, \frac{R(P - S_1) + T_2(R - P)}{R - S_1} \right)$$

In economics, this payoff is called the *Stackelberg* payoff, as it is the payoff a leader can get in a Stackelberg competition, i.e. the maximum payoff a player can get when he can choose before his opponent, knowing the opponent will choose a best reply.

With the previous notation, we can give the following definitions:

- Definition 2.** (i) A strategy  $\hat{s}_i$  is a (strict) best reply to  $s_{-i}$  if:  $\Pi_i(\hat{s}_i, s_{-i})(>) \geq \Pi_i(s_i, s_{-i}) \quad \forall s_i$ .  
(ii) A profile  $\hat{s}$  is a (strict) Nash Equilibrium (NE) if:  $\Pi_i(\hat{s})(>) \geq \Pi_i(s_i, \hat{s}_{-i}) \quad \forall s_i \neq \hat{s}_i, \quad \forall i$ .  
(iii) A profile  $\hat{s}$  is a Subgame perfect NE (SPNE) if, for all  $h_t$ ,  $\hat{s}(h_t)$  is a NE.

We denote by  $\mathcal{E} \subset F$  the set of payoff profiles that can be sustained in a NE, and by  $\mathcal{E}^p \subseteq \mathcal{E}$  the set of payoff profiles that can be sustained in an equilibrium that is also subgame perfect. In the following sections we shall characterize the sets  $\mathcal{E}$  and  $\mathcal{E}^p$  for the PD and the TG.

### 1.3.2 Memory-one strategies

In this section we consider the case in which players can only recall the outcome of the previous period, so they can only use memory-one strategies. We start with the following definitions:

**Definition 3.** Let  $h$  and  $h'$  be two histories and let  $w_{t-1} = (a_{1,t-1}, a_{2,t-1})$  and  $w'_{t-1} = (a'_{1,t-1}, a'_{2,t-1})$  be their outcomes at time  $t - 1$ . A strategy  $s_i$  is:

- (i) *memory-one* if  $w_{t-1} = w'_{t-1}$  implies  $s_i(h_t) = s_i(h'_t)$ ;
- (ii) *reactive* if  $a'_{-i,t-1} = a_{-i,t-1}$  implies  $s_i(h_t) = s_i(h'_t)$ ;
- (iii) *unconditional* if  $s_i(h_t) = s_i(h'_t) \forall h'_t$ .

As an alternative definition, we can say that a strategy is memory-one if the set of histories on which it can condition is restricted to the set  $H_m = \{CC, CD, DC, DD\}$ . Similarly, a strategy is reactive if the set of available histories is  $H_r = \{C, D\}$ , and is unconditional if  $H_u = \emptyset$ .

We denote a memory-one strategy for the first player as  $s_1 = \mathbf{p} = (p_0, p_{CC}, p_{CD}, p_{DC}, p_{DD})^9$ , and for the second player as  $s_2 = \mathbf{q} = (q_0, q_{CC}, q_{CD}, q_{DC}, q_{DD})$ . For example,  $p_{CD}$  ( $q_{CD}$ ) is the probability that the first (second) player cooperates given that in the previous period the first player chose C and the second player chose D. With this notation, reactive strategies are memory-one strategies with  $p_{CC} = p_{DC} = p_C$  and  $p_{CD} = p_{DD} = p_D$ , and unconditional strategies are memory-one strategies with  $p_{CC} = p_{CD} = p_{DC} = p_{DD} = p_0$ . A reactive strategy for the first player is then  $s_1 = (p_0, p_C, p_D)$ , where  $p_C$  ( $p_D$ ) is the probability to cooperate if the opponent cooperated (defected) in the previous round, and  $p_0$  is the probability to cooperate at the beginning of the game. Similarly,  $s_2 = (q_0, q_C, q_D)$  is a reactive strategy for the second player.

Every pair of strategies  $(s_1, s_2)$  (not necessarily memory-one) will induce a probability distribution  $\mathbf{v}$  over the possible outcomes of the game over time. Following Hilbe, Traulsen, et al. (2015), we call the probability that in period  $t$  the outcome is  $w$  as  $v_w(t)$  (which depends on the strategy profile  $s$ ), so the probability distribution over outcomes in period  $t$  is  $\mathbf{v}(t) = (v_{CC}(t), v_{CD}(t), v_{DC}(t), v_{DD}(t))$ . The limit probability of outcome  $w$ ,  $v_w$ , is the probability that in period  $t$  the outcome is  $w$ , for  $t \rightarrow \infty$ , if the limit exists. Thus,  $\mathbf{v} = (v_{CC}, v_{CD}, v_{DC}, v_{DD})$  is the limit probability induced by the strategy profile  $s$ . With this notation, the limit probability that players cooperate in period  $t$ , for  $t \rightarrow \infty$ , is:

---

<sup>9</sup>Recall that  $p_0$  is the probability to cooperate in the first round, i.e. when no history is available, and the remaining components,  $p_w$ , are the probabilities to cooperate if the outcome of the previous period was  $w$ .

$$\begin{cases} \bar{p} = v_{CC} + v_{CD} & \text{for the first player} \\ \bar{q} = v_{CC} + v_{DC} & \text{for the second player.} \end{cases}$$

To see this, just notice that the limit probability that the first (second) player cooperates is the sum of the limit probabilities of the outcomes  $CC$  and  $CD$  ( $CC$  and  $DC$ ), i.e. the outcomes in which the first (second) player cooperates.

More importantly, if one player, say the first, is using a memory-one strategy, the following holds for any opponent's strategy:

**Lemma 0** (Lemma 1 in Hilbe, Traulsen, et al. (2015)).

$$v_{CC} + v_{CD} = (1 - \delta)p_0 + \delta(v_{CC}p_{CC} + v_{CD}p_{CD} + v_{DC}p_{DC} + v_{DD}p_{DD})$$

Thus, if one player is using a memory-one strategy, his limit probability to cooperate is the weighted average (with weights  $(1 - \delta)$  and  $\delta$ ) of the probability to cooperate in the first round and the probability with which he will cooperate in the subsequent rounds. The latter probability is the weighted average of the probabilities to cooperate after each possible outcome, with weights equal to the limit probability of each outcome. Notice that this is independent from the strategy the other player is using.

With vector notation, Lemma 0 becomes :

$$(1 - \delta)p_0\mathbf{1} + \tilde{\mathbf{p}}\mathbf{v} = 0$$

where  $\tilde{\mathbf{p}} = \{\delta p_{CC} - 1, \delta p_{CD} - 1, \delta p_{DC}, \delta p_{DD}\}$ .

## Markov chains

When matched together, the behavior of two memory-one strategies can be described by a Markov chain with the following transition matrix,  $M$ :

$$M = \begin{pmatrix} p_{CC}q_{CC} & p_{CC}(1 - q_{CC}) & (1 - p_{CC})q_{CC} & (1 - p_{CC})(1 - q_{CC}) \\ p_{CD}q_{CD} & p_{CD}(1 - q_{CD}) & (1 - p_{CD})q_{CD} & (1 - p_{CD})(1 - q_{CD}) \\ p_{DC}q_{DC} & p_{DC}(1 - q_{DC}) & (1 - p_{DC})q_{DC} & (1 - p_{DC})(1 - q_{DC}) \\ p_{DD}q_{DD} & p_{DD}(1 - q_{DD}) & (1 - p_{DD})q_{DD} & (1 - p_{DD})(1 - q_{DD}) \end{pmatrix}$$

The stationary distribution of the chain,  $\mathbf{v}^* = (v_{CC}^*, v_{CD}^*, v_{DC}^*, v_{DD}^*)$ , is found by solving:

$$\mathbf{v}^* = (1 - \delta)\mathbf{v}_0 + \delta\mathbf{v}^*.M$$

where  $\mathbf{v}_0 = (p_0q_0, p_0(1 - q_0), (1 - p_0)q_0, (1 - p_0)(1 - q_0))$  is the initial condition of the chain.

If players use totally mixed strategies, the chain is irreducible, and the stationary distribution is unique. Otherwise, there are cases in which  $M$  is not ergodic, and the stationary distribution will then depend upon the initial conditions even when  $\delta = 1$ . We illustrate this point with two examples concerning well known strategies.

**Example 1.** In the terminology introduced above, TfT is a reactive strategy, since a TfT player conditions his current choice only on the choice the opponent made in the previous round. Consider a generalized version of TfT, in which each player cooperates in the initial round with a probability ranging from zero to one. Two such strategies would be represented by  $s_1 = (p_0, 1, 0)$ , and  $s_2 = (q_0, 1, 0)$ . The transition matrix and the stationary distribution (for  $\delta \rightarrow 1$ ) are:

$$M = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \mathbf{v}^* = \begin{pmatrix} p_0q_0 \\ \frac{1}{2}(p_0 + q_0 - 2p_0q_0) \\ \frac{1}{2}(p_0 + q_0 - 2p_0q_0) \\ (1 - p_0)(1 - q_0) \end{pmatrix}$$

Starting from mutual cooperation (i.e.  $p_0 = q_0 = 1$ ), players would end up in a repetition of CC; starting from mutual defection they would end up in a repetition of DD; while starting from one of the other two states will create a cycle in which outcomes CD and DC alternates.  $\triangle$

**Example 2.** Consider now the case where players use strategies of the form:  $\mathbf{p} = (p_0, 1, \frac{1}{5}, \frac{1}{5}, \frac{4}{5})$ , and  $\mathbf{q} = (q_0, \frac{1}{5}, \frac{3}{5}, \frac{2}{5}, 0)$ , The transition matrix and the stationary distribution (for  $\delta \rightarrow 1$ ), are:

$$M = \begin{pmatrix} \frac{1}{5} & \frac{4}{5} & 0 & 0 \\ \frac{2}{25} & \frac{3}{25} & \frac{8}{25} & \frac{12}{25} \\ \frac{3}{25} & \frac{2}{25} & \frac{12}{25} & \frac{8}{25} \\ 0 & \frac{4}{5} & 0 & \frac{1}{5} \end{pmatrix} \quad \mathbf{v}^* = \begin{pmatrix} \frac{5}{69} \\ \frac{26}{69} \\ \frac{16}{69} \\ \frac{22}{69} \end{pmatrix}$$

The stationary distribution in this case is unique and does not depend on the initial conditions.

The total average payoff for the first player in this case is:  $\Pi_1 = \frac{5}{69}R + \frac{26}{69}S_1 + \frac{16}{69}T_1 + \frac{22}{69}P$ , which is independent from the initial conditions.  $\triangle$

## Memory-one strategies in the payoff space

Let  $\mathbf{v}^* = (v_{CC}^*, v_{CD}^*, v_{DC}^*, v_{DD}^*)$  be the stationary distribution, which depends on players' strategies  $s_1, s_2$ , and let  $\boldsymbol{\pi}_1 = \{R, S_1, T_1, P\}$  be the payoff vector for the first player. The total average payoff obtained by the first player is then  $\Pi_1(s_1, s_2) = \boldsymbol{\pi}_1 \mathbf{v}^* = u_1$ . A similar definition holds for the second player. Now fix a strategy for one player, say the second. We can define the set of payoffs profiles that are feasible under strategy  $s_2$  as:

$$F(s_2) = \{u \in F : \exists s_1 \in S_1 \text{ s.t. } (\Pi_1(s_1, s_2), \Pi_2(s_1, s_2)) = u\}$$

In other words, this is the set of payoffs profiles that the first player can reach when the second player is choosing  $s_2$ . Given  $F(s_2)$ , a best reply for the first player is a strategy that gives him the maximum payoff inside  $F(s_2)$ .

It is easy to visualize the set  $F(s_2)$  using Monte Carlo simulations. Given a strategy  $s_2$ , choose randomly 5 probabilities that define a memory-one strategy for the first player and calculate the resulting payoff profile. Repeat this a sufficiently large number of times and obtain an estimate of the set  $F(s_2)$  (see Figure 1.4).

**Example 3.** Figure 1.4 shows the set  $F(s_2)$ , when  $\delta = 1$ , for three of the most common memory-one strategies: TfT ( $s_2 = \mathbf{q} = (1, 1, 1, 0, 0)$ ), Grim ( $s_2 = \mathbf{q} = (1, 1, 0, 0, 0)$ ), and Pavlov (or Win-Stay-Lose-Shift, WLSL, with  $s_2 = \mathbf{q} = (1, 1, 0, 0, 1)$ ). To obtain the graphs, we fixed a strategy for the second player, and we let it play with 1000 randomly generated strategies (black dots), and 100 randomly generated pure strategies (red dots) of the first player.

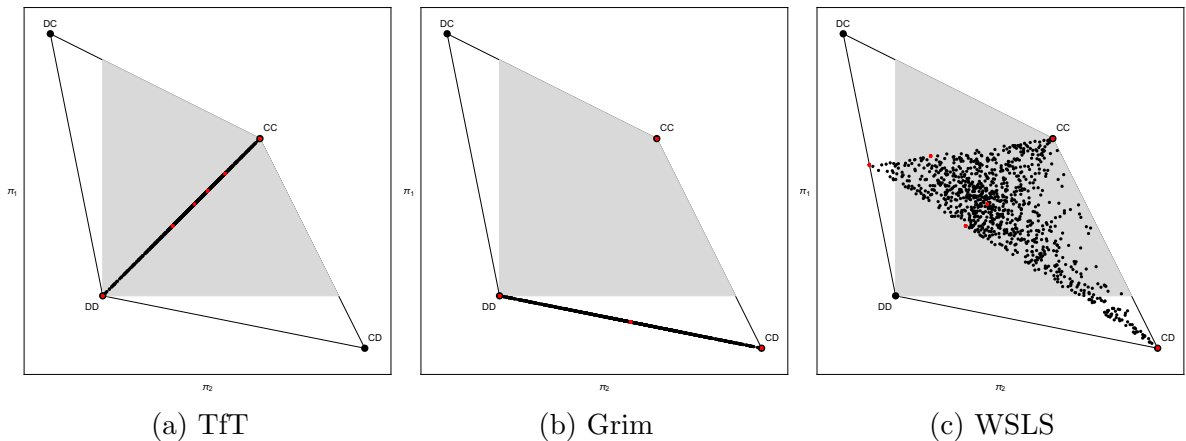


Figure 1.4: The set  $F(s_2)$  for common pure memory-one strategies

Thus, each point in the graphs corresponds to the payoff profile relative to the fixed strategy of the second player, and one random strategy of the first player. Notice that, by choosing



TfT, the second player forces the payoff profile to stay on the segment connecting the  $CC$  point with the  $DD$  point. When the second player chooses Grim, the payoffs are all contained in the segment joining  $DD$  and  $CD$ . An exception is the point  $CC$  which can be obtained when Grim plays against itself. The set of payoffs generated by WSLs is more complex.  $\triangle$

### 1.3.3 Zero Determinant Strategies

In the following we restrict our attention to the PD, so  $T_2 = T_1 = T > R$  and  $S_2 = S_1 = S < P$ . By manipulating the transition matrix  $M$ , Press and Dyson (2012) show the existence of a particular type of strategies in the repeated PD with no discounting, which they called Zero Determinant (ZD) strategies. ZD strategies are special memory-one strategies that allow players to unilaterally enforce a linear relation between payoffs. Thus, if  $s_2$  is a ZD strategy, the set  $F(s_2)$  is linear. In Figure 1.1 we showed three examples of the set  $F(s_2)$  for different types of ZD strategies.

ZD strategies with  $\delta < 1$  are described in Hilbe, Traulsen, et al. (2015)<sup>10</sup>. They noticed that the stationary vector can also be computed as the Abelian mean distribution  $\mathbf{v} = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \mathbf{v}(t)$ , where  $\mathbf{v}(t)$  is the distribution of probabilities over outcomes in period  $t$ . Thus, recalling that  $\boldsymbol{\pi}_i$  is the stage game payoff vector for player  $i$ , we can write the expected payoff in period  $t$  as  $\pi_i(t) = \boldsymbol{\pi}_i \mathbf{v}(t)$ , and the discounted average payoff as:

$$II_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \pi_i(t) = (1 - \delta) \boldsymbol{\pi}_i \sum_{t=0}^{\infty} \delta^t \mathbf{v}(t) = \boldsymbol{\pi}_i \mathbf{v}$$

Then, recalling that  $\tilde{\mathbf{p}} = \{\delta p_{CC} - 1, \delta p_{CD} - 1, \delta p_{DC}, \delta p_{DD}\}$ , we can give the following definition of a ZD strategy (for the first player):

**Definition 4.** A memory-one strategy is ZD if there exist constants  $a, b, c$  such that:

$$(1 - \delta)p_0 \mathbf{1} + \tilde{\mathbf{p}} = a\boldsymbol{\pi}_1 + b\boldsymbol{\pi}_2 + c\mathbf{1}$$

Such a ZD strategy enforces the linear payoff relation:  $aII_1 + bII_2 + c = 0$  (to see this, multiply both sides by  $\mathbf{v}$  and notice that  $II_i = \boldsymbol{\pi}_i \mathbf{v}$ ,  $\mathbf{v} \cdot \mathbf{1} = 1$ , and, from Lemma 0,  $(1 - \delta)p_0 \mathbf{1} + \tilde{\mathbf{p}} \mathbf{v} = 0$ ).

To give a better interpretation of the parameters characterizing a ZD strategy, let  $\phi = -b$ ,  $\Phi = -a/b$ , and  $k = -c/(a + b)$ . A ZD strategy then solves:

$$(1 - \delta)p_0 \mathbf{1} + \tilde{\mathbf{p}} = \phi[\Phi(\boldsymbol{\pi}_1 - k\mathbf{1}) - (\boldsymbol{\pi}_2 - k\mathbf{1})] \tag{1.1}$$

<sup>10</sup>For the original derivation of Press and Dyson (2012) when  $\delta = 1$  see Appendix 1.A.

With this notation,  $k$  is the value at which the ZD strategy intersects the diagonal in the payoff space, and is also the payoff that a ZD strategy gets against itself.  $\Phi$  is the “extortion factor”, that determines how much the ZD-player’s payoff is above (or below) the one of the opponent ( $\frac{1}{\Phi}$  is the slope of the ZD), and  $\phi$  is a normalizing constant to ensure that probabilities are between 0 and 1. As a consequence, a ZD strategy can be represented by a pair  $(\Phi, k)$ , and enforces the linear relation  $\Pi_i = \frac{1}{\Phi}\Pi_{-i} - k\frac{1-\Phi}{\Phi}$ .

Equation 1.1 can be written as:

$$\begin{cases} (1 - \delta)p_0 + \delta p_{CC} = 1 - \phi(R - k)(1 - \Phi) \\ (1 - \delta)p_0 + \delta p_{CD} = 1 - \phi(T - k - \Phi(S - k)) \\ (1 - \delta)p_0 + \delta p_{DC} = -\phi(S - k - \Phi(T - k)) \\ (1 - \delta)p_0 + \delta p_{DD} = -\phi(P - k)(1 - \Phi) \end{cases} \quad (1.2)$$

By forcing the probabilities to be between 0 and 1, we can find the combinations of  $(k, \Phi)$  that can be sustained by a ZD strategy. Specifically, strategies with positive slope exist for all  $k$  such that  $P \leq k \leq R$ , while strategies with a negative slope exist only for  $k$  such that  $\max\{P, \frac{S-T\Phi}{1-\Phi}\} \leq k \leq \min\{R, \frac{T-S\Phi}{1-\Phi}\}$ . If  $\Phi = -1$ , and  $\delta = 1$ , only the strategy with  $k = \frac{T+S}{2}$  exists. Moreover, if  $\delta < 1$ , the parameter  $\Phi$  must satisfy  $-1 < \Phi < 1$ .

**Example 4.** To better understand the restrictions needed on the pairs  $(k, \Phi)$ , consider Figure 1.5. On the left, we plotted 4 examples of ZD strategies for the second player, with  $k = \frac{T+S}{2}$ . On the right, we plotted 3 examples of ZD strategies for which  $b_2 \in F(s_2)$ .

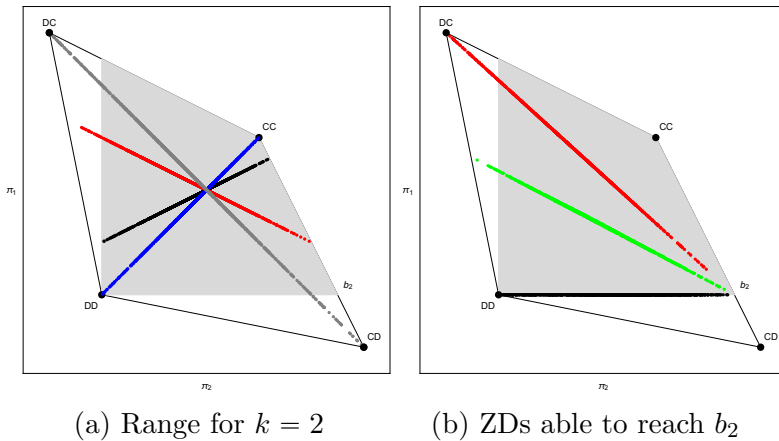


Figure 1.5: Range for ZD strategies

In Figure 1.5a, black dots correspond to a ZD strategy with a positive slope (i.e.  $\Phi > 0$ ), red dots correspond to a strategy with a negative slope (i.e.  $\Phi < 0$ ), while blue and gray dots

correspond to a strategy with  $\Phi = 1$  and  $\Phi = -1$ , respectively. Thus, for  $k = \frac{T+S}{2}$  (and only for  $k = \frac{T+S}{2}$ ), ZD strategies can have any slope between 1 and -1. Figure 1.5b, shows the range of ZD strategies that are able to reach the Stackelberg payoff profile of the second player. Those are all the ZD strategies whose set  $F(s_2)$  lies between the black and the red ones.  $\triangle$

From the expressions of ZD strategies, we can notice that, when  $\delta = 1$  (and only when  $\delta = 1$ ), TfT is a ZD strategy with  $\phi = \frac{1}{T-S}$  and  $\Phi = 1$  (see also Figure 1.4a). Grim instead is never a ZD strategy: even if, when  $\delta = 1$ , it lies on the AllD line for almost all the possible strategies of the first player, when facing a strategy with  $p_0 = p_{CC} = 1$  (i.e. any strategy that is *nice*, as intended by Axelrod (1984)), Grim gets the mutual cooperation payoff, making the relation between players' payoff non linear (see Figure 1.4b).

For a ZD to be *reactive* (so that  $p_{CC} = p_{DC} = p_C$  and  $p_{CD} = p_{DD} = p_D$ ) it must hold  $1 = \phi(R(1-\Phi) - S + \Phi T) = \phi(-P(1-\Phi) - T + \Phi S)$ , which implies  $(R+P-S-T)(1-\Phi) = 0$ . This means that, if the game has equal gains from switching (so  $R+P = T+S$ ), every ZD strategy is reactive. Otherwise, the only reactive ZD strategy is TfT (if  $\delta = 1$ ).

We shall now consider how ZD perform against themselves and against other strategies. We start with the interaction between two ZD strategies. Let  $s_1 = ZD(\Phi, k)$  and  $s_2 = ZD(\Phi, k')$  be two ZD strategies with  $k' > k$ . The resulting payoff is:

$$\Pi_1(s_1, s_2) = \frac{\Phi k + k'}{1 + \Phi} > \frac{\Phi k' + k}{1 + \Phi} = \Pi_2(s_1, s_2)$$

Thus, the strategy with the lowest  $k$  has the highest payoff. It follows that, if in a population players can only use ZD strategies, the population will evolve towards the lowest value of  $k$ , i.e.  $k = P$ , which implies that only mutual defection would be observed.

We now consider the interaction of a ZD strategy against strategies AllC and AllD. The payoff profile of a ZD strategy against an opponent who always cooperates is<sup>11</sup>:

$$\Pi(ZD, AllC) = \left( \frac{R(T-S) - (1-\Phi)k(T-R)}{(R-S) + \Phi(T-R)}, \frac{\Phi R(T-S) - (1-\Phi)k(S-R)}{(R-S) + \Phi(T-R)} \right) \quad (1.3)$$

The payoff profile of a ZD strategy against an opponent who always defects is instead:

$$\Pi(ZD, AllD) = \left( \frac{P(T-S) - (1-\Phi)k(P-S)}{(T-P) + \Phi(P-S)}, \frac{\Phi P(T-S) - (1-\Phi)k(P-T)}{(T-P) + \Phi(P-S)} \right) \quad (1.4)$$

Notice that, if  $\Phi = 0$ , it holds  $\Pi_1(AllC, ZD) = \Pi_1(AllD, ZD) = k$ . This means that, if the second player uses a ZD strategy with  $\Phi = 0$ , the first player will get the same payoff from

---

<sup>11</sup>Notice that the payoff profile do not depend on  $\delta$ . Indeed one can show that this is true for a ZD strategy interacting with *any* unconditional strategy

choosing any of his strategies. Applying simple algebra to Equations 1.3 and 1.4, it is easy to show that that, if  $\Phi > 0$ , it holds  $\Pi_1(AllC, ZD) > k > \Pi_1(AllD, ZD)$ , meaning that the payoff from playing AllC against a ZD strategy is higher than the payoff from playing AllD. The opposite is true if  $\Phi < 0$ , since it holds  $\Pi_1(AllD, ZD) > k > \Pi_1(AllC, ZD)$ .

We can then make the following:

**Remark 1.** *AllC is a best reply to every ZD strategy with  $\Phi \geq 0$ , and AllD is a best reply to every ZD strategy with  $\Phi \leq 0$ .*

The literature has singled out some notable ZD strategies. We shall describe them briefly here, as some of them will turn out to be useful in the theoretical result of the next section.

### Extortionate

Most of the interest in ZD strategies stems from the fact that Press and Dyson (2012) noticed a peculiar type of strategy, which they called extortionate. Extortionate strategies (ExS) can enforce a positive linear relation between player's payoff, and, more importantly, a player using an ExS can ensure that his payoff is never below the one of the opponent (hence the term *extortionate*).

ExS intersect the diagonal of the payoff space at the mutual defection payoff profile, implying  $k = P$ , and have a non negative slope, implying  $0 \leq \Phi \leq 1$ . By plugging those values into Equations 1.2, we can find the following expressions for a ExS:

$$\begin{cases} (1 - \delta)p_0 + \delta p_{CC} = 1 - \phi(R - P)(1 - \Phi) \\ (1 - \delta)p_0 + \delta p_{CD} = 1 - \phi(T - P - \Phi(S - P)) \\ (1 - \delta)p_0 + \delta p_{DC} = -\phi(S - P - \Phi(T - P)) \\ (1 - \delta)p_0 + \delta p_{DD} = 0 \end{cases} \quad (1.5)$$

Notice that equations in 1.5 imply  $p_0 = p_{DD} = 0$ , i.e. in the terminology of Hilbe, Traulsen, et al. (2015), ExS are *cautious*, meaning that they are never the first to cooperate. It follows that two ExS will get the mutual defection payoff when matched with each other, or when matched with any other cautious strategy (see Figure 1.6b ).

Moreover, as Press and Dyson (2012) noticed, the best reply to ExS is AllC. The resulting

payoff profile of an ExS against an AllC is:

$$\Pi(ExS, AllC) = \left( \frac{R(T - S) - (1 - \Phi)P(T - R)}{(R - S) + \Phi(T - R)}, \frac{\Phi R(T - S) - (1 - \Phi)P(S - R)}{(R - S) + \Phi(T - R)} \right)$$

Notice that  $\Pi_1(ExS, AllD) > R$  as long as  $\Phi < 1$ , so that, against AllC, ExS is able to gain more than the mutual cooperation payoff.

In Figure 1.6b, we plot the sets  $F(s_1)$  and  $F(s_2)$ , when  $s_1$  and  $s_2$  are both ExS. Since ExS have  $k = P$ , the payoff profile corresponding to any two EXS strategies is  $(P, P)$ . Notice that the black equalizer strategy in Figure 1.6a can be considered as an ExS, since the second player's payoff is never below the one of the opponent. Moreover, recall that this strategy allows the second player to reach his Stackelberg payoff (see Figure 1.5b). Nonetheless, point  $b_2$  cannot be a Nash equilibrium: if the first player fully complies by choosing AllC, the incentive for the second player to defect and get a higher payoff is too strong, and he will choose AllD.

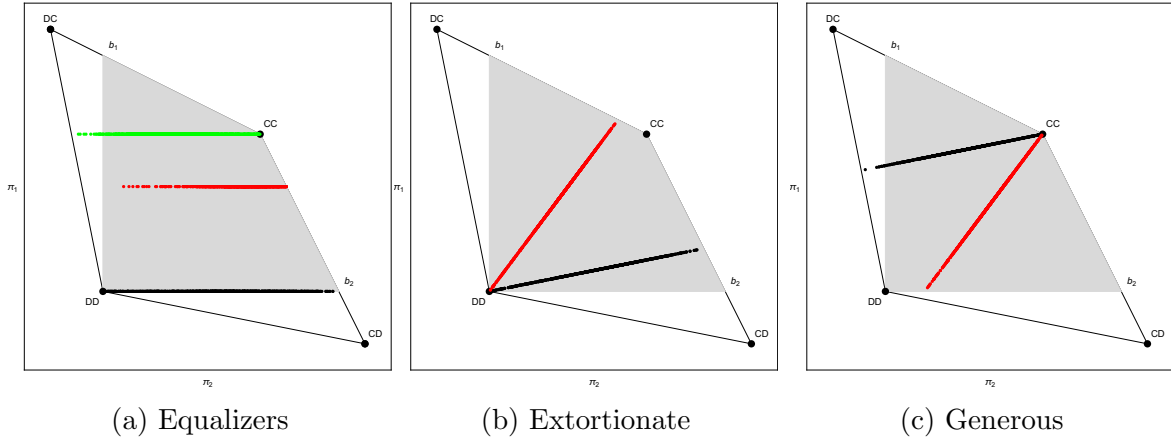


Figure 1.6: Examples of ZD strategies.

## Equalizer

Press and Dyson (2012) show that, in the repeated PD with  $\delta = 1$ , if the first player uses a strategy of the form:

$$\begin{cases} p_{CC} \\ p_{CD} = \frac{(1+p_{DD})R - p_{CC}P - (1-p_{CC}+p_{DD})T}{R-P} \\ p_{DC} = \frac{p_{DD}R + (1-p_{CC})P - (1-p_{CC}+p_{DD})S}{R-P} \\ p_{DD} \end{cases} \quad (1.6)$$

the resulting payoff for the second player is

$$\Pi_2 = \frac{(1 - p_{CC})P + p_{DD}R}{1 - p_{CC} + p_{DD}}$$

which does not depend on the second player's strategy. This implies that he will be indifferent among all his strategies, as they all give him the same payoff, hence the name *Equalizer*. Notice that the first player can fix the opponent's payoff at any level between  $P$  and  $R$ , by choosing proper values of  $p_{CC}$  and  $p_{DD}$  (ensuring that  $0 \leq p_{DC}, p_{CD} \leq 1$ ). For example, by setting  $p_{DD} = 0$ , the first player can enforce  $\Pi_2 = P$ , while by setting  $p_{CC} = 1$ , he can enforce  $\Pi_2 = R$ . For  $\delta \leq 1$ , Equalizer strategies (EqS) can be found by setting  $\Phi = 0$  in Equation 1.1 and solving

$$\tilde{\mathbf{p}} + (1 - \delta)p_0\mathbf{1} = \phi(k\mathbf{1} - \boldsymbol{\pi}_2)$$

Those strategies enforce the linear relation  $\Pi_2 = k$ , and they are of the form<sup>12</sup>:

$$\begin{cases} (1 - \delta)p_0 + \delta p_{CC} = 1 - \phi(R - k) \\ (1 - \delta)p_0 + \delta p_{CD} = 1 - \phi(T - k) \\ (1 - \delta)p_0 + \delta p_{DC} = -\phi(S - k) \\ (1 - \delta)p_0 + \delta p_{DD} = -\phi(P - k) \end{cases} \quad (1.7)$$

Thus, if  $\delta \geq \max\{\frac{T-R}{T-P}, \frac{P-S}{R-S}\}$ <sup>13</sup>, a player using an EqS can fix the opponent's payoff anywhere between the mutual defection and the mutual cooperation payoff, by setting  $P \leq k \leq R$ , and by choosing proper values of  $\phi$  to ensure that probabilities are between 0 and 1. In Figure 1.6a we plotted the set  $F(s_2)$  for three different equalizer strategies for the second player.

## Generous

Stewart and Plotkin (2012) introduced another type of ZD strategies, which they called generous. Generous strategies (GeS) are able to enforce a positive linear relation between players' payoff, and a player using a generous strategy ensures that his payoff is never above the one of the opponent (hence the term *Generous*).

GeS intersect the diagonal of the payoff space at the mutual cooperation payoff profile, implying  $k = R$ , and have a non negative slope, implying  $0 \leq \Phi \leq 1$ . By plugging those values into Equations 1.2, we have the following expressions for a GeS:

<sup>12</sup>For  $\delta = 1$ , those are the same strategies as in Equations 1.6, with  $k = \frac{P(1-p_{CC})+Rp_{DD}}{1-p_{CC}+p_{DD}}$ , and  $\phi = \frac{1-p_{CC}+p_{DD}}{R-P}$ .

<sup>13</sup>Ichinose and Masuda (2018)

$$\begin{cases} (1 - \delta)p_0 + \delta p_{CC} = 1 \\ (1 - \delta)p_0 + \delta p_{CD} = 1 - \phi(T - R - \Phi(S - R)) \\ (1 - \delta)p_0 + \delta p_{DC} = -\phi(S - R - \Phi(T - R)) \\ (1 - \delta)p_0 + \delta p_{DD} = -\phi(P - R)(1 - \Phi) \end{cases} \quad (1.8)$$

Notice that Equations in 1.8 imply  $p_0 = p_{CC} = 1$ , meaning that generous strategies are nice strategies and as such, they are able to get the mutual cooperation payoff when facing each other, or when facing any other nice strategy, as shown in Figure 1.6c.

To conclude, the following Table shows payoff profiles for the most common memory-one strategies, together with some ZD strategies, for  $R = 4, T = 6, S = 0, P = 1$ , and  $\delta = 9/10$ .

	AllC	AllD	TfT	cTfT	WSLS	sExS	mExS	GeS	EqS	EqS <sub>R</sub>
AllC	4, 4	0, <b>6</b>	4, 4	$\frac{18}{5}, \frac{21}{5}$	4, 4	$\frac{16}{7}, \frac{34}{7}$	$\frac{13}{4}, \frac{35}{8}$	4, 4	$3, \frac{9}{2}$	4, 4
AllD	6, 0	<b>1, 1</b>	$\frac{3}{2}, \frac{9}{10}$	1, <b>1</b>	$\frac{69}{19}, \frac{9}{19}$	1, <b>1</b>	1, <b>1</b>	$\frac{32}{17}, \frac{14}{17}$	$3, \frac{3}{5}$	$4, \frac{2}{5}$
TfT	4, <b>4</b>	$\frac{9}{10}, \frac{3}{2}$	<b>4, 4</b>	$\frac{54}{19}, \frac{60}{19}$	<b>4, 4</b>	$\frac{469}{370}, \frac{667}{370}$	$\frac{665}{359}, \frac{818}{359}$	<b>4, 4</b>	$3, \frac{429}{134}$	<b>4, 4</b>
cTfT	$\frac{21}{5}, \frac{18}{5}$	1, 1	$\frac{60}{19}, \frac{54}{19}$	1, 1	$\frac{681}{271}, \frac{621}{271}$	1, 1	1, 1	$\frac{1148}{359}, \frac{1004}{359}$	$3, \frac{177}{67}$	$4, \frac{148}{43}$
WSLS	4, <b>4</b>	$\frac{9}{19}, \frac{69}{19}$	<b>4, 4</b>	$\frac{621}{271}, \frac{681}{271}$	<b>4, 4</b>	$\frac{77}{47}, \frac{137}{47}$	$\frac{23}{11}, \frac{29}{11}$	<b>4, 4</b>	$3, \frac{131}{37}$	<b>4, 4</b>
sExS	$\frac{34}{7}, \frac{16}{7}$	1, 1	$\frac{667}{370}, \frac{469}{370}$	1, 1	$\frac{137}{47}, \frac{77}{47}$	1, 1	1, 1	$\frac{16}{7}, \frac{10}{7}$	$3, \frac{5}{3}$	4, 2
mExS	$\frac{35}{8}, \frac{13}{4}$	1, 1	$\frac{818}{359}, \frac{665}{359}$	1, 1	$\frac{29}{11}, \frac{23}{11}$	1, 1	1, 1	$\frac{14}{5}, \frac{11}{5}$	$3, \frac{7}{3}$	4, 3
GeS	4, <b>4</b>	$\frac{14}{17}, \frac{32}{17}$	<b>4, 4</b>	$\frac{1004}{359}, \frac{1148}{359}$	<b>4, 4</b>	$\frac{10}{7}, \frac{16}{7}$	$\frac{11}{5}, \frac{14}{5}$	<b>4, 4</b>	$3, \frac{10}{3}$	<b>4, 4</b>
EqS	$\frac{9}{2}, \mathbf{3}$	$\frac{3}{5}, \mathbf{3}$	$\frac{429}{134}, \mathbf{3}$	$\frac{177}{67}, \mathbf{3}$	$\frac{131}{37}, \mathbf{3}$	$\frac{5}{3}, \mathbf{3}$	$\frac{7}{3}, \mathbf{3}$	$\frac{10}{3}, \mathbf{3}$	<b>3, 3</b>	<b>4, 3</b>
EqS <sub>r</sub>	4, 4	$\frac{2}{5}, 4$	<b>4, 4</b>	$\frac{148}{43}, 4$	<b>4, 4</b>	2, 4	3, 4	<b>4, 4</b>	<b>3, 4</b>	<b>4, 4</b>

Table 1.4: Normal form - Example

In red are the NE, while in blue are the best replies of the second player (i.e. the column player). cTfT is cautious TfT, which starts by defecting and then plays TfT. sExS is an ExS with  $\Phi = 1/3$  while mExS is an Exs with  $\Phi = 2/3$ . GeS has  $k = R$  and  $\Phi = 2/3$ . EqS<sub>R</sub> is an EqS with  $k = R$ , while EqS has  $k = 3$ . Notice that all Nash equilibria are weak, and none of them involve an ExS.

## 1.4 Equilibria in the PD

In this section, we will give a full description of NE in the repeated PD with memory-one strategies. Our proof is divided into two parts. First, we consider payoff profiles on the Pareto frontier of the set of feasible payoffs. Our first result is that, if just one player uses a memory-one strategy, then in any efficient NE her payoff cannot be pushed below the mutual cooperation payoff  $R$ . This is true independently from the strategy played by the other player. As a corollary we obtain that if both players are constrained to use memory-one strategies, then only the mutual cooperation payoff profile  $(R, R)$  can be sustained as an efficient NE. Second, we show that, even if one restricts the attention to payoff combinations that are not Pareto efficient, when players are constrained to use memory-one strategies, equilibrium payoffs are never above the mutual cooperation payoff. This result is illustrated by Figure 1.7. We shall show that only payoff profiles within the gray square can be sustained as NE when players are restricted to use memory-one strategies.




Figure 1.7: The set of equilibrium payoffs. Only payoffs in the dark gray square can be sustained as NE with memory-one strategies.

### 1.4.1 Equilibria on the Pareto frontier

Let us start by considering payoff profiles that are on the Pareto frontier. We first observe that the best reply against a player who cooperates with probability one after a round in which he cooperated and the other defected (that is a first player with  $p_{CD} = 1$ ), is to defect with probability one. The intuition is obvious. Against a player who always cooperates after being cheated, it is optimal to always cheat.



Nonetheless, setting  $p_{CD} = 1$  can be part of a non subgame perfect NE, if the first player never cooperates on the equilibrium path, i.e. if he also sets  $p_0 = p_{DC} = p_{DD} = 0$ <sup>14</sup>:

**Observation 1.** *If  $p_{CD} = 1$  ( $q_{DC} = 1$ ), in any NE it must be  $p_0 = p_{DD} = p_{DC} = 0$  ( $q_0 = q_{DD} = q_{CD} = 0$ ).*

**Remark 2.** *In any SPNE a player will never cooperate w.p. 1 after being cheated, i.e.  $p_{CD} < 1$ .*

Recall that  $\mathcal{E} \subset F_r$  is the set of payoff profiles that can be sustained in a NE,  $\mathcal{E}_p \subseteq \mathcal{E}$  is the subset restricted to SPNE, and  $\hat{F}$  is the set of efficient payoff profiles (notice that  $\hat{F}$  contains  $b_i$ , the payoff profile that gives to player  $i$  his Stackelberg payoff). Let  $\hat{F}_i$  be the subset of efficient payoff profiles that guarantee to player  $i$  a payoff that is at least as large as the opponent's. These are the efficient payoff combinations in which we say that there is extortion. Formally, for every  $\Pi \in \hat{F}_i$ ,  $\Pi_i \geq \Pi_{-i}$ . The following proposition shows that, if the opponent of player  $i$  is using a memory-one strategy, the only payoff profile inside  $\hat{F}_i$  that can be sustained in a NE is the mutual cooperation profile, independently of which strategy player  $i$  is using.

**Proposition 1.** *For any strategy of player  $i$ , if the opponent is using a memory-one strategy,  $\hat{F}_i \cap \mathcal{E} = (R, R)$ .*

*Proof.* Assume that the first player is using a memory-one strategy. Let  $u^*$  be a NE payoff profile, and let  $u^* \in \hat{F}_2$ . Since  $u^*$  can be sustained as a NE, it must belong to the set of rational payoffs, so it must be  $v_{CD} < 1$  and  $v_{DC} < 1$ . Moreover, since  $u^* \in \hat{F}_2$ , in this equilibrium only the states  $CC$  and  $CD$  are visited, so it must be  $v_{DC} = v_{DD} = 0$ . From Lemma 0<sup>15</sup>,  $v_{DC} = v_{DD} = 0$  implies  $v_{CC} + v_{CD} = (1 - \delta)p_0 + \delta(p_{CC}v_{CC} + p_{CD}v_{CD}) = (1 - \delta)p_0 + \delta(p_{CC}v_{CC} + p_{CD}(1 - v_{CC})) = 1$ , which holds if  $p_0 = p_{CC} = p_{CD} = 1$ , or  $p_0 = p_{CC} = v_{CC} = 1$ , or  $p_0 = p_{CD} = v_{CD} = 1$ . Because of Observation 1,  $p_0 = p_{CD} = 1$  can never be an equilibrium. Then, it must be  $p_0 = p_{CC} = v_{CC} = 1$ , so that  $(R, R)$  is the only possible equilibrium outcome inside  $\hat{F}_2$ .  $\square$

From Proposition 1, it directly follows that, if both players are restricted to memory-one strategies, the only efficient equilibrium outcome is mutual cooperation:

**Corollary 1.** *If both players are using memory-one strategies,  $\hat{F} \cap \mathcal{E} = (R, R)$*

<sup>14</sup>Notice that this is the same strategy that Baek et al. (2016) found as more prevalent in a population with "intermediate" costs of cooperation, see Section 1.2.

<sup>15</sup>Recall Lemma 0 states:  $v_{CC} + v_{CD} = (1 - \delta)p_0 + \delta(v_{CC}p_{CC} + v_{CD}p_{CD} + v_{DC}p_{DC} + v_{DD}p_{DD})$ .

Notice that Corollary 1 implies that  $F_b \cap \mathcal{E} = \{(P, P), (R, R)\}$ , i.e. the only equilibria on the boundary of the payoff space are the mutual cooperation and the mutual defection payoffs. It also implies that in *any* equilibrium in memory-one strategies, the only efficient equilibrium is mutual cooperation, so that players can never get their Stackelberg payoff.

### 1.4.2 Equilibria below the Pareto frontier

Now we move to payoff profiles that are strictly within the set of feasible payoffs. We need a further piece of notation. Let  $B^{PD}$  be the set of payoffs profiles such that both players get between the mutual cooperation and the mutual defection payoff (see Figure 1.7):

$$B^{PD} = \{u \in F : P \leq u_i \leq R \quad i = 1, 2\}.$$

Our main proposition shows that, if both players are restricted to memory-one strategies, and  $\delta \geq \max\{\frac{T-R}{T-P}, \frac{P-S}{R-S}\}$ , the set of Nash equilibria is equal to the set of subgame perfect equilibria, and it is entirely described by  $B^{PD}$ :

**Proposition 2.** *If both players are constrained to memory-one strategies, then  $\mathcal{E} = \mathcal{E}^p = B^{PD}$ .*

**Corollary 2.** *If both players are constrained to pure memory-one strategies, and  $\delta \geq \max(\frac{T-R}{R-P}, \frac{P-S}{T-P})$ , then  $\mathcal{E}^p = \{(R, R), (P, P)\} \subset \mathcal{E} = B_P^{PD} \subset B^{PD}$ , where:*

$$B_P^{PD} = \{(R, R), (P, P), (\delta P + (1 - \delta)R, \delta P + (1 - \delta)R), (\frac{S + \delta T}{1 + \delta}, \frac{T + \delta S}{1 + \delta}), (\frac{T + \delta S}{1 + \delta}, \frac{S + \delta T}{1 + \delta})\}$$

To prove this proposition, we divide the set of equilibrium payoff profiles into four different sets, depending on the number of states that are visited with positive probability by the corresponding equilibrium strategies, and we will show that the union of those sets is  $B^{PD}$ .

Let  $L_j \subset \mathcal{E}$  be the set of payoff profiles that can be sustained in an equilibrium where exactly  $j$  states are visited, with  $j \in \{1, 2, 3, 4\}$  and  $\cup_{j=1}^4 L_j = \mathcal{E}$ . We shall characterize the set of NE for each of those sets, which will be a proof of Proposition 2.

Consider first the set  $L_1$ , that is the set of payoffs profiles that can be sustained in an equilibrium where only one state is visited. Such a state cannot be  $DC$  or  $CD$ , as the corresponding payoff profile gives to one of the two players less than the mutual defection payoff  $P$ . The only visited state can be either  $CC$  or  $DD$ , which implies that equilibrium strategies must be either nice ( $p_0 = p_{CC} = q_0 = q_{CC} = 1$ ) or cautious ( $p_0 = p_{DD} = q_0 = q_{DD} = 0$ ). Clearly, mutual defection

is always an equilibrium, while mutual cooperation can be an equilibrium only if  $\delta$  is sufficiently large. This implies the following:

**Lemma 1.** *If  $\delta \geq \frac{T-R}{T-P}$ :*

$$L_1 = \{\pi : (\pi_1 = \pi_2 = P) \vee (\pi_1 = \pi_2 = R)\} \subset B^{PD}.$$

Notice that  $L_1$  is also the set of equilibrium payoffs on the boundary of the payoff space:

**Remark 3.**  $F_b \cap \mathcal{E} = L_1$ .

The next Lemma deals with the set the of payoff profiles that can be sustained in an equilibrium in which two states are visited. If  $\delta$  is sufficiently large, we show that in  $L_2$  only two paths are possible: either the players alternate between  $CD$  and  $DC$  forever, or they start with mutual defection, and from the second period onward they cooperate. Thus, we can state the following:

**Lemma 2.** *If  $\delta \geq \max(\frac{T-R}{R-P}, \frac{P-S}{T-P})$ :*

$$L_2 = \{\pi : (\pi_i = \frac{S + \delta T}{1 + \delta} \wedge \pi_{-i} = \frac{T + \delta S}{1 + \delta}) \vee (\pi_1 = \pi_2 = (1 - \delta)P + \delta R)\} \subset B^{PD}$$

*Proof.* See Appendix 1.C. □

To avoid unnecessary technical complications, we will disregard the second type of equilibrium, in which players start by defecting and then cooperate forever. It has little theoretical interest, as the corresponding payoff profile approaches the mutual cooperation profile for  $\delta \rightarrow 1$ .

The following Lemma considers the set  $L_3$ , where 3 states are visited on the equilibrium path. Interestingly, in this set there are payoff profiles arising from 2 types of equilibria: in the first type, the state  $CC$  is never visited, and (at least) one of the two players' payoff is fixed at  $P$ , the mutual defection payoff; in the second type, the state  $DD$  is never visited, and (at least) one of the two players' payoff is fixed at  $R$ , the mutual cooperation payoff. Both types of equilibria have payoff profiles that are always inside the set  $B^{PD}$ :

**Lemma 3.** *If  $\delta \geq \max(\frac{P-S}{T-P}, \frac{T-R}{R-S})$ :*

$$L_3 = \{\pi : (\pi_i = P \wedge P \leq \pi_{-i} \leq \frac{S + \delta T}{1 + \delta},) \vee (\pi_i = R \wedge \frac{S + \delta T}{1 + \delta} \leq \pi_{-i} \leq R)\} \subset B^{PD}$$

*Proof.* See Appendix 1.C. □

Our final lemma deals with the set  $L_4$  of equilibria in which all states are visited. Notice that, because all states are visited with positive probability, equilibrium strategies must be optimal in every subgame of the game. This implies that every NE in  $L_4$  is also subgame perfect.

The lemma is based on a result by Ely and Valimaki (2002), which plays an important role in the analysis of repeated games with imperfect monitoring. To introduce this result, we need the notion of a *belief-free equilibrium*. Intuitively, in a belief-free equilibrium, player  $i$  chooses probabilities s.t. his opponent would not be induced to change the action he chooses even if he knows the action that player  $i$  is going to choose. Suppose, for example, that in a PD the first player is indifferent between playing C or D in the current round. Such indifference would not change if he were to know the actual action (C or D) chosen by the second player<sup>16</sup>.

Let  $X_1$  be the total average payoff for the first player if the second player chooses  $C$  in the current round, and let  $Y_1$  be the total average payoff if the second player chooses  $D$ . In a belief-free equilibrium, the first player is indifferent between  $C$  and  $D$ , so the total average payoff when playing  $C$  or  $D$  should be the same. Recall that  $\Pi_1(w)$  is the total average payoff for the first player if the current period outcome is  $w$ <sup>17</sup>. In a belief-free equilibrium it must be  $\Pi_1(CC) = \Pi_1(DC) = X_1$  and  $\Pi_1(CD) = \Pi_1(DD) = Y_1$ .

Then, finding a belief-free equilibrium strategy for the second player involves fixing values of  $X_1$  and  $Y_1$  and solving for the four probabilities<sup>18</sup>. In Appendix 1.B we show that, in a belief-free equilibrium, a strategy for the second player is characterized by probabilities  $q_w$  such that, for every  $w \in \{CC, CD, DC, DD\}$ , it holds:

$$q_w = \frac{\Pi_1(w) - ((1 - \delta)\pi_1(w) + \delta Y_1)}{\delta(X_1 - Y_1)} \quad (1.9)$$

with  $\Pi_1(w) = X_1$  if  $w \in \{CC, DC\}$ , and  $\Pi_1(w) = Y_1$  if  $w \in \{DC, DD\}$ .

In order for those probabilities to be between 0 and 1, we need:  $R \geq X_1 > Y_1 \geq P$  and  $\delta > \max\{\frac{Y_1 - S}{X_1 - S}, \frac{T - X_1}{T - Y_1}\}$ . The next remark directly follows:

---

<sup>16</sup>Notice the distinction with the indifference condition that is required if in a repeated game players play a mixed equilibrium of the stage game. In this case the first player would be indifferent only if the second player is using a mixed strategy *also* in the current round. In a belief free equilibrium instead, he would remain indifferent for any action (pure or mixed) that the second player would choose in the current round.

<sup>17</sup>Where  $w = a_1 a_2$  and  $a_1$  ( $a_2$ ) is the action taken in the previous period by the first (second) player

<sup>18</sup>Notice that there are payoff profiles that can be reached by both a belief-free equilibrium, and an equilibrium that is not belief-free. In appendix 1.B, we give an example of a belief-free equilibrium that is able to reach the mutual cooperation payoff. Clearly, this payoff can be reached also by, for example, a pair of Grim strategies, which would form a SPNE that is not belief-free.

**Remark 4.** [Result from Section 1 in Ely and Valimaki (2002)] *The payoff profile  $u^*$  can be a subgame perfect, belief-free equilibrium if and only if  $u^* \in B^{PD}$ .*

By showing that equilibria in the set  $L_4$  must be belief-free<sup>19</sup>, we are able to state that the only payoff profiles that can be sustained in SPNE where all states are visited, are in the set  $B^{PD} - \{(RR), (PP)\}$ . This is the content of the next Lemma:

**Lemma 4.** *If  $\delta \geq \max\{\frac{T-R}{T-P}, \frac{P-S}{R-S}\}$ :*

$$L_4 = B^{PD} - \{(RR), (PP)\}$$

*Proof.* See Appendix 1.C □

Together with Lemma 1-3, this implies that  $B^{PD}$  is the set of Nash equilibrium payoff profiles:

**Observation 2.**  $\mathcal{E} = \cup_{j=1}^4 L_j = B^{PD}$

Thus, in any equilibrium in memory one strategies players' payoffs are bounded from above by the mutual cooperation payoff, which is the only (strictly) efficient payoff profile.

### 1.4.3 Equilibrium strategies

We now use the results of the previous section to illustrate some key aspects of ZD strategies. Our first result characterizes the type of ZD strategies that can form a NE. We prove that they can only be equalizer, generous or defector. Our second result is more general. We show that, if attention is restricted to SPNE where both players' payoffs are different than the mutual cooperation or the mutual defection payoffs, (that is, SPNE that are strictly inside the set  $B^{PD}$ ) then, among all memory-one strategies, only equalizer can be part of an equilibrium. Finally, we characterize equilibria in reactive strategies, i.e. strategies than only condition on the opponent's last choice, showing that only mutual cooperation or mutual defection can be sustained as a NE.

---

<sup>19</sup>Dutta and Siconolfi (2010) focuses on totally mixed equilibria, and their result implies that, in the PD, a totally mixed subgame perfect equilibrium must be belief-free. We will extend their proof to prove our lemma.

**Nash Equilibria in ZD strategies** From the previous section, recall that a ZD strategy can be described by two parameters,  $k$  and  $\Phi$ , where  $k$  is the payoff a ZD strategy gets against itself, and  $\Phi$  is the slope of the strategy, i.e. how much a player's payoff is below (or above) the one of the opponent. We noticed that every ZD strategy with a non negative slope that gets the mutual cooperation payoff against itself (i.e. with  $k = R$  and  $\Phi \geq 0$ ) is a best reply to itself, and hence is a Nash equilibrium, as well as every strategy with a non positive slope that gets the mutual defection payoff (i.e. with  $k = P$  and  $\Phi \leq 0$ ), and every pair of strategies with a null slope ( $\Phi = 0$ ) (see Remark 1). We can then state the following:

**Proposition 3.** *Let  $s^* = (ZD(k_1, \Phi_1), ZD(k_2, \Phi_2))$  be a profile in ZD strategies. Then  $s^*$  is a NE if and only if, for  $i \in \{1, 2\}$ :*

- $\Phi_i = 0$  (*Equalizer*)
- $\Phi_i \leq 0$  and  $k_i = R$  (*Generous*)
- $\Phi_i \geq 0$  and  $k_i = P$  (*Defector*)

Proposition 3 implies that the full set  $B^{PD}$  of equilibrium payoffs can be sustained by Equalizer strategies. Figure 1.8 shows an example of each kind of equilibrium in ZD strategies.

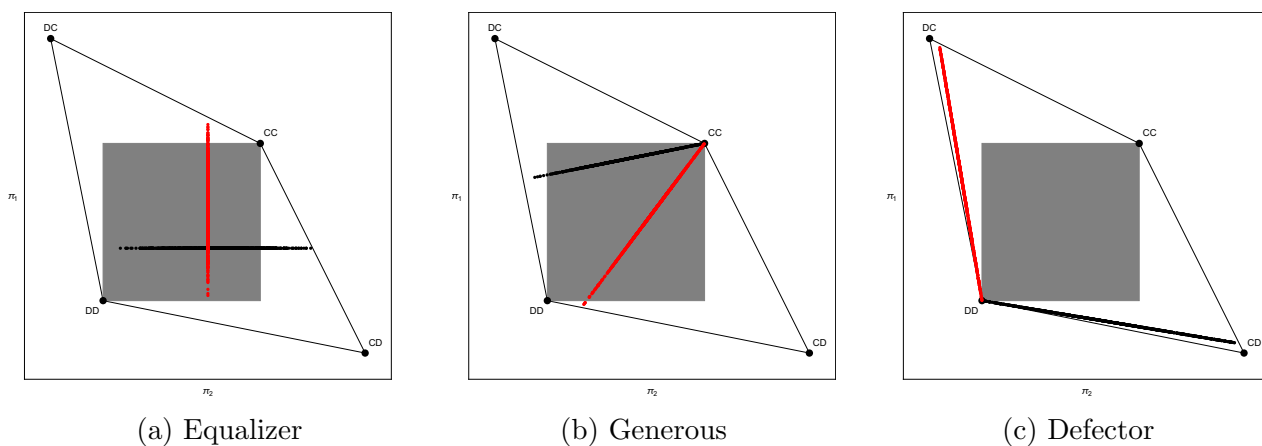


Figure 1.8: Nash equilibria in ZD strategies

Clearly, there are many other equilibria beside those in ZD strategies. We encountered one of them in the previous section, in the proof of Lemma 2. In that equilibrium, players alternate forever between  $CD$  and  $DC$ . We call the strategies that produce this pattern *alternators*. Moreover, in the proof of Lemma 3, we found another type of strategy profile (as specified by 1.22 and 1.21 in Appendix 1.C), in which a player is indifferent among his strategies only after

one or two states. We call strategies that produce this pattern *almost equalizer*. As it turns out, there are no other equilibria in memory-one strategies. This is the content of the following:

**Proposition 4.** *Let  $s^* = (\mathbf{p}, \mathbf{q})$  be a profile in memory-one strategies. Then  $s^*$  is a NE if and only if both strategies are either:*

- *nice:*  $p_0 = p_{CC} = q_0 = q_{CC} = 1$
- *cautious:*  $p_0 = p_{DD} = q_0 = q_{DD} = 0$
- *alternators:*  $p_0 = p_{CD} = q_{DC} = 0, q_0 = p_{DC} = q_{CD} = 1$
- *equalizer:* (See Equations 1.6)
- *almost equalizer:* (See Equations 1.22 and 1.21)

*Proof.* See Appendix 1.C. □

**Subgame perfect equilibria in ZD strategies** To be subgame perfect, a strategy profile should be an equilibrium after any possible history. Generous and Defective strategies can be NE because they only visit one state, either  $CC$  or  $DD$ , but they prescribe to mix after a deviation. If a deviation occurs, players do not have any incentive to punish, as they rather prefer to go back to the equilibrium path. On the contrary, Equalizer strategies, being able to fix the opponent's payoff, independently from his strategy, are by definition also SPNE. Thus, we can directly state the counterpart of Proposition 3 for SPNE in ZD strategies:

**Proposition 5.** *Let  $\mathbf{p}$  be a ZD strategy. Then  $\mathbf{p}$  is part of a SPNE if and only if is an Equalizer strategy.*

*Proof.* See Appendix 1.C □

Our next proposition generalizes this result to the full set of memory-one strategies, as a counterpart of Proposition 4. It shows that, if we exclude the special case of the one *almost equalizer* strategy profile in the set  $L_3$ , then only Equalizer strategies can form a SPNE different than mutual cooperation or mutual defection:

**Proposition 6.** *Let  $s^* = (\mathbf{p}, \mathbf{q})$  be a profile in memory-one strategies, in which the outcome is different than mutual cooperation or mutual defection. Then  $s^*$  is a SPNE if and only if  $\mathbf{p}$  and  $\mathbf{q}$  are Equalizer strategies.*

*Proof.* See Appendix 1.C □

**Reactive strategies** We now discuss the relation between the results in Section 1.4 and reactive strategies. We will show that, for generic PD, if players are restricted to reactive strategies, the only payoff profiles that can be sustained in a Nash equilibrium are the mutual cooperation and the mutual defection payoffs i.e.  $\mathcal{E} = \{(R, R), (P, P)\}$ .

Recall that a reactive strategy is described by three probabilities:  $(p_0, p_C, p_D)$ , where  $p_0$  is the probability to cooperate in the first period, and  $p_j$  is the probability to cooperate if the opponent's previous move was  $j$ , with  $j \in \{C, D\}$ . First observe that, similarly to Observation 1, the best reply against a player who cooperates after a defection is to defect with probability one, thus in equilibrium we must have  $p_D < 1$ . This implies that the results in Proposition 1 and its corollary still hold for reactive strategies<sup>20</sup>, i.e. if  $u^*$  is a NE, and if it belongs to the boundary of the payoff space, then it must be either  $u^* = (R, R)$ , or  $u^* = (P, P)$ .

The next proposition shows that, with reactive strategies, in all equilibria only one state is visited, i.e. the set of equilibrium payoffs is reduced to those in the set  $L_1$ . This is a consequence of the limited possibility of punishment that reactive strategies allow to players, on which were based the other NE in the general case.

**Proposition 7.** *If players are constrained to reactive strategies,  $\delta < 1$  and  $P - S \neq T - R$ , then the only payoff profiles that can be sustained in a Nash equilibrium are  $(R, R)$  and  $(P, P)$ . If  $P - S = T - R$ , the set of equilibrium payoff profiles is  $B^{PD}$ .*

*Proof.* See Appendix 1.C □

In Appendix 1.D we report the case considered by Baklanov (2018) in which  $\delta = 1$ . We show that, even if in this case it might be possible to reach equilibrium payoffs higher than mutual cooperation, the maximum payoff a player can get falls far below his Stackelberg payoff.

---

<sup>20</sup>Lemma 0 with reactive strategies is  $v_{CC} + v_{CD} = (1 - \delta)p_0 + \delta(p_C(v_{CC} + v_{CD}) + p_D(v_{DC} + v_{DD}))$ .



## 1.5 The Trust Game

### 1.5.1 The repeated Trust Game

In this section, we consider an asymmetric version of the PD, the Trust Game (TG). In a TG, the first player chooses whether to cooperate or not, and the second player makes his decision after having observed the first player's choice. To be consistent with the rest of the literature, we sometime call Trust (T) and Not Trust (NT) the cooperative and non-cooperative actions of the first player and Reward (R) and Not Reward (NR) the actions of the second player. Similarly, the first player will sometime be called Trustor and the second player Trustee. Notice that in a standard TG the second player has nothing to decide if the first player chooses NT. In this sense the TG differs from a standard sequential PD. Table 1.5 shows the normal form of the TG. When dealing with the repeated version of the TG, we make the simplifying assumption that the first player can observe the choice of the second player, even when he plays NT. This may be unrealistic in some cases, but not in others. Kreps and Wilson (1982) and Mailath and Samuelson (2006) discuss the "quality choice" game, in which a customer needs to trust a firm, that can produce a good of high or low quality. There are cases in which a customer may observe the quality chosen by the firm, even when he decides not to buy the good.

	R	NR
T	R,R,	S,T
NT	P,P	P,P

Table 1.5: Stage game payoffs in the TG

The stage game has a single pure NE (NT,NR). There are also many Nash equilibria where the first player chooses  $NT$ , and the second player plays  $R$  with a probability smaller than his Stackelberg probability  $q^s = \frac{P-S}{R-S}$ . Notice that this is the probability with which the second player should play  $R$  in order to make the first player indifferent between Trust and Not Trust. Hence, there is a compact set of NE in which both players get their minmax payoff,  $P$ .

For the first player the (pure and mixed) Stackelberg payoff is  $R$ . For the second player, the pure Stackelberg payoff is  $R$ , while the mixed Stackelberg payoff is  $\pi^s = Rq^s + T(1 - q^s) > R$ .

Figure 1.9 represents the set of feasible payoffs for the TG. When the game is repeated between two long-run players, with perfect monitoring over past actions, the Folk Theorem applies and any feasible profile can be an equilibrium outcome provided that it yields both players a payoff

larger than their minmax payoff  $P$ . In the picture, the shaded area represents the set of payoff profiles that can be sustained as subgame perfect NE for sufficiently patient players.




Figure 1.9: The set of feasible payoffs in the TG.

In this section we will characterize all NE for the repeated TG, when players are constrained to use memory-one strategies. The analysis will be similar to the one presented in the previous section with respect to the PD. We will show that, as in the PD, no player can get in equilibrium more than what he can get by mutual cooperation. In particular, there are no NE that sustain the Stackelberg payoff combination  $(P, \pi^s)$ . The set of payoff profiles that can be sustained by NE is in fact much smaller than in the PD. We shall show that  $(R, R)$  is the only payoff profile that can be achieved in an equilibrium in which both players get more than  $P$ .

In the next section we will consider an asymmetric contest in which a long-run second player interacts with a sequence of short-run first players, who can observe some of his past choices. We consider two specifications of this basic setting. In the first, each short-run player can observe the entire history of past choices made by the second player against other short-run players. In the second, the short-run player only observes how the long-run player choose in the previous round. We will show that, when the information about the choices made by the long-run player is sufficiently rich, the long-run player can get his mixed Stackelberg payoff. However, when the short-run player can only observe the last period outcome of the game, the second player can at most secure himself  $R$ . This follows from the fact that when only his last choice is observed, the long-run player can only build two types of reputations, either never Reward or always Reward. To build a reputation in which Reward is played with a positive probability, but smaller than one, short-run players should be able to observe more than one outcome. Perhaps counterintuitively, we prove that providing short-run players more information about the long-run players' past choices may reduce their payoff in equilibrium.

### 1.5.2 Zero Determinant and Unconditional strategies

From the previous section, recall that a ZD strategy can be described by two parameters,  $k$  and  $\Phi$ , where  $k$  is the value at which the ZD strategy intersects the diagonal of the payoff space, and  $\Phi$  is the slope of the strategy, which determines by how much a player's payoff is below (or above) the one of the opponent. A ZD strategy for the first player solves:

$$(1 - \delta)p_0\mathbf{1} + \tilde{\mathbf{p}} = \phi[\Phi(\boldsymbol{\pi}_1 - k\mathbf{1}) - (\boldsymbol{\pi}_2 - k\mathbf{1})]$$

where  $\tilde{\mathbf{p}} = \{\delta p_{CC} - 1, \delta p_{CD} - 1, \delta p_{DC}, \delta p_{DD}\}$ . Expanding the previous equation, we have:

$$\begin{cases} (1 - \delta)p_0 + \delta p_{CC} = 1 - \phi(R - k)(1 - \Phi) \\ (1 - \delta)p_0 + \delta p_{CD} = 1 - \phi(T - k - \Phi(S - k)) \\ (1 - \delta)p_0 + \delta p_{DC} = -\phi(P - k)(1 - \Phi) \\ (1 - \delta)p_0 + \delta p_{DD} = -\phi(P - k)(1 - \Phi) \end{cases}$$

Those probabilities are between 0 and 1 only if  $P \leq k \leq R$ . Figure 1.10a shows some ZD strategies for the first player. In every outcome of the TG, the first player never obtains a larger payoff than the second player. It follows that the only ZD strategies that guarantee to the first player a payoff at least as large as the one of the opponent, are those in which he always defects, i.e. strategies characterized by  $p_0 = p_{DD} = p_{DC} = 0$ . Like in the PD, the first player can use equalizer and generous strategies. In the first case, he can set the second player payoff anywhere between  $P$  and  $R$ , as shown in Figure 1.10a. Notice that, when  $\Phi = 1$ , all ZD strategies for the first player are payoff equivalent to AllD, which implies that he can only get close to the left boundary of the payoff set, but he cannot pick a strategy that would allow equal payoffs above the mutual defection payoff.




Figure 1.10: ZD strategies for the TG.

A ZD strategy for the second player solves:

$$\tilde{\mathbf{q}} + (1 - \delta)q_0\mathbf{1} = \phi[\Phi(\boldsymbol{\pi}_2 - k\mathbf{1}) - (\boldsymbol{\pi}_1 - k\mathbf{1})]$$

where  $\tilde{\mathbf{q}} = \{\delta q_{CC} - 1, \delta q_{DC} - 1, \delta q_{CD}, \delta q_{DD}\}$ . Expanding the previous equation, we have:

$$\begin{cases} (1 - \delta)q_0 + \delta q_{CC} = 1 - \phi(R - k)(1 - \Phi) \\ (1 - \delta)q_0 + \delta q_{CD} = -\phi(S - k - \Phi(T - k)) \\ (1 - \delta)q_0 + \delta q_{DC} = 1 - \phi(P - k)(1 - \Phi) \\ (1 - \delta)q_0 + \delta q_{DD} = -\phi(P - k)(1 - \Phi) \end{cases} \quad (1.10)$$

Figure 1.10b shows an example of several ZD strategies for the second player. Notice that in every outcome of the TG the second player's payoff is never below the one of the opponent. It follows that the only ZD strategies that allow the second player to obtain a payoff that is never above the one of the opponent are those in which he always cooperates, i.e. strategies characterized by  $q_0 = q_{CC} = q_{DC} = 1$  (which implies  $\Phi = 1$ ).

Our first result is that, when  $\delta < 1$ , the second player has no ZD strategies:

**Proposition 8.** *When  $\delta < 1$ , the set of ZD strategies for the second player is empty. When  $\delta = 1$ , all ZD strategies for the second player have either  $k = P$  or  $\Phi = 1$ .*

*Proof.* Notice that, since  $(1 - \delta)q_0 + \delta q_{DC} = 1 - \phi(P - k)(1 - \Phi)$ , and  $(1 - \delta)q_0 + \delta q_{DD} = -\phi(P - k)(1 - \Phi)$ , we have  $q_{DC} = \frac{1}{\delta} + q_{DD}$ . Then,  $q_{DC} \leq 1$  implies  $\delta = 1$ ,  $q_{DD} = 0$ , and  $q_{DC} = 1$ . Finally, to have  $q_{DD} = 0$ , it must be either  $k = P$  or  $\Phi = 1$ .  $\square$

Proposition 8 implies that the family of ZD strategies of the second player is represented by segments of different slopes emanating from  $(P, P)$ . By choosing  $\Phi$ , the second player sets the *extortion factor*, i.e. by how much his payoff will be above the one of the opponent.

Figure 1.10c shows an example of an equalizer strategy for the first player (in red) and for the second player (in black). In fact, this is the only EqS that the second player can use, and it is payoff equivalent to the Stackelberg unconditional strategy  $q^S$ . Indeed, each ZD strategy of the second player enforces the same relation between player's payoffs as an unconditional strategy, meaning that every unconditional strategy with  $q \geq q^S$  is payoff equivalent to an ExS strategy.

**Remark 5.** *Every ZD strategy of the second player is payoff equivalent to an unconditional strategy with*

$$q = \frac{P(1 - \Phi) - S + \Phi T}{R(1 - \Phi) - S + \Phi T}$$

### 1.5.3 Equilibrium payoffs

In this section we will characterize the set of NE for the TG. We follow the same road as in the PD. We first show that within the Pareto frontier, only the mutual cooperation payoff profile can be sustained as an equilibrium, if the first player is constrained to use a memory-one strategy. Second, we show that only a small subset of the set of feasible payoffs strictly below the Pareto frontier can be sustained as a SPNE, even when players are infinitely patient. In particular, in no equilibrium the second player can obtain his Stackelberg equilibrium payoff.

Before moving on, notice that in the TG the cooperative action is a best reply for the first player, if the second player cooperates also. Thus, all the conditions on  $\delta$  for the existence of fully cooperative equilibria are only needed for the second player. This is important because the efficient equilibria that we shall find in this section would remain equilibria also in a context in which the first player would be myopic, that is if he just played a best response to the current choice of the second player.

**Equilibria on the Pareto frontier** Observe that, as in the PD, if the first player cooperates after being cheated (i.e. if he sets  $p_{CD} = 1$ ), the second player will always defect. This can be part of a (not subgame perfect) NE only if the first player never cooperates, i.e. if he sets  $p_0 = p_{DC} = p_{DD} = 0$ . Contrary to the PD, however, the same reasoning does not apply to the second player. If he sets  $q_{DC} = 1$ , i.e. if he cooperates after being cheated, the first player has an incentive to cooperate, rather than to defect.

**Observation 3.** *In any equilibrium, if  $p_{CD} = 1$ , it must be  $p_0 = p_{DC} = p_{DD} = 0$ .*

Recall that  $\mathcal{E}$  is the set of NE,  $\mathcal{E}_p$  is the set of SPNE, and that  $\hat{F}$  is the set of Pareto efficient payoff profiles. What follows is the counterpart of Proposition 1 for the TG.

**Proposition 9.** *If the first player is constrained to use a memory-one strategy,  $\hat{F} \cap \mathcal{E} = (R, R)$ .*

*Proof.* To obtain a payoff profile  $u^* \in \hat{F}$  the equilibrium strategies must be such that  $v_{DD} = v_{DC} = 0$  (otherwise the payoff profile would not be on the Pareto frontier). This implies  $v_{CC} + v_{CD} = 1$ , which, by Lemma 0<sup>21</sup>, holds if either  $p_0 = p_{CC} = v_{CC} = 1$ , or  $p_0 = p_{CD} = v_{CD} = 1$ , or  $p_0 = p_{CC} = p_{CD} = 1$ . From Observation 3, if  $p_0 = p_{CD} = 1$ ,  $u^*$  cannot be a Nash equilibrium. If instead  $v_{CC} = 1$ , then  $u^* = (R, R)$ .  $\square$

<sup>21</sup>Recall Lemma 0 states:  $v_{CC} + v_{CD} = (1 - \delta)p_0 + \delta(v_{CC}p_{CC} + v_{CD}p_{CD} + v_{DC}p_{DC} + v_{DD}p_{DD})$ .

**Equilibria below the Pareto frontier** We now characterize NE that are strictly below the Pareto frontier. Let  $B^{TG}$  be the set of payoff profiles such that the payoff of the first player is  $P$ , and the payoff of the second player is at most  $R$ :

$$B^{TG} = \{u \in F : u_1 = P \text{ and } P \leq u_2 \leq R\}.$$

The next proposition shows that in the TG the only payoff profiles that can be sustained in an equilibrium in memory-one strategies are either  $(R, R)$  or elements of  $B^{TG}$  (See Figure 1.11)<sup>22</sup>.

**Proposition 10.** *If players are constrained to memory-one strategies,  $\mathcal{E} = \mathcal{E}_p = (B^{TG} \cup (R, R))$*




Figure 1.11: NE payoff profiles in the TG with memory-one strategies

As with the PD, we will proceed by describing the sets  $L_j$ , that is the set of payoff profiles that can be sustained in equilibria in which exactly  $j$  states are visited.

Consider first the set  $L_1$ . Notice that, as for the PD, the repetition of the outcome  $CD$  cannot be an equilibrium. Moreover, neither can the repetition of the outcome  $DC$  be an equilibrium: even if both players' payoffs are equal to their minmax payoffs, the path would imply  $q_{DC} = 1$  and  $p_{DC} = 0$ , but if the second player is willing to cooperate, the first player has an incentive to set  $p_{DC} = 1$  and choose  $C$ . Thus, the only visited outcome can be either  $CC$  or  $DD$ . Mutual defection is always an equilibrium, while mutual cooperation is an equilibrium only if  $\delta$  is large enough:

**Lemma 5.** *If  $\delta \geq \frac{T-R}{T-P}$*

$$L_1 = \{\pi : (\pi_1 = \pi_2 = P) \vee (\pi_1 = \pi_2 = R)\}$$

<sup>22</sup>We omit here the payoff profile  $u_1 = u_2 = (1 - \delta)P + \delta R$ , i.e. the payoff resulting from the equilibrium where both players defect in the first period, and cooperate forever from the second period onward. Notice that, as in the PD, this profile tend to  $R, R$  when  $\delta \rightarrow 1$ , and to  $P, P$  when  $\delta \rightarrow 0$

Consider now the set  $L_2$ , where on the equilibrium path players visit 2 states. Recall that in the PD there are only 2 equilibria: either players alternate between CD and DC, or they start by mutual defection, and from the second period they cooperate forever. In the TG, only the latter is an equilibrium. To see this, consider that alternating between CD and DC would yield to the first player a payoff below his minmax payoff  $P$ . The next Lemma shows that there are no other equilibria beside the one in which only CC is visited in the long-run.

**Lemma 6.** *If  $\delta \geq \frac{T-R}{R-P}$*

$$L_2 = \{\pi : (\pi_1 = \pi_2 = (1 - \delta)P + \delta R)\}$$

*Proof.* See Appendix 1.C. □

As with the PD, we omit this equilibrium from the following analysis, as the corresponding payoff profile approaches  $(R, R)$  when  $\delta \rightarrow 1$ .

The next Lemma considers the set  $L_3$ , where only one state is never visited on the equilibrium path. Recall that, in the PD, in this set there are two types of equilibrium paths: one where the state  $CC$  is never visited, and one where the state  $DD$  is never visited. In the TG, neither of those paths can be an equilibrium. It is easy to see why the state  $CC$  must be visited. In any path in which only the other states are visited, the first player's payoff is bound to be below  $P$ , and hence it cannot be a NE. In the Appendix 1.C we will show that there are neither equilibria in which only the state  $DD$  is never visited. In fact, the next Lemma shows that the set  $L_3$  is empty:

**Lemma 7.**

$$L_3 = \emptyset$$

*Proof.* See Appendix 1.C. □

Our final Lemma considers the set  $L_4$ , that is the set of equilibria that visit every state. Recall that, since every state is visited on the equilibrium path, equilibria in this set must be subgame perfect. Moreover, in the PD, the proof of Lemma 4 implies that every equilibrium in  $L_4$  must be belief-free, and that a strategy for the second player in a belief equilibrium is characterized by probabilities  $q_w$  such that, for every  $w \in \{CC, CD, DC, DD\}$ , it holds:

$$q_w = \frac{\Pi_1(w) - ((1 - \delta)\pi_1(w) + \delta Y)}{\delta(X - Y)} \tag{1.11}$$

with  $\Pi_1(w) = X$  if  $w \in \{CC, DC\}$ , and  $\Pi_1(w) = Y$  if  $w \in \{DC, DD\}$ , where  $X$  and  $Y$  are the total average payoffs for the first player if the second player today cooperates or defects, respectively, i.e.  $X = \Pi_1(CC) = \Pi_1(DC)$  and  $Y = \Pi_1(CD) = \Pi_1(DD)$ . We will show that, in the TG, equilibrium profiles in  $L_4$  are characterized by strategies as in 1.11 for the first player, and by strategies that are either as in 1.11, or equal to the unconditional Stackelberg strategy, for the second player. This implies that any equilibrium in  $L_4$  gives to the first player his minmax payoff  $P$ , and gives to the second player at most the mutual cooperation payoff  $R$ :

**Lemma 8.** *If  $\delta \geq \max\{\frac{T-R}{T-P}, \frac{P-S}{R-S}\}$ ,*

$$L_4 = \{\pi \in F : \pi_1 = P \quad \text{and} \quad P < \pi_2 \leq R\}$$

*Proof.* See Appendix 1.C □

Thus, Lemma 8 shows that in the set  $L_4$  there are all payoffs profiles in  $B^{TG}$  except for the mutual defection payoff. Together with Lemma 5-7, this implies that the set of Nash equilibrium payoff profiles is  $B^{TG} \cup (R, R)$ .

**Observation 4.**  $\cup_{j=1}^4 L_j = \mathcal{E} = (B^{TG} \cup (R, R))$

### 1.5.4 Equilibrium strategies

As we did for the PD, we now use the previous findings to describe equilibrium strategies in the TG. We first characterize the set of NE in ZD strategies when  $\delta = 1$ . As with the PD, a NE profile can only be in Generous, Defective or Equalizer strategies, while a SPNE can only be in Equalizer strategies. Contrary to the PD, when we restrict attention to SPNE, only payoffs profiles inside  $B^{TG}$  can be sustained in an equilibrium in ZD strategies. The intuition is that it is not possible for the second player to use a nice strategy, and at the same time to make the first player indifferent. The second result is that, when  $\delta < 1$ , if we restrict attention to equilibria different than mutual cooperation or mutual defection, then, among all memory one strategies, the only SPNE strategy for the second player is his Stackelberg strategy. Finally, we show that, as in the PD, when we restrict attention to reactive strategies, only mutual cooperation and mutual defection can be sustained as a NE (not necessarily subgame perfect).



**Nash Equilibria in ZD strategies** As we saw in Section 1.5.2, Proposition 8, when  $\delta < 1$  the set of ZD strategies for the second player is in fact empty. Thus, equilibria in ZD strategies only exist when  $\delta = 1$ . Just like in the PD, AllC is a best reply to any ZD strategy with  $\Phi \geq 0$ , and AllD is a best reply to any ZD strategy with  $\Phi \leq 0$  (we omit the trivial proof of this fact). Thus, as in the PD, the TG admits three types of NE strategy profiles: Equalizer, Generous and Defectors strategies. Figure 1.12 represents each of these types of equilibria (red dots represent  $s_1$ , black dots  $s_2$ ). The main difference with the PD is that in any equilibrium of the TG, the first player's payoff can be either  $R$  (as in Figure 1.12a ) or  $P$  (as in Figures 1.12b and 1.12c). The discussion above is summarized by the following:

**Remark 6.** Fix  $\delta = 1$  and let  $s^* = (ZD(k_1, \Phi_1), ZD(k_2, \Phi_2))$  be a profile in ZD strategies.

Then

$s^*$  is a NE if and only if:

- $\Phi_1 = \Phi_2 = 0$  and  $k_2 = P$  (Equalizer)
- $\Phi_2 = 1, \Phi_1 \geq 0$  and  $k_1 = k_2 = R$  (Generous)
- $\Phi_2 \leq 0, k_2 = P$  and  $\Phi_1 = 1$  or  $k_1 = P$  (Defector)

From the previous remark, and as we can see in Figure. 1.12, all equilibria in equalizer strategies have  $\Pi_1 = P$ , and  $P \geq \Pi_2 \geq R$ . Moreover, the mutual cooperation payoff profile  $(R, R)$  can be sustained only if equilibrium strategies are Generous, and if  $\delta \geq \max\{1 - \phi(R - P), \phi(T - R)\}$ , with  $\phi < 1/(T - R)$ .




Figure 1.12: ZD equilibria in the TG

We can state the equivalent of Proposition 3, regarding SPNE in ZD strategies. As with the PD, only equalizer strategies can be part of a SPNE equilibrium. This has the interesting implication that mutual cooperation cannot be sustained by a SPNE with ZD strategies.

**Proposition 11.** *Fix  $\delta = 1$  and let  $s^* = (ZD(k_1, \Phi_1), ZD(k_2, \Phi_2))$  be a profile in ZD strategies. Then  $s^*$  is a SPNE if and only if those are Equalizer strategies:  $\Phi_1 = \Phi_2 = 0$  and  $k_2 = P$*

The proof of this proposition is similar to the analogous Proposition 3 and it is omitted.

**Nash Equilibria in memory-one strategies** Clearly, there are other equilibria beside those in ZD strategies. For example, the same payoff profiles implied in Remark 6 can be attained in equilibria where the second player uses an unconditional strategy, by setting  $q = 1$  (generous),  $q = 0$  (defector), or  $q = q^S$  (Equalizers). Moving to the full set of memory one strategies, from the proof of Proposition 10 we learned that any NE in memory-one strategies either visits only one state, or it visits all the states. Moreover, we know that, if an equilibrium only visits one state, equilibrium strategies must be either nice or cautious. If instead the equilibrium visits all the states, then the first player must use an Equalizer strategy, while the second player must use an unconditional strategy, or a ZD-Equalizer (if  $\delta = 1$ ). We can then state the following:

**Remark 7.** *Let  $s^* = (\mathbf{p}, \mathbf{q})$  be a profile in memory-one strategies. Then  $s^*$  is a NE if and only if strategies are either:*

- *both nice:  $p_0 = p_{CC} = q_0 = q_{CC} = 1$*
- *both cautious:  $p_0 = p_{DD} = q_0 = q_{DD} = 0$*
- *both equalizer (if  $\delta = 1$ )*
- *$\mathbf{p}$  equalizer and  $\mathbf{q} = q^S$  unconditional*

The fourth type of equilibrium deserves a special attention. In this equilibrium, the first player chooses a strategy  $\mathbf{p}$  in which the conditional probabilities of playing Trust after each outcome are such that at every round the second player is indifferent between playing Reward and Not Reward. The second player mixes between Reward and Not Reward at every round with his Stackelberg probability  $q^S$  in order to make the first player indifferent between Trust and Not Trust at every round. This equilibrium is thus of the type:

$$\mathbf{q}^* = q^S \quad \mathbf{p}^* = \begin{cases} (1 - \delta)p_0 + \delta p_{CC} = 1 - \phi(R - k) \\ (1 - \delta)p_0 + \delta p_{CD} = 1 - \phi(T - k) \\ (1 - \delta)p_0 + \delta p_{DC} = -\phi(P - k) \\ (1 - \delta)p_0 + \delta p_{DD} = -\phi(P - k) \end{cases} \quad (1.12)$$

Our next proposition shows that, when  $\delta < 1$ , only the strategy profile  $s = (\mathbf{p}^*, \mathbf{q}^*)$  can form a SPNE different than mutual defection or mutual cooperation.

**Proposition 12.** *Fix  $\delta < 1$  and let  $s$  be a SPNE profile different from mutual cooperation or mutual defection. Then  $s$  is a SPNE if and only if  $s = (\mathbf{p}^*, \mathbf{q}^*)$ .*

Proposition 12 implies that the first player's payoff in any mixed equilibrium is fixed at  $P$ , and that the maximum payoff the second player can get in a mixed equilibrium is bounded from above by the mutual cooperation payoff. Figure 1.13 shows the mixed equilibrium in which the second player is able to get the mutual cooperation payoff (recall that black dots represent the second player's strategy, and red dots represent the first player strategy).




Figure 1.13: Mixed SPNE: Maximum payoff for the second player - TG

Finally, we consider the case of reactive strategies, i.e. strategies that only condition on what the opponent did the previous round. Notice that the proof of Proposition 7 for the PD holds also for the case of a TG. This implies that, when players are restricted to reactive strategies, only mutual defection and mutual cooperation can be NE, and only mutual defection is also SPNE:

**Remark 8.** *If players are restricted to reactive strategies and  $\delta < 1$ ,  $\mathcal{E} = \{(R, R), (P, P)\}$  and  $\mathcal{E}_P = (P, P)$*

Summarizing the results, when players are restricted to memory-one strategies, not all payoffs combinations can be part of an equilibrium. Proposition 10 shows that in any equilibria different than mutual cooperation or mutual defection, the first player's payoff is  $P$ , and second player's best payoff is  $R$ , the same payoff he would have get by playing a pure strategy. Thus, like in the PD, there is no gain in mixing and mutual cooperation is the only Pareto efficient payoff.

## 1.6 Short-run vs Long-run: does it pay to extort?

As we noticed in Section 1.5.3, in the repeated TG efficient equilibria exist provided that the second player is sufficiently patient. These would remain equilibria, even if the first player were to be completely myopic. As we shall see in the next chapter, this is important because it allows us to provide an empirical test for the predictions of this model. It suffices to observe the interaction between a single long-run second player who plays with a sequence of short-run first players who can observe his past choices.

In this section we consider a model in which a single long-run (LR) second player interacts with a sequence of short-run (SR) first players, in the spirit of Fudenberg and Levine (1989) (see Mailath and Samuelson (2006) for a textbook presentation of these models). The type of situations we have in mind is a single shop which serves sequentially a large group of consumers who can observe the way in which previous consumers have been treated. The Trust Game is thus played by a short-lived agent (the consumer) in the role of the trustor, against a long-lived agent (the shop) in the role of the trustee. In what follows we shall always refer to the first player as the short-run player and to the second player as the long-run player. Notice that in this setting the long-run player cares about the way his current choices influence future short-run players' behavior. Short-run players do not have such a concern, as they only play once. New equilibria emerge in this setting, because a sufficiently patient long-run player will prefer to play cooperatively today to induce short-run players to play cooperatively in the following rounds. Notice that this is a point in which the Trust Game differs from the PD. In a PD, letting one short-run player to observe the past behavior of the long-run player does not affect equilibria, as the short-run player will defect no matter what he learns about the long-run player's past choices. So the only equilibrium would be a repetition of mutual defection, just like in the one-shot game.

### 1.6.1 Short-run and long-run players with memory-one

Let us start with the model in which each SR player only observes the previous period's outcome of the game. A strategy for the SR player is a mapping from the four possible outcomes of the game to the set of mixed strategies. The SR player is thus forced to play a memory-one strategy as defined in 1.3.2. Notice that in this case no possibility arises for the LR player to form a reputation for playing a mixed strategy, as the SR player can only condition on pure

actions. We can thus apply the results obtained in Section 1.5.3. In particular, if the SR player can observe the choices made in the previous round both by the long and the (previous) short run player, then by proposition 10 in any SPNE the maximum payoff the LR player can obtain in equilibrium is  $R$ . If the SR player can only observe the choice made by the LR player, but not the one made by the (previous) SR player, then she is forced to use a reactive strategy. In this case, Remark 8 applies, and in any NE only outcomes  $CC$  and  $DD$  are observed.

Finally, in the presence of commitment types, we can directly apply the results in Fudenberg and Levine (1989), implying that the only equilibrium outcome in memory one strategies (and in reactive strategies), is the mutual cooperation profile. To see this, notice that when the SR player is restricted to memory one strategies, the LR player can commit only to two types, both pure: he can commit to be either an AllC type, or an AllD type. Then, results in Fudenberg and Levine (1989) implies that the LR player will commit to the strategy that gives him the maximum payoff, which, in this context, implies full cooperation.

### 1.6.2 Observable mixtures

We shall now consider a repeated game with short-run, long-run players with observable mixtures. In particular, we shall assume that at each round the long-run player chooses a probability  $q_t \in [0, 1]$ , which determines the probability with which he will play Reward at  $t$ .  $q_t$  will be observed by the short-run player at  $t + 1$ . Notice that we assume that the short-run player observes  $q_{t-1}$ , but not the *action* chosen by the long-run player. This model is thus analogous to the one discussed by Barlo et al. (2009), in which memory-one strategies are considered in a context in which mixed actions are observed.

We restrict the strategy set  $\mathcal{Q}$  of the long-run player by assuming that the probability  $q_t$  with which he plays Reward may depend on  $t$ , but not on the previous history of the game. A strategy of this type is thus a sequence  $Q = \{q_t\}_{t=1}^{\infty}$ .

The short-run player's strategy may condition on  $t$ , and, if  $t > 1$ , on the mixed strategy  $q_{t-1}$  chosen by the long-run player in  $t - 1$ . A strategy of this type is thus a sequence  $F = \{f_t(r)\}_{t=1}^{\infty}$ , where  $r = \emptyset$  if  $t = 1$ , and  $r = q_{t-1}$  if  $t > 1$ .  $f_1(\emptyset) \in [0, 1]$  is the probability of playing  $C$  at the first round, when no previous history is available. For  $t > 1$ ,  $f_t$  is a map,  $f_t : [0, 1] \rightarrow [0, 1]$ , and  $f_t(r)$  is the probability of choosing  $C$  at round  $t$  given the mixed strategy  $q_{t-1} = r \in [0, 1]$  chosen by the long-run player in  $t - 1$ .  $\mathcal{F}$  is the set of strategies for the short-run player.

A strategy profile for this game is a pair  $(F, Q)$  where  $Q \in \mathcal{Q}$  is the strategy chosen by the long-run player and  $F \in \mathcal{F}$  is the strategy of the short-run player. Given a strategy profile  $(F, Q)$ , players' payoff at time  $t > 1$  are:

$$\pi_{1t}(F, Q) = (1 - f_t(q_{t-1}))P + f_t(q_{t-1})(Rq_t + S(1 - q_t)) \quad (1.13)$$

$$\pi_{2t}(F, Q) = (1 - f_t(q_{t-1}))P + f_t(q_{t-1})(Rq_t + T(1 - q_t)) \quad (1.14)$$

As is customary in this literature, we shall assume that the short-run player breaks an eventual indifference by choosing the strategy that the long-run player prefers him to choose. In the TG, when the first player is indifferent between C and D, he will play C. With this assumption, the short-run player best response function is

$$BR_1(r) = \begin{cases} 1, & \text{if } r \geq q^S \\ 0, & \text{if } r < q^S \end{cases} \quad (1.15)$$

where  $q^S = \frac{P-S}{R-S}$  is the Stackelberg mixed strategy for the long-run player.

A strategy profile  $(F, Q)$  is a Nash equilibrium for this game if  $f_t(q_{t-1}) \in BR_1(q_t) \quad \forall t$ , that is if the short-run player plays a (myopic) best response to the strategy chosen by the long-run player, *at round t*. At the same time,  $Q$  maximizes the inter-temporal payoff of the long-run player. That is,  $Q$  solves:

$$\max_{Q \in \mathcal{Q}} \sum_{t=1}^{\infty} \delta^{t-1} \pi_{2t}(F, Q)$$

We shall consider *stationary* NE, that is equilibria in which the long-run player chooses the same mixed action at each round  $t$  ( $q_t = \hat{q} \in [0, 1] \quad \forall t$ ) and the short run-player chooses the same map ( $f_t = \hat{f} \quad \forall t > 1$ ). Then,  $\hat{q}$  and  $\hat{f}(\hat{q})$  are the probabilities with which the cooperative action is chosen at each round. Notice that, because of the way in which we defined the best response correspondence, in equilibrium no short-run player will use a mixed strategy.

Given a stationary strategy  $\hat{f}$ , we shall denote with  $\underline{q}(\hat{f})$  the smallest value of  $q$  such that  $\hat{f}(q) = 1$ . That is,  $\underline{q}(\hat{f})$  is the smallest probability of cooperative behavior by the second player that induces the first player to play cooperatively. If  $\hat{f}(q) = 0$  for every  $q$ , then  $\underline{q}(\hat{f}) = \emptyset$ .

Our main result is that stationary equilibria are grouped into two sets. In the first set there are all equilibria in which the long-run player cooperates with a probability  $\hat{q} < q^S$  and the short-run player sets  $\hat{f}(q) = 0$  for all  $q$ . These are all the non-cooperative equilibria corresponding to the payoff profile  $(P, P)$ . The second set of NE only exists if the long-run player is sufficiently

patient. It contains all strategy profiles such that  $\hat{q} \geq q^S$  and  $\underline{q}(\hat{f}) = \hat{q}$ . These are the equilibria in which short-run players are willing to cooperate provided that the long-run player cooperates with a probability at least as large as  $\underline{q}(\hat{f})$ , and the long-run player chooses exactly this probability. The payoff profiles associated to these equilibria are represented by the thick gray line in the figure below.




Figure 1.14: Set of payoff profiles that can be sustained in stationary equilibria with observable mixtures, that is the union of the dark gray line on the right boundary of the payoff space, and the red point  $DD$

This is the content of the following

**Proposition 13.** *Let  $(\hat{f}, \hat{q})$  be a stationary equilibrium. Then either  $\hat{q} < q^S$  and  $\underline{q}(\hat{f}) = \emptyset$ , or  $\hat{q} \geq q^S$  and  $\underline{q}(\hat{f}) = \hat{q}$ .*

*Proof.* See Appendix 1.C □

A direct consequence of Proposition 13 is that the only equilibrium surviving the elimination of weakly dominated strategies has  $\hat{q} = q^S$  and  $\underline{q}(\hat{f}) = q^S$ , i.e. is the one that gives to the second player his Stackelberg payoff.

Proposition 13 should be contrasted with Proposition 10 and Remark 8 in Section 1.5.3. We noticed that, since the best reply of the first player when the second player cooperates is to cooperate also, the mutual cooperation equilibrium is the only efficient equilibrium also when the first player is myopic, i.e. when he only care about the *current period's* payoff. Thus, when the short-run player can only condition on the *realized* action of the long-run player in the previous round, no extortion can take place in equilibrium. When instead he can condition on the mixture that the long-run player used in the previous round, new equilibria emerge in which the long-run player can extort a larger payoff.

## 1.7 Discussion and Conclusions

In this chapter we developed a theoretical model for the repeated PD when players have limited memory. Motivated by the rise in interest for the so called Zero-Determinant strategies, which allow players to reach extortionate outcomes, we characterized the full set of equilibrium strategies and payoffs. Specifically, our first result is that mutual cooperation is the only outcome on the Pareto frontier in an equilibrium in memory one strategies. Our second result is that in *any* equilibrium in memory one strategies, players' payoffs are bounded from above by the mutual cooperation payoffs. This implies that, with memory one strategies, allowing players to use mixed actions does not change the maximum payoff they can get in equilibrium. Finally, we described equilibrium strategies profiles, showing that, except for two peculiar cases, they must be or nice, or cautious, or equalizer. Our final result is that the only equilibrium outcomes that can be sustained in reactive strategies, are mutual cooperation and mutual defection.

In the second part of the chapter, we showed that similar results hold also when considering the repeated TG. Specifically, mutual cooperation is the only strictly efficient equilibrium, and extortionate outcomes can never be sustained in equilibrium.

In the last part of the chapter, we applied the model to an asymmetric setting where one long-run second player interacts with a sequence of short-run first players. We call the information available to the short-run player about the long-run player's previous behavior the *reputation* of the long-run player. We show that giving more information to the short-run player is equivalent to expanding the set of possible reputations, and consequently of equilibria, so that also extortionate outcomes become equilibria. Specifically, when only the previous period is disclosed to the short-run player, there is no extortion in equilibrium, and the *best* reputation the second player can build is the fully cooperative one. If instead the first player has access to richer information, the *best* reputation the second player can build is the extortionate one. This finding confirms that mutual cooperation is only one out of the many possible outcomes arising from the repetition of the game. Whether it will also be selected will depend on the type of incentives that players have. In situations like the TG, where the incentive for the long-run second player is to reward the trust often enough, but not always, the long-run player will need a further incentive to be fully cooperative, i.e. reputation is not enough.



# Appendix

## 1.A Zero Determinant strategies with no discounting (Press and Dyson, 2012)

In this section we report the derivation of ZD strategies, as shown in Press and Dyson (2012).

Recall that, when players are using memory-one strategies, we can describe the development of the game over time with a Markov chain with the following transition matrix:

$$M = \begin{pmatrix} p_{CC}q_{CC} & p_{CC}(1 - q_{CC}) & (1 - p_{CC})q_{CC} & (1 - p_{CC})(1 - q_{CC}) \\ p_{CD}q_{CD} & p_{CD}(1 - q_{CD}) & (1 - p_{CD})q_{CD} & (1 - p_{CD})(1 - q_{CD}) \\ p_{DC}q_{DC} & p_{DC}(1 - q_{DC}) & (1 - p_{DC})q_{DC} & (1 - p_{DC})(1 - q_{DC}) \\ p_{DD}q_{DD} & p_{DD}(1 - q_{DD}) & (1 - p_{DD})q_{DD} & (1 - p_{DD})(1 - q_{DD}) \end{pmatrix}$$

We can then find the stationary vector by solving

$$\mathbf{v}M = \mathbf{v} \tag{1.16}$$

Now let  $M' = M - I$ , where  $I$  is the identity matrix. Applying Cramer's rule to  $M'$ , we have:

$$\text{Adj}(M')M' = \det(M')I = 0, \tag{1.17}$$

where  $\text{Adj}(M')$  is the matrix of minors of  $M'$ .

Since, from 1.16, we have  $\mathbf{v}M' = 0$ , it follows that every row of  $\text{Adj}(M')$  is proportional to  $\mathbf{v}$ .

This result does not change if we manipulate  $M'$  by adding the first column to the second and third column, and substituting the last column with a general vector  $\mathbf{f} = (f_1, f_2, f_3, f_4)$ .

We then obtain the following matrix:

$$\hat{M}(\mathbf{f}) = \begin{pmatrix} -1 + p_{CC}q_{CC} & -1 + p_{CC} & -1 + q_{CC} & f_1 \\ p_{CD}q_{CD} & -1 + p_{CD} & (1 - q_{CD}) & f_2 \\ p_{DC}q_{DC} & p_{DC} & -1 + q_{DC} & f_3 \\ p_{DD}q_{DD} & p_{DD} & (1 - q_{DD}) & f_4 \end{pmatrix}$$

And, by 1.16 and 1.17, we know that  $\mathbf{v}$  is proportional to each row of  $\text{Adj}(\hat{M}(f))$ .

Notice that two of the columns of  $\hat{M}(f)$  depend only on one player strategy: the second column of  $\hat{M}(f)$  is  $\tilde{p} = \{p_{CC} - 1, p_{CD} - 1, p_{DC}, p_{DD}\}$  and the third is  $\tilde{q} = \{q_{CC} - 1, q_{CD}, q_{DC} - 1, q_{DD}\}$ . Thus, by setting  $\tilde{\mathbf{p}} = f$ , the first player is able to make  $\det(\hat{M}(f)) = 0$ .

Now, the dot product of any vector  $\mathbf{f}$  with the stationary vector  $\mathbf{v}$  is :

$$\mathbf{v}\mathbf{f} = v_1f_1 + v_2f_2 + v_3f_3 + v_4f_4 \propto d_1f_1 + d_2f_2 + d_3f_3 + d_4f_4,$$

where  $\mathbf{d} = (d_1, d_2, d_3, d_4)$  is the last row of  $\text{Adj}(\hat{M}(f))$ , i.e.  $d_i$  are the determinants of the 3x3 matrices formed by the first three columns of  $\hat{M}(f)$ , leaving out one row at the time.

Then, we can express the dot product of any vector  $\mathbf{f}$  with the stationary vector  $\mathbf{v}$  as:

$$\mathbf{v}\cdot\mathbf{f} = \frac{\det(\hat{M}(f))}{\det(\hat{M}(1))},$$

where the denominator is the proportionality constant to ensure  $\mathbf{v}\cdot\mathbf{1} = 1$ .

Recalling that  $\boldsymbol{\pi}_i$  is the stage game payoff vector for player  $i$ , the total average payoff is

$$\Pi_i = \mathbf{v}\boldsymbol{\pi} = \frac{\det(\hat{M}(\boldsymbol{\pi}_i))}{\det(\hat{M}(1))}.$$

It follows that:

$$a\Pi_1 + b\Pi_2 + c = a \frac{\det(\hat{M}(\boldsymbol{\pi}_1))}{\det(\hat{M}(1))} + b \frac{\det(\hat{M}(\boldsymbol{\pi}_2))}{\det(\hat{M}(1))} + c = \frac{\det(\hat{M}(a\boldsymbol{\pi}_1 + b\boldsymbol{\pi}_2 + c\mathbf{1}))}{\det(\hat{M}(1))}$$

.

We finally arrived to the important point: if the first player sets

$$\tilde{\mathbf{p}} = a\boldsymbol{\pi}_1 + b\boldsymbol{\pi}_2 + c\mathbf{1}, \tag{1.18}$$

we have  $\det(\hat{M}(a\boldsymbol{\pi}_1 + b\boldsymbol{\pi}_2 + c\mathbf{1})) = 0$ , which implies  $a\Pi_1 + b\Pi_2 + c = 0$ , i.e. the first player is able to unilaterally enforce a linear relation between payoffs.

Solutions to 1.18 are called ZD strategies.

For example, by setting  $a = 0$  in 1.18 (i.e. by using an equalizer strategy), the first player is able to enforce  $b\Pi_2 + c = 0$ , i.e. to set the opponent's payoff at  $-\frac{c}{b}$ .

## 1.B Belief-free equilibria (Ely and Valimaki, 2002)

Let  $B^{PD}$  be the set of payoffs profiles such that both players payoff are between the mutual cooperation and the mutual defection payoff (see Figure 1.7):

$$B^{PD} = \{u \in F : P \leq u_i \leq R \text{ for all } i\}$$

Ely and Valimaki (2002) show that any payoff profile inside  $B^{PD}$  can be the outcome of a subgame perfect equilibrium in memory-one strategies:

**Remark 4.** [Result from Section 1 in Ely and Valimaki (2002)] *The payoff profile  $u^*$  can be a subgame perfect, belief-free equilibrium if and only if  $u^* \in B^{PD}$ .*

To construct the equilibrium, Ely and Valimaki (2002) proceed as follows: player  $i$  chooses probabilities s.t., if player  $i$  today chooses  $C$ , his opponent is indifferent between  $C$  and  $D$  and obtains a continuation value of  $X$ , and if player  $i$  today chooses  $D$ , his opponent is indifferent between  $C$  and  $D$  and obtains continuation value of  $Y$ .

Ely, Hörner, et al. (2005) call this type of equilibrium *belief-free* since the optimal action *in one period* do not depend on players' beliefs over which action the opponent will choose in *that period*, so that players are indifferent no matter which action the opponent is currently playing.

Notice that this is a condition for indifference that is different from the usual one of the shot mixed equilibrium: in a belief-free equilibrium, continuation probabilities are s.t. a player is indifferent no matter which action the opponent is currently playing, while the repetition of the one-shot equilibrium does not guarantee that a player is indifferent also *today*, as he would be indifferent today only if the opponent is playing the mixed equilibrium also *today*.

Let  $\rho_i(w)$  be the continuation payoff for player  $i$  if today outcome is  $w$  and, with a little abuse of notation, let  $\pi_i(w)$  be the stage game payoff for player  $i$  corresponding to outcome  $w$ . The total average payoff ( $\Pi_i(w)$ ), if today's outcome is  $w$ , is the weighted average of today's payoff ( $\pi_i(w)$ ) and the continuation payoff relative to  $w$  ( $\rho_i(w)$ ), with weights  $(1 - \delta)$  and  $\delta$ :

$$\Pi_i(w) = (1 - \delta)\pi_i(w) + \delta\rho_i(w)$$

Recalling that  $p_w$  is the probability that the first player cooperates tomorrow if today's outcome is  $w$  (similarly is  $q_w$  for the second player), then  $p_w q_w$  is the probability that tomorrow's outcome is  $CC$  if today's outcome is  $w$ , i.e. the probability that both players cooperate tomorrow if today outcome is  $w$ , and similarly for the other possible outcomes.

Then, we can write the continuation payoff  $\rho_i(w)$  as the average of the total payoffs for each of tomorrow's possible outcomes:

$$\rho_i(w) = p_w q_w \Pi_i(CC) + p_w(1 - q_w) \Pi_i(CD) + (1 - p_w) q_w \Pi_i(DC) + (1 - p_w)(1 - q_w) \Pi_i(DD)$$

In a belief-free equilibrium, if the first player cooperates, the second player total payoff should be fixed, no matter which action he chooses. Thus, the total average payoff of the second player after outcomes  $CC$  and  $CD$ , i.e. after outcomes in which the first player cooperates, should be the same:  $\Pi_2(CC) = \Pi_2(CD) = X$ . Similarly, the second player's payoff after outcomes  $DC$  and  $DD$ , i.e. after outcomes in which the first player defects, should be the same:  $\Pi_2(DC) = \Pi_2(DD) = Y$ . In this case, the continuation payoff of the second player reduces to  $\rho_2(w) = p_w X + (1 - p_w) Y$ , which does not depend on second player's strategy, hence the name *belief-free*.

The total average payoff of the second player if today outcome is  $w$  is then:

$$\Pi_2(w) = (1 - \delta) \pi_2(w) + \delta \rho_2(w) = (1 - \delta) \pi_2(w) + \delta (p_w X + (1 - p_w) Y).$$

This must hold for every  $w \in \{CC, CD, DC, DD\}$ . We have then 4 equations, that we can solve for each  $p_w$  and get:

$$p_w = \frac{\Pi_2(w) - ((1 - \delta) \pi_2(w) + \delta Y)}{\delta(X - Y)} \quad (1.19)$$

where  $\Pi_2(w) = X$  if  $w \in \{CC, CD\}$ , and  $\Pi_2(w) = Y$  if  $w \in \{DC, DD\}$ .

In order for those probabilities to be between 0 and 1, we need:  $R \geq X > Y \geq P$  and  $\delta > \max\{\frac{Y-S}{X-S}, \frac{T-X}{T-Y}\}$ . Moreover, Equations 1.19 imply:  $0 < p_{CC} \leq 1$ ,  $0 < p_{CD}, p_{DC} < 1$ , and  $0 \leq p_{DD} < 1$ .

If the first player is using a strategy as specified by Equations 1.19, the second player payoff is  $p_0(X - Y) - Y$ . At any period of the game, including the first one, if the first player's realized action is C (D), the second player's total average payoff is  $X$  ( $Y$ ), implying that the first player can set the second player's payoff anywhere between  $P$  and  $R$ .

To be subgame perfect, a pair of memory-one strategies must be a Nash equilibrium after each possible outcome  $w \in \{CC, CD, DC, DD\}$ . This is to say that an equilibrium is subgame perfect if it is still a Nash equilibrium when the game starts at  $w$ , and players should play  $p_w$  and  $q_w$ , i.e. a game with  $p_0 = p_w$  and  $q_0 = q_w$ .

If the first player is using a strategy as in Equations 1.19, after any outcome  $w$  the total average payoff of the second player is  $p_w(X - Y) - Y$ , which does not depend on the second player's

strategy. If also the second player is using this type of strategy, neither the first player's payoff depends on the first player's strategy, meaning that those strategies form a subgame perfect Nash equilibrium. This implies that all (and only) payoff profiles in  $B^{PD}$  can be the outcome of this type of equilibrium. Notice that the values of  $X$  and  $Y$  are not required to be the same for the two players.

**Example 5.** It is possible for a player to get the mutual cooperation payoff if the opponent is using strategies as in Equations 1.19, with  $X = R$ . Consider the case where the both players are using strategies as in Equation 1.19. Then, if the first player sets  $X = R$  and starts with  $C$ , strategies with any  $Y < R$  can form a s.p.e. For example, by setting  $Y = P$ , equilibrium probabilities are:

$$\begin{cases} p_{CC} = \frac{X - \delta Y - (1 - \delta)R}{\delta(X - Y)} = 1 \\ p_{CD} = \frac{X - \delta Y - (1 - \delta)T}{\delta(X - Y)} = \frac{\delta(T - P) - (T - R)}{\delta(R - P)} \\ p_{DC} = \frac{Y - \delta Y - (1 - \delta)S}{\delta(X - Y)} = \frac{(1 - \delta)(P - S)}{\delta(R - P)} \\ p_{DD} = \frac{Y - \delta Y - (1 - \delta)P}{\delta(X - Y)} = 0 \end{cases} \quad (1.20)$$

Figure 1.B.1, on the left, shows an example of equilibrium where both players start with  $C$  and then use those kind of strategies. On the right, there is instead an example of an asymmetric equilibrium, where the first player payoff's is fixed at 2, and second player's payoff is fixed at 3. Notice that also this equilibrium can be sustained by several strategies. For example, since we set  $R = 4$  and  $P = 1$ , the first player can fix the second player payoff at 3 by setting  $p_0 = 1$ ,  $X = 3$  and any  $Y < 3$ , or by setting  $p_0 = 0$ ,  $Y = 3$ , and any  $X > 3$ . Similarly for the second player (who sets the first player payoff at 2).




Figure 1.B.1: Belief free equilibria - Examples. Numerical values are  $T = 6$ ,  $S = 0$  and  $\delta = 9/10$ .

The gray square in the figure represent the set  $B^{PD}$ , i.e. the set of payoff that can be sustained in a belief-free equilibrium in memory-one strategies.

## 1.C Proofs

**Lemma 2.** *If  $\delta \geq \max(\frac{T-R}{R-P}, \frac{P-S}{T-P})$ :*

$$L_2 = \left\{ \pi : \left( \pi_i = \frac{S + \delta T}{1 + \delta} \wedge \pi_{-i} = \frac{T + \delta S}{1 + \delta} \right) \vee (\pi_1 = \pi_2 = (1 - \delta)P + \delta R) \right\} \subset B^{PD}$$

*Proof.* Notice that, from Remark 3, strategies in this set have to be pure, as it cannot be that a player is choosing either AllC or AllD and the other player is mixing. Thus, in  $L_2$  two cases are possible: either the equilibrium strategies alternate forever between two states, or the initial state is different from the state in the second period, which is then repeated forever. We shall consider them in turn.

Let us start with strategies that alternate forever between two states. Because of Remark 3, in equilibrium two strategies can only alternate between  $CC - DD$  and  $CD - DC$ . To alternate between  $CC$  and  $DD$  requires that players always cooperate after mutual defection, and always defect after mutual cooperation. This implies  $q_{DD} = p_{DD} = 1$  and  $q_{CC} = p_{CC} = 0$ . Two strategies with these characteristics cannot form an equilibrium. To see this, consider the toughest punishment that players can enforce in this case, that is to set  $p_{CD} = q_{DC} = 0$ . When the first player sets  $p_{DD} = 1$ , the second player has an incentive to choose AllD, which would lead to a path in which players alternate between outcomes  $DD$  and  $CD$ , and clearly this cannot be an equilibrium.

Alternating between states  $CD$  and  $DC$  can be an equilibrium, with corresponding payoff profile  $(\frac{S+\delta T}{1+\delta}, \frac{T+\delta S}{1+\delta})$  (when the first player starts by cooperating) or  $(\frac{T+\delta S}{1+\delta}, \frac{S+\delta T}{1+\delta})$  (when the second player starts by cooperating). This is an equilibrium if players do not have an incentive to defect when they are supposed to cooperate. Considering again the maximum punishment, that is to set  $p_{DD} = q_{DD} = 0$ , we need  $\delta \geq \frac{P-S}{T-P}$ .

Consider now the case in which the initial state is different from the state in the second period, which is then repeated forever. From Observation 1, we know that neither of those states can be  $CD$  or  $DC$ . A situation where the initial state is mutual defection, and from the second period onward there is mutual cooperation can happen, for example, with a pair of WSLS strategies that start by defecting: the induced path is  $DD, CC, CC, \dots$ , which is an equilibrium if  $2R > P + T$  and  $\delta \geq \frac{T-R}{R-P}$ , and which visits only the states  $CC$  and  $DD$ . The opposite situation can never be an equilibrium, as players would defect from the very first round.

Thus, we have  $L_2 = \left\{ \left( \frac{S+\delta T}{1+\delta}, \frac{T+\delta S}{1+\delta} \right), \left( \frac{T+\delta S}{1+\delta}, \frac{S+\delta T}{1+\delta} \right), ((1 - \delta)P + \delta R, (1 - \delta)P + \delta R) \right\} \subset B$ .  $\square$

**Lemma 3.** If  $\delta \geq \max(\frac{P-S}{T-P}, \frac{T-R}{R-S})$ :

$$L_3 = \{\pi : (\pi_i = P \wedge P \leq \pi_{-i} \leq \frac{S + \delta T}{1 + \delta}), \vee (\pi_i = R \wedge \frac{S + \delta T}{1 + \delta} \leq \pi_{-i} \leq R)\} \subset B^{PD}$$

*Proof.* As with  $L_2$ , in  $L_3$  there are several types of equilibrium paths: (i) the first period outcome is different from the second, and the second is different from the third, which then repeats forever; (ii) the first outcome is different from the second, and from the second on it alternates between two other outcomes; finally (iii) all 3 outcomes have positive probability in the long-run. Notice that, while strategies of the first two types has to be pure, strategies of the third type can be mixed.

From Observation 1 we know that there cannot be equilibria in strategies of type (i): since the long-run outcome can be neither  $CD$  or  $DC$ , one of those outcomes have to be visited in the first or the second period (if both outcomes are visited, we would end up in  $L_2$ ). If the outcome  $CD$  is followed by  $DD$ , the player who is supposed to cooperate would defect immediately. If instead the outcome  $CD$  is followed by  $CC$ , it would imply  $p_{CD} = 1$ , which cannot be an equilibrium.

Considering the type of strategies profiles leading to (ii), Observation 1 and Lemma 2 imply that the first period outcome can be either  $CC$  or  $DD$ , followed by an alternation between  $CD$  and  $DC$ . Neither of those can be an equilibrium: if the outcome in the first period is  $CC$ , the player that is supposed to cooperate in the second round has an incentive to defect from the beginning, as this would not change the induced path, but would give him a payoff of  $T$  instead of  $R$  at the very beginning. Similarly, if the outcome in the first period is  $DD$ , the player who is supposed to cooperate in the second period has an incentive to cooperate from the beginning, as this would give him a payoff of  $\frac{S+\delta T}{1+\delta}$ , which, for  $\delta \geq \frac{P-S}{T-P}$ , is greater than  $\frac{(1-\delta^2)P+\delta(S+\delta T)}{1+\delta}$ .

Now we move to the third case, where players might use mixed strategies, and in the long-run only one state is never visited. Let  $w'$  be the that state.

Recall that the transition matrix is :

$$M = \begin{pmatrix} p_{CC}q_{CC} & p_{CC}(1-q_{CC}) & (1-p_{CC})q_{CC} & (1-p_{CC})(1-q_{CC}) \\ p_{CD}q_{CD} & p_{CD}(1-q_{CD}) & (1-p_{CD})q_{CD} & (1-p_{CD})(1-q_{CD}) \\ p_{DC}q_{DC} & p_{DC}(1-q_{DC}) & (1-p_{DC})q_{DC} & (1-p_{DC})(1-q_{DC}) \\ p_{DD}q_{DD} & p_{DD}(1-q_{DD}) & (1-p_{DD})q_{DD} & (1-p_{DD})(1-q_{DD}) \end{pmatrix}$$

Let  $\alpha(w_1, w_2)$  be a generic element of  $M$ , that is the probability to go from state  $w_1$  to state  $w_2$ .

The proof is based on the following conditions, which must hold if in equilibrium the state  $w'$  is never visited:

1.  $\sum_{w_i \neq w'} \alpha(w_i, w') = 0$ : the probability to go to state  $w'$  from any other state must be 0.
2.  $\alpha_{w,w} < 1$  for any  $w \neq w'$ : the long-run path cannot be a repetition of only one state.
3. if  $\alpha_{w_1, w_2} = 1 \rightarrow \alpha_{w_2, w_1} < 1$  for any  $w_1, w_2 \neq w'$ : the long-run path cannot be a repetition of only two states.
4.  $w_0 \neq w'$ : the equilibrium path cannot start in  $w'$
5.  $p_{CD}, q_{DC} < 1$ : a player should never cooperate after being exploited (see Observation 1)
6.  $(p_{DD} = 0 \rightarrow p_{DC} > 0) \wedge (p_{DC} = 0 \rightarrow p_{DD} > 0)$ : in equilibrium no player can use the strategy AllD.
7.  $p_{w'} = q_{w'} = 0$  (assumption on punishment, not a necessary condition): if players deviate from the equilibrium path, and the state  $w'$  is visited, they will use the maximum punishment possible, that is, to defect after  $w'$ . Notice that this condition does not change the equilibrium payoff profiles, since the state  $w'$  is never visited.

Consider for example the case in which the mutual cooperation outcome is never visited, so  $w' = CC$ . From Condition 1, it must be that  $p_{CD}q_{CD} = p_{DC}q_{DC} = p_{DD}q_{DD} = 0$ . Moreover, from Condition 6,  $q_{DD} = 0$  implies  $q_{CD} > 0$ , (otherwise the second player would be always defecting), which in turn implies  $p_{CD} = 0$ . Thus, it must be either  $q_{DD} = p_{CD} = 0$ , or  $p_{DD} = q_{DC} = 0$  (if  $p_{DD} = q_{DD} = 0$  Condition 2 is not satisfied). Moreover,  $q_{CD} = 0$  implies  $p_{DD} = 0$  and  $p_{DC} = 0$  implies  $q_{DD} = 0$ .

We can proceed in a similar way for all other states and, by finding the necessary conditions on strategies to avoid one particular state, we can check whether those can form an equilibrium.

Thus, the conditions that allow a path to visit every state in the long-run except for  $CC$  are:

1.  $q_{DD} = p_{CD} = p_{DC} = 0$
2.  $q_{DD} = p_{CD} = q_{DC} = 0$
3.  $p_{DD} = q_{DC} = p_{CD} = 0$
4.  $p_{DD} = q_{DC} = q_{CD} = 0$



Consider first case 1 (the case 4 is specular). The transition matrix in this case is :

$$M = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & q_{CD} & 1 - q_{CD} \\ 0 & 0 & q_{DC} & 1 - q_{DC} \\ 0 & p_{DD} & 0 & 1 - p_{DD} \end{pmatrix}$$

In this case, the second player has an incentive to set  $q_{DC} = 0$ , that is to avoid the outcome  $CD$  and always defect, and this cannot be an equilibrium.

Consider now the case 3 (the case 2 is specular). The transition matrix is:

$$M = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & q_{CD} & 1 - q_{CD} \\ 0 & p_{DC} & 0 & 1 - p_{DC} \\ 0 & 0 & q_{DD} & 1 - q_{DD} \end{pmatrix}$$

Notice that, if  $p_{DC} = 1$ , the second player best reply would be  $q_{CD} = 1$ , violating Condition 3. If  $p_{DC} = 0$ , condition 6 would be violated, as the first player is choosing AllD. Then, it must be  $0 < p_{DC} < 1$ , i.e. the first player must mix after  $DC$ . This means that the payoff of choosing  $D$  has to be the same as the payoff of choosing  $C$ , given that the opponent is choosing  $D$  (since  $q_{DC} = 0$ ). Recalling that  $\Pi_i(w)$  is the total average payoff of player  $i$  if today's outcome is  $w$ , it must be:

$$\Pi_1(CD) = \Pi_1(DD)$$

This holds if:

$$q_{DD} = \frac{(P - S) - \delta q_{CD}(T - P)}{\delta(T - S)}, \quad q_{CD} < \frac{P - S}{\delta(T - P)}, \quad \delta \geq \frac{P - S}{T - P}$$

Thus, we have:  $0 < q_{DD} < q_{CD} < 1$ , meaning that the second player should mix after both outcomes  $DD$  and  $CD$ . In turn, the second player is indifferent if the payoff of choosing  $D$  is the same as the payoff of choosing  $C$ , given that the opponent is choosing  $D$  (since  $p_{CD} = p_{DD} = 0$ ):

$$\Pi_2(DC) = \Pi_2(DD)$$

which holds if  $p_{DC} = \frac{P-S}{\delta(T-P)}$  and  $\delta \geq \frac{P-S}{T-P}$ .

Thus, we have a first type of equilibrium (we call it  $D\beta$ -eq) that visits 3 states in the long-run, formed by strategies of the type:

$$\mathbf{p} = (p_0, 0, 0, p_{DC}, 0) \quad \mathbf{q} = (q_0, 0, q_{CD}, 0, q_{DD}). \quad (1.21)$$

Notice that, in all those equilibria, the second player's payoff is  $P$ , and the first player payoff is between  $P$  and  $\frac{S+\delta T}{1+\delta}$ , meaning that the first player payoff is bounded from above by the payoff he would get in the equilibrium where states  $CD$  and  $DC$  alternate.

Consider now the case where the state  $DD$  is never visited. Condition 1 implies that it must be:  $(1 - p_{CC})(1 - q_{CC}) = (1 - p_{CD})(1 - q_{CD}) = (1 - p_{DC})(1 - q_{DC}) = 0$ . Moreover, Condition 5, implies  $p_{CD}, q_{DC} < 1$  and  $p_{DC} = q_{CD} = 1$ . Thus, we can have only two cases:

1.  $p_{CC} = p_{DC} = q_{CD} = 1$
2.  $q_{CC} = p_{DC} = q_{CD} = 1$

Consider case 1 (case 2 is specular). The transition matrix in this case is:

$$M = \begin{pmatrix} q_{CC} & 1 - q_{CC} & 0 & 0 \\ p_{CD} & 0 & 1 - p_{CD} & 0 \\ q_{DC} & 1 - q_{DC} & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

As before, Condition 2 implies  $0 < p_{CD} < 1$ , so the first player should get the same payoff when choosing  $C$  and when choosing  $D$ , given that the opponent chooses  $C$  (since  $q_{CD} = 1$ ):  $\Pi_1(CC) = \Pi_1(DC)$ . Solving for  $q_{CC}$  and  $q_{DC}$ , we get:

$$q_{CC} = \frac{-\delta q_{DC}(R - S) + \delta(R - S) + R - T}{\delta(S - T)}, \quad 0 \leq q_{DC} < 1 - \frac{T - R}{\delta(R - S)} \quad \delta > \frac{T - R}{R - S}$$

Thus, it must be  $0 < q_{CC} < 1$ , meaning that the second player is mixing after outcome  $CC$ , implying  $\Pi_2(CC) = \Pi_2(CD)$ . Solving for  $p_{CD}$  we finally have:

$$p_{DC} = 1 - \frac{T - R}{\delta(R - S)}, \quad \delta > \frac{T - R}{R - S}$$

This is the second (and last) type of equilibrium (we call it  $C\beta$ -eq) in the set  $L_3$ , with equilibrium strategies:

$$\mathbf{p} = (p_0, 1, p_{CD}, 1, 0) \quad \mathbf{q} = (q_0, q_{CC}, 1, q_{DC}, 0) \quad (1.22)$$

In all those equilibria, the second player's payoff is fixed at  $R$ , while the first player's payoff is between  $\frac{S+\delta T}{1+\delta}$ , and  $R$ , meaning that the first player's payoff is bounded from below by the payoff he would get in the equilibrium where states  $CD$  and  $DC$  alternate.

Finally, consider the case where the state  $CD$  is never visited (the case where the state  $DC$  is never visited is specular). Condition 1 implies:  $p_{CC}(1 - q_{CC}) = p_{DC}(1 - q_{DC}) = p_{DD}(1 - q_{DD}) = 0$ .

Because of Condition 2, we have  $q_{DC} < 1$ , implying  $p_{DC} = 0$ ,  $p_{DD} > 0$ , and  $q_{DD} = 1$ . Thus, we have again two cases:

1.  $p_{DC} = p_{CC} = 0, q_{DD} = 1$
2.  $p_{DC} = 0, q_{CC} = q_{DD} = 1$

In the first case the transition matrix is:

$$M = \begin{pmatrix} 0 & 0 & q_{CC} & 1 - q_{CC} \\ 0 & 0 & 0 & 1 \\ 0 & 0 & q_{DC} & 1 - q_{DC} \\ p_{DD} & 0 & 1 - p_{DD} & 0 \end{pmatrix}$$

Notice that the second player in this case has an incentive to set  $q_{CC} = q_{DC} = 0$ , meaning that the first player will set  $p_{DD} = 0$ , and this cannot be an equilibrium, as the first player would be always defecting.

In the second case the transition matrix is:

$$M = \begin{pmatrix} p_{CC} & 0 & 1 - p_{CC} & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & q_{DC} & 1 - q_{DC} \\ p_{DD} & 0 & 1 - p_{DD} & 0 \end{pmatrix}$$

Also in this case the second player has an incentive to set  $q_{DC} = 0$ , meaning that he is using the strategy  $\mathbf{q} = (p_0, 1, 0, 0, 1)$ , that is, WSLS. If  $2R > T + P$ , the first player best reply is to set  $p_{CC} = 1$ , and this cannot be an equilibrium in  $L_3$ .

Thus, we have shown that in the set  $L_3$  there are only two types of equilibria: D3-eq and C3-eq, whose payoffs are inside the set  $B$ .

Notice that only the D3-eq are also s.p.e: indeed, if the path reach the state  $DD$ , it is optimal for a player to defect, given that the opponent will defect also. This is not the case with type C3-eq: after the outcome  $CC$ , a player might have an incentive to deviate and suffer a lower payoff for one period, but ensuring that the path does not get stuck in mutual defection. If  $\delta > \frac{P-S}{R-S}$ , the second player always have an incentive to cooperate instead of punishing the first player, implying that C3-eq are not s.p.e.  $\square$

**Lemma 4.** If  $\delta \geq \max\{\frac{T-R}{T-P}, \frac{P-S}{R-S}\}$ :

$$L_4 = B^{PD} - \{(RR), (PP)\}$$

*Proof.* First notice that the first part of the proof of lemma 3 implies that also in the set  $L_4$  there are no equilibria in pure strategies. Moreover, we know that equilibria in  $L_4$ , since they visit every state, must be s.p.e., and that belief-free equilibria a la Ely are all in the set  $B$ . Thus, if we show that an equilibrium in  $L_4$  are belief-free, automatically we prove also that the set of equilibrium payoffs is a subset of  $B$ . Recall that, in a belief-free equilibrium, it must hold  $\Pi(CC) = \Pi(DC)$  and  $\Pi(CD) = \Pi(DD)$  (we dropped the subscripts from  $\Pi_i$  for convenience), so that a player is indifferent between his two actions  $C$  and  $D$ , no matter which action the opponent is currently choosing. Moreover, since in the set  $L_4$  equilibria cannot be pure, in equilibrium at least one player is mixing after at least one state.

Let  $w'$  be the state after which the first player is mixing.

Since the first player is mixing after  $w'$ , he must be indifferent between playing  $C$  or  $D$ .

The payoff from playing  $C$  after outcome  $w'$  is:  $q_{w'}\Pi(CC) + (1 - q_{w'})\Pi(CD)$ , while the payoff from playing  $D$  is:  $q_{w'}\Pi(DC) + (1 - q_{w'})\Pi(DD)$ .

Those payoffs are equal if:  $q_{w'}(\Pi(CC) - \Pi(CD) + \Pi(DD) - \Pi(DC)) = \Pi(DD) - \Pi(CD)$ .

Notice that solving this expression for  $q_{w'}$  would give a probability outside the interval  $[0, 1]$ , since the PD does not have a one shot mixed equilibrium.

So, the payoffs are equal if either  $\Pi(DD) = \Pi(CD)$  and  $\Pi(CC) = \Pi(DC)$ , or if  $\Pi(DD) = \Pi(CD)$  and  $q_{w'} = 0$ , or if  $\Pi(CC) = \Pi(DC)$  and  $q_{w'} = 1$ .

Dutta and Siconolfi (2010) stops here, by noticing that in a totally mixed equilibrium only the first condition is relevant, which is exactly the condition for a belief-free equilibrium. When we allow the equilibrium to not be totally mixed, the last two conditions are the relevant ones.

Assume first that  $\Pi(CC) = \Pi(DC)$  and  $q_{w'} = 1$ . Then, for any  $w$ , the first player will choose  $C$  after  $w$  whenever  $q_w(\Pi(DD) - \Pi(CD)) > \Pi(DD) - \Pi(CD)$ .

Let  $\Pi(DD) > \Pi(CD)$ . This implies that the first player will choose  $C$  whenever  $q_w > 1$ , that is, never, unless  $q_w = 1$ , in which case he is indifferent. Thus, after any  $w$ , either  $q_w < 1$  and  $p_w = 0$ , or  $q_w = 1$ . Since this must hold for any  $w$ , this implies that the probability to go to the outcome  $CD$  is zero, which contradicts our assumption that every state is visited.

This same reasoning can be applied to all other cases, and the conclusion is that, if in equilibrium at least one player is mixing, then either the equilibrium is belief-free, or not every state is visited, meaning that the resulting equilibrium payoff would not be in the set  $L_4$ .

Let  $\Pi(DD) < \Pi(CD)$ . This implies that the first player will choose  $C$  whenever  $q_w < 1$ . Thus, after any  $w$ , either  $q_w < 1$  and  $p_w = 1$ , or  $q_w = 1$ . Since this must hold for any  $w$ , this implies that the probability to go to the outcome  $DD$  is zero, which contradicts our assumption that every state is visited.

Now assume that  $\Pi(DD) = \Pi(CD)$  and  $q_{w'} = 0$ . Then, for any  $w$ , the first player will choose  $C$  after  $w$  whenever  $q_w(\Pi(CC) - \Pi(DC)) > 0$ .

Let  $\Pi(CC) > \Pi(DC)$ . This implies that the first player will choose  $C$  whenever  $q_w > 0$ . Thus, after any  $w$ , either  $q_w > 0$  and  $p_w = 1$ , or  $q_w = 0$ . Since this must hold for any  $w$ , this implies that the probability to go to the outcome  $DC$  is zero, which contradicts our assumption that every state is visited.

Let  $\Pi(CC) < \Pi(DC)$ . This implies that the first player will choose  $C$  whenever  $q_w < 0$ , i.e. never, unless  $q_w = 0$ . Thus, after any  $w$ , either  $q_w > 0$  and  $p_w = 0$ , or  $q_w = 0$ . Since this must hold for any  $w$ , this implies that the probability to go to the outcome  $CC$  is zero, which contradicts our assumption that every state is visited.

Thus, we have that, if  $\Pi(CC) = \Pi(DC)$ , it must be  $\Pi(CC) = \Pi(DC)$ , and if  $\Pi(CC) = \Pi(DC)$ , it must be  $\Pi(CC) = \Pi(DC)$ .

This, in turn, implies that, if in a subgame perfect equilibrium at least one player is mixing, and all states are visited with positive probability, the equilibrium must be belief-free.  $\square$

**Proposition 4.** *Let  $s^* = (\mathbf{p}, \mathbf{q})$  be a profile in memory-one strategies. Then  $s^*$  is a NE if and only if both strategies are either:*

- *nice:*  $p_0 = p_{CC} = q_0 = q_{CC} = 1$
- *cautious:*  $p_0 = p_{DD} = q_0 = q_{DD} = 0$
- *alternators:*  $p_0 = p_{CD} = q_{DC} = 0, q_0 = p_{DC} = q_{CD} = 1$
- *equalizer:* (See Equations 1.6)
- *almost equalizer:* (See Equations 1.22 and 1.21)

*Proof.* Recall that in the set  $L_1$  there can only be pairs of strategies that are either nice or cautious (notice that Generous and Defective equilibria belong to this set, while Equalizer strategies can belong to any set except for  $L_3$ ). In the set  $L_2$  there are only strategies that

alternate between outcomes  $CD$  and  $DC$ . In the set  $L_3$  we found only two equilibria, in *almost* equalizer strategies. In those equilibria, one player is indifferent after one outcome, and the opponent is indifferent after the other two outcomes. Finally, in the set  $L_4$ , there can only be Equalizer strategies: as we saw in the proof of Lemma 4, strategies in this set should form a belief-free equilibrium. A belief-free equilibrium is possible only if a player is indifferent among all his strategies, no matter which action the opponent is currently using. This implies that the relation between players' payoffs must be linear. Ichinose and Masuda (2018) show that, among all memory-one strategies, only ZD and unconditional strategies can enforce a linear relation between players payoff. This implies that, in the PD, only Equalizer strategies are able to fix the opponent payoff, and thus only Equalizer strategies can form equilibria in the set  $L_4$ .  $\square$

**Proposition 5.** *Let  $\mathbf{p}$  be a ZD strategy. Then  $\mathbf{p}$  is part of a SPNE if and only if is an Equalizer strategy.*

*Proof.* (if) Recall that, in a belief-free equilibrium,  $X_i$  is the total average payoff for player  $i$  if the opponent chooses  $C$  in the current round, and  $Y_i$  is the total average payoff in the case the opponent chooses  $D$  in the current round. One can notice, recalling that equalizer strategies are of the form as in Equations 1.7, that strategies in a belief-free equilibrium are ZD-equalizer strategies with  $k = p_0(X_1 - Y_1) + Y_1$ , and  $\phi = \frac{1-\delta}{X_1 - Y_1}$ .

Alternatively, we can say that any pair of equalizer strategies forms a belief free equilibrium where the second player's total payoff in each period is  $X = k + \frac{(1-p_0)(1-\delta)}{\phi}$  if the first player chooses C, and  $Y = k - \frac{(1-\delta)p_0}{\phi}$  if the first player chooses D. This implies that every pair of Equalizer strategies form a subgame perfect Nash equilibrium.

(only if) This is not the case with either generous or defector strategies. To see why, consider first an equilibrium in generous strategies, so  $s_1 = ZD(\Phi_1, R)$  and  $s_2 = ZD(\Phi_2, R)$ . The total average payoff for the first player after any outcome  $w \neq CC$ , given that from tomorrow both players will follow their equilibrium strategies, is:

$$\Pi_1(w) = R - \frac{(1-\delta)(\Phi_1\phi_1(1-q_w) + \phi_2(1-p_w))}{\phi_1\phi_2(1-\Phi_1\Phi_2)}$$

In this case, for any value of  $q_w, \phi_i, \Phi_i$ , the payoff of setting  $p_w = 1$  is always higher than the payoff of setting  $p_w < 1$ , meaning that in any period the first player is willing to play C even when he is supposed to mix, and thus, this cannot be a subgame perfect equilibrium (a similar argument applies to defector strategies).  $\square$

**Proposition 6.** *Let  $s^* = (\mathbf{p}, \mathbf{q})$  be a profile in memory-one strategies, in which the outcome is different than mutual cooperation or mutual defection. Then  $s^*$  is a SPNE if and only if  $\mathbf{p}$  and  $\mathbf{q}$  are Equalizer strategies.*

*Proof.* Recall that in the set  $L_1$  there are pairs of strategies that are either nice or cautious, and that the only payoffs inside this set are the mutual defection and the mutual cooperation payoff profiles. Even if Generous and Defective strategies are not SPNE, there are Equalizer strategies that are in the set  $L_1$ , i.e. those with  $k = P$  or  $k = R$ . Certain Equalizer strategies belong to the set  $L_2$ , where players alternate between outcomes  $CD$  and  $DC$ , and they punish deviations with probabilities in order to make the opponent indifferent between deviating, or sticking to the equilibrium. Moreover, this is the only SPNE in the set  $L_2$ . In the set  $L_3$  there are only the two types of equilibria with strategies as specified in 1.22 and 1.21. Only the first type of equilibrium, which never visit the state  $CC$ , is also a SPNE. Finally, the proof of Proposition 5 implies that every strategy in a belief-free equilibrium is an Equalizer strategy. Since equilibria in the set  $L_4$  must be belief-free, it follows that, among all memory-one strategies, only Equalizer can form SPNE in the set  $L_4$ .  $\square$

**Proposition 7.** *If players are constrained to reactive strategies,  $\delta < 1$  and  $P - S \neq T - R$ , then the only payoff profiles that can be sustained in a Nash equilibrium are  $(R, R)$  and  $(P, P)$ . If  $P - S = T - R$ , the set of equilibrium payoff profiles is  $B^{PD}$ .*

*Proof.* First notice that, with reactive strategies, the set  $L_2$  is empty: one equilibrium path ( $DD$ , and then  $CC$  forever) implies  $p_D = 1$ , which cannot be an equilibrium. The other equilibrium path (which alternates between  $CD$  and  $DC$ ), implies  $p_0 = p_C = q_C = 1$  and  $q_0 = p_D = q_D = 0$ , i.e. those are TtT strategies starting from different actions. As long as  $2R > T + S$ , this cannot be an equilibrium, as players have an incentive to deviate to the mutual cooperation equilibrium. Now consider the set  $L_4$ . The proof of Proposition 4 shows that strategies in this set can only be Equalizer, but we saw from equations 1.1 that, unless the game has equal gains from switching, that is unless  $P - S = T - R$ , reactive strategies are never Equalizer, which implies that the set  $L_4$  is empty. Thus, if  $P - S \neq T - R$ , the set of Nash equilibria payoffs in the PD with reactive strategies is  $\mathcal{E} = \{(R, R), (P, P)\}$ , and the only subgame perfect equilibrium is AllD. If instead  $P - S = T - R$ , every ZD strategy is reactive, (Hilbe, Nowak, and Traulsen, 2013) implying that the results with general memory-one strategies hold also in this case.  $\square$

**Lemma 6.** If  $\delta \geq \frac{T-R}{R-P}$

$$L_2 = \{\pi : (\pi_1 = \pi_2 = (1 - \delta)P + \delta R)\}$$

*Proof.* In  $L_2$  two cases are possible: either two states are visited in the long-run, or the initial state is different from the state in the second period, which is then repeated forever. We shall consider them in turn.

Let us start with strategies that visit two states in the long-run. Proposition 10 implies that those states cannot be  $CC$  and  $CD$ , as this would imply  $p_{CD} = 1$ , and the second player would have an incentive to choose D after  $CC$ .

If the states are  $DD$  and  $DC$ , then the first player is always defecting, the resulting payoff profile is  $(P, P)$ . Nonetheless, this cannot be an equilibrium, since, if  $q_{DD} = 1$ , the first player has an incentive to choose C after  $DD$ .

If the states are  $CC$  and  $DC$ , the second player is always cooperating, i.e.  $q_{CC} = q_{DC} = 1$ . This cannot be an equilibrium, as the first player should play D after mutual defection, while he has an incentive to play C.

Finally, if the two states are  $CD$  and  $DC$ , or  $CC$  and  $DD$ , it must be that players are using pure strategies. None of those cases can be an equilibrium: in the first case, the first player payoff would be lower than P. In the second case, the second player has an incentive to choose D also after  $DD$  (see proof of Lemma 2).

Consider now the case in which the initial state is different from the state in the second period, which is then repeated forever. From Observation 3 and Lemma 5, we know that neither of those states can be  $CD$  or  $DC$ . As with the PD, there is an equilibrium where the initial state is mutual defection, and from the second period onward there is mutual cooperation, if  $2R > P + T$  and  $\delta \geq \frac{T-R}{R-P}$ , and which visits only the states  $CC$  and  $DD$ . The opposite situation instead can never be an equilibrium, since the second player would defect from the very first round.

Thus, we have  $L_2 = \{\pi_i = \pi_{-i} = (1 - \delta)P + \delta R\}$ . □

**Lemma 7.**

$$L_3 = \emptyset$$

*Proof.* As with  $L_2$ , in  $L_3$  there are several types of equilibrium paths: (i) the first period outcome is different from the second, and the second is different from the third, which then



repeats forever; (ii) the first outcome is different from the second, and from the second on two outcomes are visited in the long-run; finally (iii) all 3 outcomes have positive probability in the long-run.

From lemma 6, we know that there cannot be equilibria of type (i) and (ii): in the first case, the long-run outcome can only be  $CC$ . Then,  $CD$  and/or  $DC$  have to be visited in the first or the second period.

If  $CD$  is visited in the second period, it must be  $p_{CD} = 1$ , which cannot be an equilibrium. If  $DC$  is visited in the second period, the first player would have an incentive to choose  $C$  (instead of  $D$ ) in the second period, to anticipate the stream of mutual cooperation payoffs.

So, it must be that the state  $DD$  is visited in the second period. But then, in the first period outcomes can only be  $CD$  or  $DC$ . But then, the first player would defect immediately if the outcome is  $CD$ , and he would cooperate immediately if the outcome is  $DC$ .

Considering the type of strategies profiles leading to (ii), lemma 2 implies that there are no equilibria where in the long-run two states are visited, thus there are no equilibria of this type.

Finally, the third type of equilibrium is the one where one state is never visited in the long-run. Specularly with the proof of 3, let  $w'$  be the that state.

As with the proof of Lemma 3, we consider the possible cases one by one. Since the transition matrix is the same as in the PD, conditions 1-6 in the proof of Lemma 3 hold also in the TG. Condition 5 is instead modified to take into account the observation 1: if in an equilibrium 3 states are visited in the long-run, it must be:  $p_{CD}, q_{CD} < 1$ .

We start by noticing that, if  $w' = CC$ , the first player payoff would be below  $P$ , and this cannot be an equilibrium.

Consider now the case  $w' = DD$ , so the state  $DD$  is never visited. Then, one of the following conditions must hold:

1.  $p_{DC} = p_{CD} = q_{CC} = 1$
2.  $q_{DC} = q_{CD} = p_{CC} = 1$
3.  $p_{DC} = q_{CD} = q_{CC} = 1$
4.  $p_{DC} = q_{CD} = p_{CC} = 1$

If  $q_{CC} = 1$ , the best reply of the first player is to set  $p_{CC} = 1$ , so that only the outcome  $CC$  is visited in the long-run.

If  $q_{CD} = 1$ , the best reply of the first player is to set  $p_{CD} = 1$ , which can never be an equilibrium.

Consider now the case  $w' = CD$ , i.e. the case where the state  $CD$  is never visited. Then, one of the following conditions must hold:

1.  $q_{DC} = 1, p_{CC} = 0, p_{DD} = 0$
2.  $q_{DC} = 1, p_{CC} = 0, q_{DD} = 1$
3.  $p_{DC} = 0, p_{CC} = 0, q_{DD} = 1$
4.  $p_{DC} = 0, q_{CC} = 1, q_{DD} = 1$

Case 1: if  $q_{DC} = 1$ , the first player will set  $p_{DC} = 1$ , and the second player has an incentive to set  $q_{CC} = q_{DD} = 1$ , meaning that  $p_{CC} = 0$  cannot be a best reply for the first player.

$$M = \begin{pmatrix} 0 & 0 & q_{CC} & 1 - q_{CC} \\ 0 & 0 & 0 & 1 \\ p_{DC} & 0 & 1 - p_{DC} & 0 \\ 0 & 0 & q_{DD} & 1 - q_{DD} \end{pmatrix}$$

Case 2: if  $q_{DC} = q_{DD} = 1$ , the best reply for the first player is to set  $p_{DC} = p_{DD} = 1$ , but then the second player would have an incentive to defect after outcomes  $DC$  and  $DD$ . Thus, this cannot be an equilibrium.

$$M = \begin{pmatrix} 0 & 0 & q_{CC} & 1 - q_{CC} \\ 0 & 0 & 0 & 1 \\ p_{DC} & 0 & 1 - p_{DC} & 0 \\ p_{DD} & 0 & 1 - p_{DD} & 0 \end{pmatrix}$$

Case 3: if  $q_{DD} = 1$ , the first player has an incentive to set  $p_{DD} = 1$ , and if  $p_{DC} = p_{CC} = 0$ , the second player has an incentive to defect after outcomes  $CC$  and  $DC$ , so he will set  $q_{CC} = q_{DC} = 0$ . But this means that players alternate between outcomes  $CC$  and  $DD$ , which cannot be an equilibrium.

$$M = \begin{pmatrix} 0 & 0 & q_{CC} & 1 - q_{CC} \\ 0 & 0 & 0 & 1 \\ 0 & 0 & q_{DC} & 1 - q_{DC} \\ p_{DD} & 0 & 1 - p_{DD} & 0 \end{pmatrix}$$

Case 4: if  $q_{CC} = q_{DD} = 1$ , the first player would set  $p_{CC} = 1$ , and this cannot be an equilibrium in  $L_3$ .

$$M = \begin{pmatrix} p_{CC} & 0 & 1 - p_{CC} & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & q_{DC} & 1 - q_{DC} \\ p_{DD} & 0 & 1 - p_{DD} & 0 \end{pmatrix}$$

Consider finally the case  $w' = DC$ , i.e. the case where the state  $DC$  is never visited. Then, one of the following conditions must hold:

1.  $p_{CD} = 1, q_{CC} = 0, q_{DD} = 0$
2.  $p_{CD} = 1, q_{CC} = 0, p_{DD} = 1$
3.  $q_{CD} = 0, q_{CC} = 0, p_{DD} = 1$
4.  $q_{CD} = 0, p_{CC} = 1, p_{DD} = 1$

Cases 1 and 2 cannot be equilibria, since  $p_{CD} = 1$ .

Consider now Case 3. The first player has an incentive to set  $p_{CD} = p_{DC} = 0$ , that is to avoid the outcome  $CD$ . The second player then has an incentive to set  $q_{DD} = 0$ , that is to defect after outcome  $DD$ . But if  $q_{CD} = q_{DD} = 0$ , the second player is always defecting, and this cannot be an equilibrium inside  $L_3$ .

$$M = \begin{pmatrix} 0 & p_{CC} & 0 & 1 - p_{CC} \\ 0 & p_{CD} & 0 & 1 - p_{CD} \\ 0 & 0 & 0 & 1 \\ q_{DD} & 1 - q_{DD} & 0 & 0 \end{pmatrix}$$

Finally, consider the case 4. Also in this case the first player has an incentive to set  $p_{CD} = 0$ . But then, if  $\delta \geq \frac{T-R}{R-P}$ , the second player has an incentive to set  $q_{CC} = q_{DD} = 1$ , so neither this can be an equilibrium inside  $L_3$ .

$$M = \begin{pmatrix} q_{CC} & 1 - q_{CC} & 0 & 0 \\ 0 & p_{CD} & 0 & 1 - p_{CD} \\ 0 & 0 & 0 & 1 \\ q_{DD} & 1 - q_{DD} & 0 & 0 \end{pmatrix}$$

It follows that the set  $L_3$  is empty. □

**Lemma 8.** If  $\delta \geq \max\{\frac{T-R}{T-P}, \frac{P-S}{R-S}\}$ ,

$$L_4 = \{\pi \in F : \pi_1 = P \text{ and } P < \pi_2 \leq R\}$$

*Proof.* First notice that the first part of the proof of Lemma 7 implies that also in the set  $L_4$  there are no equilibria in pure strategies. Thus, in equilibrium at least one player is mixing after at least one state. Let  $w'$  be the state after which the first player is mixing.

Since the first player is mixing after  $w'$ , he must be indifferent between playing  $C$  or  $D$ .

We can write the payoff from playing  $C$  after outcome  $w'$  as:  $q_{w'}\Pi_1(CC) + (1 - q_{w'})\Pi_1(CD)$ , and the payoff from playing  $D$  as:  $q_{w'}\Pi_1(DC) + (1 - q_{w'})\Pi_1(DD)$ . Those payoffs are equal if:

$$q_{w'}(\Pi_1(CC) - \Pi_1(CD) + \Pi_1(DD) - \Pi_1(DC)) = \Pi_1(DD) - \Pi_1(CD)$$

If  $\Pi_1(CC) - \Pi_1(CD) + \Pi_1(DD) - \Pi_1(DC) \neq 0$ , we can solve for  $q_{w'}$  and get:

$$q_{w'} = \frac{\Pi_1(DD) - \Pi_1(CD)}{\Pi_1(CC) - \Pi_1(CD) + \Pi_1(DD) - \Pi_1(DC)}$$

Notice that  $q_{w'}$  does not depend on  $w'$ , meaning that the second player is using an unconditional strategy, and we know that this strategy can only be his Stackelberg strategy,  $q^S$ . If  $q = q^S$ , the second player is mixing after every state, so he must be indifferent after every state. We can then apply the same reasoning as in the proof of Lemma 4, and prove that it must hold  $\Pi_2(DD) = \Pi_2(DC)$  and  $\Pi_2(CC) = \Pi_2(CD)$ , i.e. the first player must use a *belief-free* strategy.

If  $\Pi_1(CC) - \Pi_1(CD) + \Pi_1(DD) - \Pi_1(DC) = 0$ , then it must hold  $\Pi_1(DD) = \Pi_1(CD)$  and  $\Pi_1(CC) = \Pi_1(DC)$ , meaning that the second player must use a belief-free strategy.

However, any belief-free strategy is an equalizer strategy (see the proof of Proposition 5), so that, if an equalizer strategy do not exists, neither a belief-free strategy exists. It directly follows that, if  $\delta < 1$ , there are no belief-free strategies that the second player can use, so all equilibria in  $L_4$  involve the unconditional Stackelberg strategy of the second player, and a *belief-free* memory-one strategy for the first player. If  $\delta = 1$ , the second player can also use an Equalizer strategy that fixes the opponent's payoff at  $P$ .

This implies that in any equilibrium, the first player payoff is fixed at  $P$ , while the second player payoff can be anywhere between  $P$ (excluded) and  $R$ (included):

$$L_4 = \{\pi \in F : \pi_1 = P \text{ and } P < \pi_2 \leq R\} \subset B$$

□

**Proposition 13.** *Let  $(\hat{f}, \hat{q})$  be a stationary equilibrium. Then either  $\hat{q} < q^S$  and  $\underline{q}(\hat{f}) = \emptyset$ , or  $\hat{q} \geq q^S$  and  $\underline{q}(\hat{f}) = \hat{q}$ .*

*Proof.* Since in equilibrium each short-run player is playing a best reply, either  $\hat{f}(\hat{q}) = 0$  or  $\hat{f}(\hat{q}) = 1$ . We shall first consider the case in which there is at least a probability of playing Reward that induces the first player to play Trust, that is  $\underline{q}(\hat{f}) \in [0, 1]$ . First notice that in equilibrium it must be the case that  $\hat{q} = \underline{q}(\hat{f})$ . This is an immediate consequence of the fact that the long-run player chooses a best response to the choices made by the short-run player. On the other hand, since  $\hat{f}(\hat{q}) \in BR_1(\hat{q})$ ,  $\hat{f}(\hat{q}) = 1$  implies that  $\hat{q} \geq q^S$ . It remains to consider the case in which  $\underline{q}(\hat{f}) = \emptyset$ . Since the short-run player will never cooperate, no matter what the strategy chosen by the long-run player may be, playing any strategy  $\hat{q} \in [0, 1]$  yields the same payoff to the long-run player. If  $\underline{q}(\hat{f}) = \emptyset$  to be an equilibrium it must be the case that  $\hat{q} < q^S$ . This completes the proof.  $\square$

## 1.D Reactive strategies with no discounting (Baklanov, 2018)

In this section we report the case of  $\delta = 1$ , as considered in Baklanov (2018). When  $\delta = 1$ , Lemma 0 implies that the limit probabilities to cooperate in round  $t$  for  $t \rightarrow \infty$  are:

$$\begin{cases} \bar{p} = v_{CC} + v_{CD} = p_C(v_{CC} + v_{DC}) + p_D(v_{CD} + v_{DD}) = p_D + (p_C - p_D)\bar{q} \\ \bar{q} = v_{CC} + v_{DC} = q_C(v_{CC} + v_{CD}) + q_D(v_{DC} + v_{DD}) = q_D + (q_C - q_D)\bar{p} \end{cases} \quad (1.23)$$

So that:

$$\bar{p} = \frac{p_D + q_D(p_C - p_D)}{1 - (p_C - p_D)(q_C - q_D)} \quad \text{and} \quad \bar{q} = \frac{q_D + p_D(q_C - q_D)}{1 - (p_C - p_D)(q_C - q_D)} \quad (1.24)$$

Nowak and Sigmund (1988) show that when  $\delta = 1$  the limit probabilities that players cooperate,  $\bar{p}$  and  $\bar{q}$ , are independent. Considering the outcome CC, this means that the limit probability that both players cooperate is the product of the limit probabilities that each player cooperate, i.e.  $v_{CC} = \bar{p}\bar{q}$ , and similarly for the other outcomes.

Players payoffs are then:

$$\Pi_1 = \bar{p}\bar{q}R + \bar{p}(1 - \bar{q})S_i + (1 - \bar{p})\bar{q}T + (1 - \bar{p})(1 - \bar{q})P.$$

$$\Pi_2 = \bar{p}\bar{q}R + \bar{p}(1 - \bar{q})T + (1 - \bar{p})\bar{q}S + (1 - \bar{p})(1 - \bar{q})P.$$

Notice that we can use Equations 1.23 and write the payoff of the first player as a function only of  $\bar{p}$ ,  $q_C$  and  $q_D$  (and similarly for the second player). This function can then be maximized w.r.t.  $\bar{p}$ , to find the first player's best reply to any  $q_C$  and  $q_D$ . Baklanov (2018) noticed that in this case the set of totally mixed Nash equilibria is characterized by:

**Remark 9.** (*Totally mixed Nash equilibria in reactive strategies when  $\delta = 1$ , Theorem 2 in Baklanov (2018)*)

$$\left\{ \begin{array}{l} \frac{D\Pi_1}{\bar{p}} = 0, \frac{D\Pi_2}{\bar{q}} = 0, \frac{D_2\Pi_1}{\bar{p}} \leq 0, \frac{D_2\Pi_2}{\bar{q}} \leq 0 \\ \bar{p} = \frac{p_D + q_D(p_C - p_D)}{1 - (p_C - p_D)(q_C - q_D)}, \bar{q} = \frac{q_D + p_D(q_C - q_D)}{1 - (p_C - p_D)(q_C - q_D)} \\ 0 < \bar{p} < 1, 0 < \bar{q} < 1 \end{array} \right. \quad (1.25)$$

In those equilibria, players are using mixed strategies, which are not belief-free, as they are indifferent among only a subset of their strategy space, unless  $T - R = P - S$ , in which case all equilibrium strategies are equalizer strategies<sup>23</sup>.

If  $T - R > P - S$ , it is possible that a player gets more than the mutual cooperation payoff. The maximum payoff that the first player can get corresponds to the equilibrium with the minimum (maximum) probability that the first (second) player cooperates:

$$(\bar{p}^{b1}, \bar{q}^{b1}) = \left( \frac{P - S}{P + R - 2S}, \frac{T - P}{2T - R - P} \right)$$

Over all strategies that allow players to reach  $(\bar{p}^{b1}, \bar{q}^{b1})$ , i.e. that satisfy Equations 1.24, the pair of strategies that are also best reply to each other, and thus form a Nash equilibrium, are:

$$(p_C, p_D) = \left( \frac{(P - S)(2T - P - R)}{(T - P)(P + R - 2S)}, 0 \right) \quad \text{and} \quad (q_C, q_D) = \left( 1, \frac{(R - P)(T - S)}{(R - S)(2T - R - P)} \right)$$

In this equilibrium, the first player's payoff, which is also the maximum payoff he can get in an equilibrium, can be higher than  $R$ , but it is strictly lower than his Stackelberg payoff. Indeed, to reach the Stackelberg payoff we need  $q_D = 1$ , but  $1 - q_D = \frac{(T - R)(P + R - 2S)}{(R - S)(2T - R - P)} > 0$

---

<sup>23</sup>Notice that, if  $p_C = p_D = p$  is a Nash equilibrium of the one-shot game, all equilibrium conditions are trivially satisfied.

## Chapter 2

# Be nice, but not always: Extortion in the Lab

with Luciano Andreozzi and Marco Faillo (University of Trento)

In this chapter we report the results from an experiment designed to assess the impact that disclosing information about players' previous behavior might have on equilibrium selection in a repeated Trust Game. From a theoretical point of view, the model developed in Chapter One implies that we should observe more fully cooperative behavior in a setting in which minimal information (i.e. about only the previous choice of the opponent) is available, compared to a setting in which richer information (i.e. the entire history of play) is disclosed. In the latter case, also *extortionate* profiles (i.e. profiles in which one player can get more than the mutual cooperation payoff) can be sustained in equilibrium, so we expect less-than-fully cooperative strategies to be more profitable. Our experimental results are broadly in line with these predictions. In a setting of minimal information, no player is able to gain more than the mutual cooperation payoff. In a setting with richer information, there is a significant number of second players who were nice, but not always, and managed to "exploit" the first players' trust. Overall, our data show that increasing the amount of information was effective in allowing first players to discriminate between mild abusers (who play Reward frequently enough to make Trust the optimal strategy) and hard abusers (who mostly play Not Reward). However, first players seem unwilling to tolerate a high level of abuse of trust, even when it would be in their interest to do so.

## 2.1 Introduction

The Trust Game has been widely studied in economics since it is suitable for representing several situations in which market failures are likely to appear. A consumer buying a product before knowing its quality, a bank giving a loan before knowing whether it will be paid back, or an employer hiring a worker before knowing worker's dedication to the job (Kreps, 1996).

Those settings share the same basic structure:

at the beginning of the game, the first player has to decide whether or not to trust the second player. If the first player chooses Not Trust, the game ends and both players get their outside option,  $P$ . If the first player chooses Trust, the second player has to decide whether to reward her trust or not. If he chooses Reward, both players get  $R > P$ ; if he chooses Not Reward, the second player gets  $T > R$ , and the first player gets  $S < P$ .

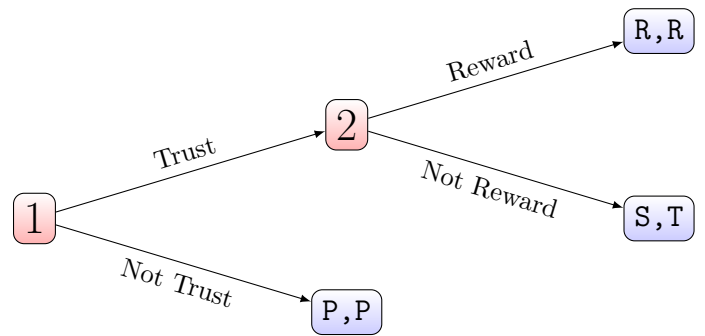


Figure 2.1: The Trust Game

This creates a problem of incentives. Compared to the outside option, both players are better off when trust is rewarded; however, the second player has an incentive to choose Not Reward. Anticipating this, the first player chooses Not Trust, leading to the inefficient outcome (Not Trust, Not Reward). Some incentives has to be given to the second player if the efficient (Trust, Reward) outcome is to be reached. One possibility would be to change the payoffs associated to the different actions. For example, a fine may be introduced in case of not rewarding, or a prize in case of rewarding. One may wonder whether the outcome (Trust, Reward) can be obtained while maintaining the basic structure of the game. This is important because in many situations the possibility to directly punish or award the counterpart is limited. In everyday life, trust is commonly given and rewarded even when explicit incentives are virtually non existent.

Several models have been proposed to explain cooperation in the Trust Game. When the game is repeated, reputation concerns may give to the second player the right incentives to play Reward (Kreps and Wilson, 1982; Mailath and Samuelson, 2001). Notice that repetition *per se* will not solve the problem. It is a well know result that when a game with a single NE is repeated for a finite number of rounds, no new NE is formed which is different from a repetition of the stage game NE. This is a straightforward implication of backward induction.



When the end of the game is *unknown*, players cannot use backward induction, and the Folk Theorem applies: every pair of payoffs that Pareto-dominates the outside option can be obtained (or approximated) in equilibrium, if the probability of going to the next period is sufficiently large. The question of whether the mutual cooperation outcome can be sustained in equilibrium becomes then a problem of equilibrium selection.

Indeed, most of the literature has focused on strategies that are able to generate a stream of Trust and Reward. This type of equilibrium is sustained by the belief that the first player will stop playing Trust after having observed a single deviation from Reward. For example, reputation models in Fudenberg and Levine (1989) or in Ely and Välimäki (2003), assume that there can be different *types* of second players, which are committed to play a certain strategy. If the first player's belief that the second player is a rewarding-type is high enough (i.e. larger than 0), the efficient outcome (Trust, Reward) can be sustained in equilibrium.

The experimental literature has followed a similar path. Several laboratory experiments have found a general tendency of subjects to play Trust and Reward more often when the game is repeated, compared to one-shot settings (Bornhorst et al., 2004; Kanagaretnam et al., 2010), or when information about past behavior is available, compared to when it is not (Charness et al., 2011). However, in all of these experiments the frequencies of the outcome (Trust, Reward) falls far below the 100% predicted by the fully cooperative equilibrium. This phenomenon is usually attributed to noise and hence little attempts have been made to check to what extent it could have a theoretical explanation. The idea that experimental subjects coordinate on different equilibria of the repeated game has attracted no attention so far. Our experiment fills this gap. Specifically, we want to test whether the possibility of building several types of reputation can lead subjects to (try to) coordinate on different equilibria. With this aim, we design an experiment where subjects play a repeated TG with an unknown ending. Players keep their role throughout the game, and they are randomly rematched in every period. We force the second player to build a reputation, in the sense that, before choosing, each first player receives some information about the opponent's past behavior.

We have two treatments: in the first treatment (called Last-treatment), only the action taken by the second player in the previous period is disclosed to the first player. In the second treatment (called Full-treatment), the first player is informed about the frequency with which the opponent chose each action in the previous rounds. The two settings allow for different equilibria. A strategy that prescribes to choose Not Reward once in a while can be an equilibrium only when

the information given to the first player is rich enough. Otherwise, if the first player only knows the previous action, the second player would be better off by always choosing Reward.

In this sense, this paper relates to the literature on Extortionate strategies (ExS). Those are strategies that, by conditioning only on the last outcome of the game, are able to induce full cooperation and at the same time to get a higher payoff than the opponent. As such, ExS can never be an equilibrium if the opponent only recalls the last outcome of the game. They might instead have a chance if the opponent has longer memory.

Thus, we have two clear predictions over subjects' behavior: fully-cooperative second players are more frequent in the Last-treatment, and second players who are "nice, but not always" enjoys a higher payoff in the Full-treatment.

Overall, our results confirm these predictions. When looking at the second player's average frequency of reward throughout the game, we find more subjects choosing a fully-cooperative strategy in the Last treatment. Moreover, not only those who were "nice but not always" were able to enjoy a higher payoff in the Full-treatment, but they were also able to gain more than the fully-cooperative subjects.

Those results suggest that the repetition of the game is not enough to ensure that the fully cooperative equilibrium will be selected, even if the game has an unknown ending. Indeed, when the information is sufficiently rich, there are (many) other equilibria which can guarantee to the second player a higher payoff. We show that subjects are indeed able to recognize the different settings implied by the different information disclosed, and are ready to "extort" the first player, if given the opportunity to do so. Nonetheless, the first players in our experiment were less willing to give in to Extortion than the theory predicts.

## 2.2 Trust in the Lab

Several studies used laboratory experiments to examine behavior in the repeated Trust Games. Some of them focused on the effects that disclosing information may have on the frequency of trust and reward observed. Even if providing information is less effective than introducing an incentive system (Bigoni et al., 2014), or allowing for partner matching (Bolton, Ockenfels, et al., 2011), there is an overwhelming evidence that disclosing information is successful in promoting cooperation, compared to the no-information case.

For example, Bracht and Feltovich (2009) show that, in finitely repeated games, giving information about the opponent's last action was successful in promoting Trust and Reward, compared to the case of no-information. Similar results were also found when information was provided about the opponent's last five moves (Minozzi, 2015) or about the full history of past choices (Bohnet et al., 2005). Moreover, providing information increased the rate of cooperation in the repeated image scoring game (Bolton, Katok, et al., 2005), and in a Trust Game where the first player can choose the amount of information he wants to have (Charness et al., 2011).

Few studies tried to compare the effects on cooperative behavior of varying the amount of information disclosed to the first player about the second player's previous behavior. The main finding seems to be that, when increasing the amount of information, trust is usually higher. For example, Keser (2003) considered a finitely repeated Trust Game in which first players could provide feedbacks concerning the trustworthiness of the second players. They found that giving information about the distribution of previous ratings, increases the average frequencies of both Trust and Reward, compared to the case where only the most recent rating is available.

The study which is closer to the one we present below is Duffy et al. (2013). The authors consider a repeated Trust Game with unknown ending, comparing, among other conditions, the case where only the last action taken by the second player is disclosed to the opponent (minimal information), to the one where also the past frequencies and the detailed information about the last 10 periods' choices are available. On the one hand, compared to the no-information case, disclosing minimal information increased the level of Reward (given that the first player chose Trust), but not the average level of Trust. This seems to suggest that disclosing information to the first player triggers a response from the second player, and is thus effective in promoting cooperative behavior. On the other hand, compared to the detailed-information case, disclosing minimal information decreases both levels of trust and reward. Thus, they observed more cooperative behavior when more information was disclosed to the first player.

Compared to their findings, we found higher levels of Trust and Reward in the case of minimal information, but lower levels of both variables in the case of *more* information<sup>1</sup>. However, our treatment with *more* information is different from the one considered by Duffy et al. (2013). In their setting, the first player knows the *detailed* history of the previous rounds. This may crucially change our equilibrium predictions, which are based on the assumption that the first player knows the *scrambled* sequence of the second player's past actions.

---

<sup>1</sup>Notice that we considered a simultaneous TG, while they used a sequential TG.

It is instructive to compare the theoretical explanation Duffy et al. (2013) propose of their experimental results to the one we shall present below. The main aim of Duffy et al. (2013) is to test Kandori (1992)'s model of "contagious equilibrium". This type of equilibrium arises in settings in which a game is repeatedly and anonymously played by a group of players, and *no* information is provided on the other player's past behavior. Using the Prisoner's Dilemma as an example, Kandori (1992) proved that even in this case there are cooperative equilibria, provided that players are sufficiently patient. To sustain such equilibria, players are divided into two types: c-types, who are players who have never defected or experienced a defection in previous rounds, and d-types, who are all the others. All c-types are supposed to play C while all d-types play D. Whenever a c-type meets a d-type, he becomes a d-type himself (notice that d-types never revert to be c-types). To see why these strategies can form an equilibrium, suppose that initially all players are c-types so that cooperation is observed at every round. Switching to D triggers a contagious decline of cooperation, as every individual who meets a defector will switch to defection and increase the speed at which defection spreads. If players are sufficiently patient, they will prefer not to trigger the decline of cooperation. In a recent paper, Xie and Y.-J. Lee (2012) extended this result to the Trust Game.

The contagious equilibrium is obviously very fragile, as a single defection suffices to bring about an irreversible collapse of cooperation. Focusing on the repeated Trust Game, Duffy et al. (2013) study a variant of this model, in which players are allowed to see the past history of the individual they are matched with. In this setting, strategies as in Kandori (1992) are not an equilibrium anymore. The intuition is straightforward: when meeting a c-type, the first player will always play Trust, even if he is a d-type. Duffy et al. (2013) show that in this case there is an equilibrium where second players play a strategy as in Kandori (1992), while first players use a variant that prescribes to choose Trust only when meeting a c-type, regardless of one own's type. Notice that this equilibrium is easier to sustain, the more information on previous behavior is available. Nonetheless, it has a weak element in that first players are assumed to play Trust only if the opponent has constantly played Reward in the past.

We instead allow first players to trust a second player even if he played Not Reward in the past, provided that the frequency of Reward is sufficiently large. The second player may then have an incentive to form a reputation for playing a mixed strategy. The empirically verifiable consequence of this model is that reputations in mixed strategies are not possible when only the second player's last choice is observed. This is the main point of our experimental design.

## 2.3 Should you be nice? Equilibria in the Trust Game

In our experiment we consider the simultaneous version of the Trust Game, as in Table 2.1.

	$R$	$NR$
$T$	$R, R$	$S, T$
$NT$	$P, P$	$P, P$

Table 2.1: Stage game payoffs in the TG

The game has one pure Nash equilibrium, mutual defection, and a compact set of equilibria where the first player chooses Not Trust and the second player chooses Reward with a probability smaller than  $q^S = \frac{P-S}{R-S}$ , that is the probability that would make the first player indifferent between Trust and Not Trust (this is also called the second player's Stackelberg strategy, as is it the probability he would like to commit to play, if having the chance to do so).

Because of the Folk-Theorem, when players are sufficiently patient there is a large set of efficient, more cooperative equilibria in which they may get different payoffs. Figure 2.1 represents the set of feasible payoff profiles, in which the payoff of the first player is on the y-axis and the payoff of the second player is on the x-axis. When players are sufficiently patient, all points that Pareto dominate the  $NT$  payoff profile can be sustained by a profile of strategies that forms a SPNE. One of those is the Stackelberg payoff profile  $b_2$ . This is the profile in which the second player's payoff is maximized, under the assumption that the first player gets at least her outside option payoff  $P$ . We call the corresponding payoff for the second player his Stackelberg equilibrium payoff  $\pi^S$ , as it is the payoff he would get if he could commit himself to choose a (mixed) strategy before the first player makes his choice. Notice that  $\pi^S = Rq^S + T(1 - q^S)$ .

In Figure 2.1 we plotted some examples of unconditional strategies for the first and for the second player. Unconditional strategies for the first player are the dotted parallel lines, while unconditional strategies for the second player are dashed lines originating from the point  $NT$ .

Consider an equilibrium for which the resulting payoff profile is in the blue area. This means that, on the equilibrium path, the frequency with which the second player chooses Reward is lower than the frequency with which the first player chooses Trust. Moreover, the second player is able to enjoy a payoff that is larger than the mutual cooperation payoff.

When in equilibrium the realized payoff combination lies in this area, we say that the second player *extorts* the first player, and we call these equilibria *extortionate*.

Notice that any efficient equilibrium different from mutual cooperation lies in the blue area, as it requires that, on the equilibrium path, the second player mixes between Reward and Non Reward, while the first player always trusts. Thus, according to our definition, any efficient equilibrium is extortionate, except for the mutual cooperation payoff. We can then distinguish between mild-extortion and hard-extortion, according to whether the resulting payoff profile is closer to the mutual cooperation or the Stackelberg payoff profile. Indeed, any profile in between can be sustained in equilibrium, provided that the frequency with which the second player plays Reward is sufficiently large as to make playing Trust a best reply.




Figure 2.1: The set of feasible payoff profiles in the repeated Trust Game.

In a seminal paper, Press and Dyson (2012) suggested that in a repeated Prisoner’s Dilemma a player could extort a larger payoff than the opponent by using simple memory-one strategies they call ”extortionate strategies”. These strategies are memory-one in the sense that the probability with which they cooperate at each round (which is always smaller than one) only depends on the outcome of the previous round. Extortionate strategies are obtained by setting the probability of cooperation after each outcome in such a way that the opponent cannot do better than cooperate at every round. This implies that any strategy playing against an extortionate strategy will be forced to accept a smaller payoff than the opponent.

An implication of Press and Dyson (2012) result seems to be that what we call extortionate equilibria can be obtained by players using simple memory-one strategies. This intuition is incorrect: as long as players are constrained to use memory-one strategies, no extortionate equilibrium exists in this sense (see Chapter One). This conclusion is quite surprising in the light of the current literature. Not only memory-one strategies cannot sustain extortionate equilibria, but, when one player uses a memory-one strategy, the maximum payoff that his opponent can

get is the mutual cooperation payoff. This implies that there are no extortionate equilibria against a player who uses a memory-one strategy, no matter how complex his opponent's strategy is.

When considering the Trust Game, the set of payoffs that can be sustained in equilibrium is even smaller. If the first player is constrained to use a memory-one strategy, not only the second player payoff is bounded from above by  $R$ , (implying that there are no extortionate equilibria, like in the Prisoner's Dilemma), but in any equilibrium different from mutual cooperation, the first player's payoff is fixed at  $P$  (see Chapter One).

Although some of the proofs are fairly involved, the intuition behind these results is straightforward. In any extortionate equilibrium, the second player is required to use a mixed strategy, and to mix *on* the equilibrium path. When the first player only conditions on the previous outcome of the game, no deviation from a mixed strategy can be observed and hence punished. Thus, in any equilibrium in which players are mixing on the equilibrium path, the indifference condition implies that the the payoff of the second player is bounded from above by  $R$ , and the payoff of the first player is fixed at  $P$ , so that no extortionate outcome can be an equilibrium. Notice that this result doesn't prove that no extortion can ever take place in equilibrium. The results only imply that for an extortionate equilibrium to exist, one player (the one who suffers the extortion) must be able to condition her choice on at least the *frequency* with which the other player plays cooperatively in the previous rounds. Indeed, if the player is able to condition not only on the opponent's realized action in the previous period, but also on the mixture that he used, extortionate outcomes could be sustained in equilibrium. In this context, Barlo et al. (2009) prove a version of the Folk-theorem. They show that if players can condition on the mixed action chosen by the opponent in the previous round, then all payoff combinations that Pareto dominate the non-cooperative equilibrium can be sustained as subgame perfect equilibrium. For the Trust Game, this result directly implies that extortionate equilibria exist in memory-one strategies with observable mixtures (see Chapter One).

There is an interesting observable implication of the results we just summarized. In experimental settings in which the first player can only condition on the choice made by the second player in the previous round, no extortion should be observed. When richer information about the second player's past choices is available, then at least some extortion should take place. Notice that this is somehow counterintuitive, as one might be lead to think that, in a context of minimal information, we should observe more extortion, as it is easier for the second player

to *clean* his history, by playing Reward only once. However, if first players punish accordingly, the second player would be better off by playing Reward at every round.

Specifically, we expect subjects in our experiment to be able to distinguish between these two situations, so that they might choose more often a mixed (Extortionate) strategy when it is convenient to do so. To the best of our knowledge, we are the first to test Extortion in the laboratory. Previous attempts have been made, but they involved subjects playing against a computer who was programmed to play an Extortionate strategy. Interestingly, when participants are aware they are facing a computer, they learn to play a best reply to the Extortionate strategy, and they fully comply (Wang et al., 2016). However, when they are not aware that they are facing a computer, they don't give in to Extortion, resulting in a poor performance for the Extortion strategy, and a lower payoff for the first players (Hilbe, Röhl, et al., 2014). In this respect, our results are coherent: also the first players in our experiment were less willing to give in to Extortion than the theory predicts.

## 2.4 Extortion in the Lab

### 2.4.1 Experimental Design

There are obvious difficulties in devising an experimental setting in which our model can be tested. The main difficulty is common to all experiments on repeated games, and is due to the fact that the repeated game strategies are not observable (see Dal Bo and Frechette, 2018). Even when observing mutual cooperation, it is difficult to ascertain whether a player uses TfT or WSLS or any other *nice* strategy, i.e. a strategy that is never the first to defect. A further difficulty, which is more specific to the setting we have in mind, stems from the fact that it seems impossible to determine the type of information on which subjects will condition their decisions. A way must be found to force subjects to play memory-one strategies in one setting, and more sophisticated strategies in others.

Our experimental design is meant to get round these difficulties. We use a variation of the experiment proposed by Duffy et al. (2013), in which the amount of information given to first players about second players' past behavior varies.

Upon arrival, subjects in the experiment are randomly assigned to a role, either of "Partecipante Uno" (first player) or "Partecipante Due" (second player). Subjects keep their role throughout



the experiment. At the beginning of each period subjects are rematched with a new opponent and play a simultaneous Trust Game, with payoffs as in Table 2.1.

	<i>R</i>	<i>NR</i>
<i>T</i>	4, 4	0, 6
<i>NT</i>	1, 1	1, 1

Table 2.1: Payoffs in the experiment

Notice that in the simultaneous version of the TG, second players are asked to take their decision even if the opponent chooses NT. This choice was motivated by several considerations. First, with simultaneous moves we are able to have observations also in periods when the first player chooses not to trust. Second, in the simultaneous version of the game the second player chooses in a “cold state”, i.e. without knowing whether his opponent is trustful or not<sup>2</sup>. Finally, and most importantly, in a sequential game a likely outcome would be that many first players would have chosen not to trust at the beginning of the session and this would have prevented many second players from forming a good reputation. Most observations would then have been on the inefficient equilibrium (Not Trust, Not Reward).

In every period after the first one, we disclose to the first player information about the current opponent’s past behavior. We have two treatments, denoted as Last and Full. In the Last-treatment the first player is informed about the last action (either Reward or Not Reward) taken in the previous period by the second player with whom he is currently matched. In the Full-treatment, the frequency with which Reward was played by the second player in all previous rounds is computed and shown to the current first player.

The motivation behind this experimental design is that in the *Last-treatment* we (try to) force the first player to condition his choice on just the last move of the second player, as it would be the case if he were constrained to use a memory-one strategy. In the *Full-treatment*, we (try to) force the first player to use a strategy that conditions on the overall frequency of past cooperative behavior of the second player, which can be considered as a (rough) approximation of the case of observable mixtures.

The experimental design is devised to minimize confusion. In the *Full* treatment we only reveal the frequency with which Reward was chosen by the second player in the past, but not his last choice, to avoid that the first player conditions on the frequency *and* on the last outcome.

---

<sup>2</sup>Indeed, this is the explanation given by Guth et al. (1997) to the finding that second players tend to reward slightly more with sequential, rather than simultaneous, moves.

## 2.4.2 Predictions

From a theoretical point of view, the Last-treatment is meant to be a proxy for a model in which a game is played repeatedly, but players can only recall the last action taken by the opponent. The Full-treatment approximates a setting in which the frequency of Reward is recalled, as it would be the case in a repeated game with observable mixtures. The results we summarized in the previous section imply that these settings have very different implications for the payoff profiles that can be sustained in equilibrium.

Consider again Figure 2.1, which represents the set of feasible payoff profiles for the repeated Trust Game. In our experiment, an observation is a point in that triangle, whose coordinates are the average payoff earned by one of the second players (on the x-axis) and the average payoff earned by the all the first players that have been matched with her.

To simplify the present analysis, we restrict attention to *stationary* equilibria in which the second player uses an unconditional strategy, that is a strategy that prescribes the same action at every period. Imagine for example a second player who, in either treatment, plays at every round Reward and Not Reward with the same probability. When matched against this player, all payoff combinations are forced to lie on the dashed line that connects the NT point to the point labeled 0.5 in Figure 2.1. To see this, notice that the resulting profile is going to be a linear combination of the profile NT, (i.e. the profile when all first players decide to play Not Trust), and the profile labeled 0.5 (i.e. the profile when they all decide to play Trust).

We shall denote as  $f$ -Reward the unconditional strategy that plays Reward with probability  $f$  at every round. Notice that, by setting different values of  $f$ , the second player can force the payoff profile to lie on a different line passing through NT, whose slope is determined by  $f$ .

Recall that  $q^S$  is the probability that makes the first player indifferent between choosing Trust and Not Trust, which determines the second player's Stackelberg payoff. Then, in the terminology of the previous section, every  $f$ -Reward strategy with  $1 > f > q^S$ , is an extortionate strategy. Indeed, any such strategy is able to induce full Trust by the first players, and it is thus able to reach extortionate payoff profiles. Clearly, extortion is milder the higher the value for  $f$ . In the payoff space, this implies that the larger the probability  $f$  of playing Reward, the steeper the line on which the payoff profiles will lie. In Figure 2.1, we plot four such lines, corresponding to  $f$  equal to 0.9, 0.75, 0.5, and 0.25. Notice that the latter one is the Stackelberg strategy,  $q^S$ , and it connects the profile NT with the Stackelberg profile  $b_2$ .

While it is true for both treatments that by choosing a value of  $f$  the second player can force the payoff to lie on a straight line passing through  $NT$ , from our theoretical analysis we know that only in the Full-treatment  $f < 1$  can be sustained as an efficient Nash equilibrium.

The theoretical results summarized in the previous section imply that, when players can only condition on the previous move, it doesn't pay to mix. Moreover, as Press and Dyson (2012) noticed, if a player has limited memory, his opponent cannot gain from using a longer memory strategy. This implies that, in our Last-treatment, also the second player will use a memory one strategy, and we can apply the results from Chapter One, shown in Figure 2.1a. Thus, in the Last-treatment, the only (strictly) efficient equilibrium has  $f = 1$ , with payoff profile (4,4).

When mixed strategies are observable, it is easy to prove that the only equilibria are those on the Pareto frontier. To see this, notice that when the first player can observe  $f$ , his best reply is to choose Trust whenever  $f \geq q^S$ , and Not Trust otherwise. This implies that every  $f \geq q^S$  can form a NE able to induce full trust, so that the corresponding payoff profile will lie on the Pareto frontier. Among those, there is also the profile  $b_2$ , i.e. the profile in which the second player gets his Stackelberg payoff, corresponding to the strategy  $f = q^S$ . Notice that this is the only equilibrium that survives the elimination of (weakly) dominated strategies. Thus, in the Full-treatment, there is a compact set of NE in which  $f$  ranges from 1 to the Stackelberg probability  $q^S = \frac{1}{4}$ . The payoffs profiles associated to these NE are on the part of the Pareto frontier in which the first player's payoff is at least equal to 1, as shown in Figure 2.1b.

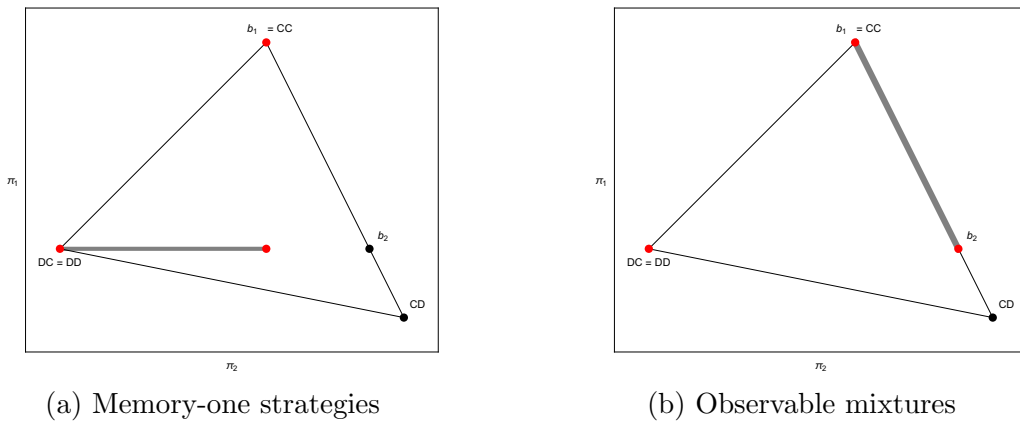


Figure 2.1: NE payoff profiles in the repeated TG (red dots and dark gray line).

We can summarize the discussion above by saying that both in the Last and in the Full-treatments there are NE in which the first player never trusts and efficient equilibria in which he always trusts. As a consequence, we do not expect the overall level of trust to depend

on the treatment. Nonetheless, we expect to find more second players willing to choose a fully cooperative strategy in the Last-treatment, and more second players willing to choose an extortionate strategy in the Full-treatment. Finally, we expect extortionate strategies to be more profitable in the setting in which they are able to form an efficient equilibrium, i.e. in the Full-treatment. This provides us with two hypotheses to test:

**Prediction 1** The strategy that prescribes to always play Reward is more frequent in the Last-treatment, while Extortionate strategies are more frequent in the Full-treatment.

**Prediction 2** The second player, by using an Extortionate strategy, is able to gain more than the mutual cooperation payoff only in the Full-treatment.

When looking at Figure 2.1, those predictions imply that payoff profiles observed in the Last-treatment are all in the green area, while in the Full-treatment we expect at least some second players to force an extortionate equilibrium in the blue area.

## 2.5 The experiment

### 2.5.1 Setting

The experiment was run between November 2015 and May 2016, at CEEL (University of Trento, Italy). A total of 112 subjects participated to the experiment, all students from the University of Trento. The experiment was conducted using the software Z-tree (Fischbacher, 2007). Participants were randomly divided into two groups, first and second players, and they kept their role throughout the experiment. The instructions were given in a neutral way, with actions labeled as “Left”, “Right”, “Top” and “Bottom”.

The experiment consisted of two parts. In the first part, we ran the standard repeated Trust Game. To avoid too short sessions, subjects were randomly matched for at least 20 periods. After that, a coin was tossed at the end of every period to determine the end of the game. In the second part, we asked subjects to choose a strategy to play the repeated game. Second players had to specify a probability (from 0 to 1) of choosing Reward at every round, i.e. they had to pick a (possibly mixed) unconditional strategy. First players had to specify an action (i.e. either Trust or Not Trust), conditional on the possible information they would receive

about the previous behavior of the current opponent. For example, in the Last-treatment, first players had to choose whether or not to trust a second player who chose either Reward or Not Reward in the previous match. The strategies chosen by all subjects were then used as inputs for a computer simulation which we ran at the end of the experiment. Each subject's final payoff was the sum of the payoff he earned in the first part plus the payoff the strategy he choose earned in the computer simulation. The results of the second part of the experiment will be presented in a forthcoming paper.

In the first part of the experiment, the way information was presented was meant to be as intuitive as possible. In the *Full-treatment*, on the computer screen the first players could see a histogram indicating the frequency with which each action was chosen in the past by their current opponent (see Appendix 2.B). On the right-hand side of the screen they could see a bar-graph indicating their own average payoff up to the current period. Second players could see a bar-graph representing their own payoff up to that period and on the left-hand side a histogram indicating the frequency with which they had chosen each action in the past. In the *Last-treatment*, the setting was the same, with the only exception that the first players could only see the action chosen by the current second player in the previous round, instead of the histogram with the frequencies of the past choices.

In both treatments, the screen was not showing the number of periods elapsed, and it is very implausible that subjects counted periods by themselves. There is no reason to believe that the fact that the random termination of the game only begun at the twentieth round played any role in their decisions. Also, our data do not reveal any endgame effect.

We ran 3 sessions per treatment, one session with 16 subjects, two with 18 subjects and all the others with 20 subjects (so that a total of 29 second players where involved in the Full-treatment, and 27 in the Last-treatment). Thus, each player was interacting with at least 8 opponents. Our group size is quite large compared to similar experiments (see for example Duffy et al., 2013). We decided to run sessions with a large number of participants to minimize the chance of two players interacting again with each other<sup>3</sup>.

Each subject was involved in only one treatment, and plays only one repeated game. The shorter game lasted 20 periods, and the longer 26 periods. Monetary payoffs were equal to the sum of payoffs received during the game, at the exchange rate of 1 token= 0.30 Euro. The average payoff was 10 euro, including a show-up fee of 3 euro.

---

<sup>3</sup>Subjects knew they could meet again the same opponent, but only after at least 8 rounds.

## 2.5.2 Results

The main result of our experiment is summarized by Figure 2.1. Each point corresponds to the combination of the average payoff of a second player, and the average payoff of all the first players that have been matched with him (we call this the *induced* payoff). Blue points correspond to the Full-treatment, while red points correspond to the Last-treatment.

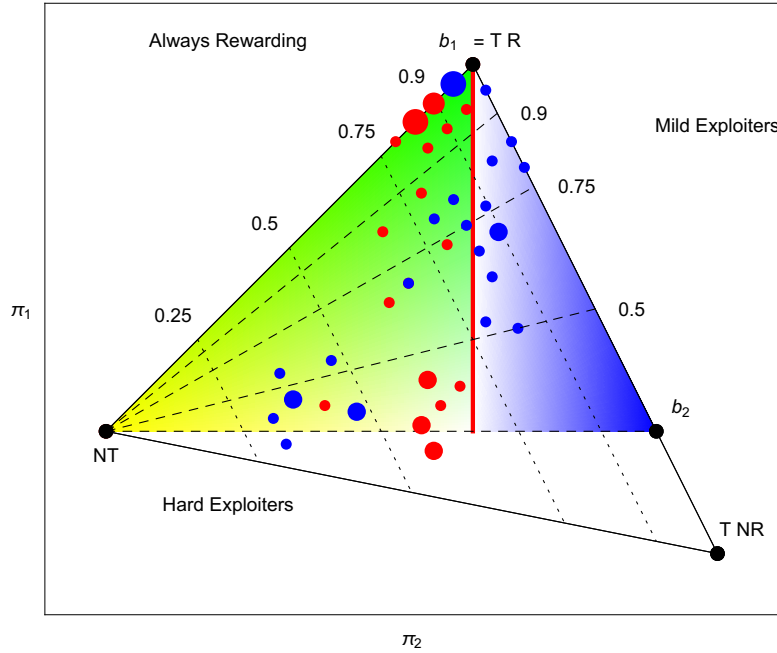


Figure 2.1: Results - Average payoffs in the 20<sup>th</sup> period.

According to Prediction 1 second players should choose more often a fully-cooperative strategy in the Last-treatment, and more often an extortionate strategy in the Full-treatment. According to Prediction 2, only in the Full-treatment second players are able to gain more than the mutual cooperation payoff. These predictions imply that the majority of red points should be on the left of the red vertical line, while the majority of blue points should be on the right.

The experimental results are broadly in line with the theoretical predictions. Particularly striking is the fact that the second players were able to reach a payoff larger than 4 only in the Full-treatment. This proves that at least some of the subjects involved were able to understand the logic of extortionate equilibria and play accordingly. At the same time, the ability to extort a larger payoff looks rather limited, as no second player succeeded in getting close to her Stackelberg equilibrium payoff. In fact, our data reveal that playing Reward with a frequency between 0.25 and 0.5 was not enough to induce an high level of trust, although the theory predicts the existence of such equilibria.

Notice that there is a clear analogy with the extensive evidence on the Ultimatum game. Some subjects may have interpreted frequencies of Rewards below 0.5 as "unfair", and decided to play Not Trust as a form of punishment. The extent to which some form of social preferences may explain this phenomenon is an open question that deserves further investigation.

In the Last-treatment, dots are mostly found around two areas. There is an area of cooperative equilibria, in which the most common outcome is (T,R), and an area of almost-extortionate equilibria, where second players try to extort the opponent (in the middle of the triangle). Notice that those points are around a (non-efficient) mixed equilibrium in memory one strategies (i.e. the red point in the middle of the payoff space in Figure 2.1a). Even if those players were able to gain more than players choosing Reward with similar frequencies in the Full-treatment, they were not able to get more than the mutual cooperation payoff, as our theory predicts.

We will see in the following sections that second players can be divided into 3 groups: those who always cooperate (always rewarding), those who cooperate, but not always (mild exploiters) and those who defect, but not always (hard exploiters, see Figure 2.1).

Our data reveal that Always rewarding types are more common in the Last-treatment, while Mild exploiter types are more common in the Full-treatment. Moreover, trust was the highest for mild exploiters in the Full-treatment, who were the only ones able to get more than the mutual cooperation payoff. Both our predictions are thus confirmed.

Those results imply that the effect of information on players' behavior is twofold: on the one hand, increasing information allow second players to successfully extort first players, as they are able to gain more than the mutual cooperation payoff. On the other hand, providing information helps first players to discriminate between mild and hard abusers, choosing to trust only the firsts. Interestingly, we find that first players tend to trust a second player who chose Not Reward in the previous period often enough to allow him to get a payoff higher than mutual defection, but not enough to allow him to get more than the mutual cooperation.

Overall, we can say that changing the amount of information disclosed to the first player was effective in triggering second players to adopt different strategies, as they tend to be more mild exploiters in the Full-treatment, that is when extortionate strategies are more profitable.

In the following, we will give a deeper description of our data. After an overview, we will describe second players' *observable* strategies, and we will analyze more in detail the role of information.

## An overview of the data

We shall now give a brief description of the main features of the data we collected. Our analysis considers the first 20 periods of each session. We may refer to the 20<sup>th</sup> period as the last, or the end, but keep in mind that it was the last period only in 2 sessions, and that the average length of play was of 22 periods.

Table 2.1 shows the average frequencies of actions Trust (T) and Reward (R), together with average frequencies of the different outcomes: Trust and Reward (T,R), Trust and Not Reward (T,NR), Not Trust and Reward (NT,R), and Not Trust and Not Reward (NT,NR). Values are just the sum of the different choices (or outcomes) divided by the total number of observations.

	T	R	T,R	T,NR	NT,R	NT,NR
Full	0.74	0.70	0.55	0.19	0.15	0.11
Last	0.74	0.68	0.55	0.18	0.14	0.13

Table 2.1: Average frequencies - Frequencies are computed considering periods 2-20. Sample size (subjects for each role) is  $N = 27$  for the Last-treatment and  $N = 29$  for the Full-treatment

The averages are almost exactly the same across treatments, a finding in sharp contrast with Duffy et al. (2013), who found a significant difference in the average frequencies of both Trust and Reward. Specifically, they found a lower level of Trust (46%) and a similar level of Reward (71%) in their min-treatment (corresponding to our Last-treatment), and a higher frequency of both Trust (81%) and Reward (94%) in their info-treatment (but recall that, contrary to our Full-treatment, they also gave to first players the detailed history of the previous 10 rounds of the game; moreover, Duffy et al. (2013) consider smaller groups and shorter interactions). The results from the Last-treatment might suggest that, contrary to previous findings, when moving from sequential to simultaneous games the average trust may increase in settings in which only minimal information is available.

**Finding 1.** *There is no treatment effect as long as the average frequencies of Trust and Reward are concerned.*

We now move to choices across periods. The average frequencies of Reward do not change over time, ranging between 0.6 and 0.8 in both treatments, while the average frequency of Trust do change over time, increasing in the Last-treatment (from 0.68 to 0.74) and decreasing in the Full-treatment (from 0.8 to 0.71).



Table 2.2 shows the average frequencies of outcomes (T,R) and (T,NR) in the early (2-7) and late (15-20) periods of the game. Even if at the beginning of the game the outcome (T,R) is more frequent in the Full-treatment, the opposite is true at the end of the game. At the same time, the frequency of the outcome (T,NR) is higher in the Full-treatment both at the beginning and the end of the game, although not significantly.

Outcome Periods	T,R		T,NR	
	2-7	15-20	2-7	15-20
Full	0.57	0.5	0.22	0.21
Last	0.48	0.55	0.20	0.18

Table 2.2: Average frequencies for early and late periods. Only the differences for the outcome (T,R) are significant (Fisher test, significant difference between the treatments at the beginning of the game (p-value  $\leq 0.05$ ), and across each treatment comparing the begin with the end (p-values  $\leq 0.01$ )).

**Finding 2.** *The average frequency of the outcome (T,R) is increasing over time in the Last-treatment, and decreasing in the Full-treatment. The frequency of the outcome (T,NR) is always higher in the Full-treatment, even if not significantly.*

Next, we consider the frequency with which each subject played Reward in the firsts 20 periods. Recall that we expect this frequency to be equal to 1 in the Last-treatment, and smaller than one in the Full-treatment. Figure 2.A.1 in the Appendix shows that this is not the case. Moreover, even if the median choice is almost the same between treatments, the distribution of choices is more dispersed in the Last-treatment.

A different picture emerges when one looks at the way specific players behaved. We shall now compute the same variables, restricting the attention to those subjects who played Reward at the first period. These subjects are particularly interesting, because they are those who are most likely to be interested in building a good reputation. Fig. 2.A.2 in the Appendix shows the average frequencies of Reward considering only those subjects. Not only the median in the Last-treatment is higher than the median in the Full-treatment, but also the distribution in the Last-treatment is less dispersed than when all subjects were considered. Thus, there is scope to say that those subjects who signal in the first period their willingness to build a good reputation are more likely to exploit the first player (i.e. they have, on average, a lower frequency of Reward) in the Full-treatment, that is when extortionate equilibria are possible.

Table 2.3 shows the average frequencies of actions and outcomes, as in Table 2.1, considering again only the subjects who choose Reward in the first period. These data are interesting as

they reveal whether a second player who alternates between Reward and Not Reward is able to induce first players to play Trust. This seems in fact to be the case as the outcome (T,NR) is significantly more frequent in the Full-treatment than in the Last.

	T	R**	T,R**	T,NR***	NT,R	NT,NR
Full	0.78	0.72	0.58	0.20	0.13	0.09
Last	0.78	0.79	0.66	0.12	0.13	0.09

Table 2.3: Average frequencies considering only subjects starting with R - Sample size (subjects for each role) is  $N = 19$  for the Last-treatment and  $N = 24$  for the Full-treatment. Significance: \* 0.1; \*\* 0.05; \*\*\* 0.01.

This time we find a (significant) difference in the frequency of Reward, as well as of the outcomes (T,R) and (T,NR), while the average frequencies of Trust and of the other outcomes are still equal across treatments. Thus, those second players who choose to Reward in the first period choose Reward less often in the Full than in the Last-treatment. Apparently, they were also able to induce the first players to play Trust, which is compatible with the outcome (T,NR) being significantly more likely in the Full than in the Last-treatment. This is in accordance with our Prediction 2 and is summarized by the following:

**Finding 3.** *When considering only those players who choose Reward in the first period, the average Reward is (significantly) higher in the Last-treatment, while the average Trust is the same between treatments, in accordance with Prediction 1.*

Finally, Table 2.4 shows 3 different ratios, which will be useful in this first step of the analysis.

The first is the ratio of the number of outcomes (T,R) over the total outcomes in which the second players played Reward:  $r_1 = \frac{\#(T,R)}{\#R}$ . This ratio should be equal to 1 in both treatments, if the second player succeeds in building a good reputation.

The second is the ratio of the number of outcomes (T,NR) over the total outcomes in which Not Reward was played:  $r_2 = \frac{\#(T,NR)}{\#NR}$ . This ratio represents the relative frequency with which a Not Reward is Trusted. If our predictions are correct, it should be equal to 1 in the Full-treatment and to 0 in the Last-treatment.

Ratios  $r_1$  and  $r_2$  can be interpreted as the (ex-post) probabilities of being trusted when playing Reward and when playing Not Reward.

The last is the ratio of the number of outcomes (NT,R) over the total outcomes in which Not Trust is played:  $r_3 = \frac{\#(NT,R)}{\#NT}$ . This ratio tells how often the first player could have been better

off by trusting the second player. Our theoretical results imply that the first player should choose Not Trust only in the non-cooperative equilibria, meaning that  $r_3$  should be equal to zero in both treatments.

	All			Reward First		
	$r_1$	$r_2$	$r_3$	$r_1$	$r_2^{**}$	$r_3$
Full	0.72(29)	0.69(25)	0.56(27)	0.76(24)	0.77(20)	0.62(22)
Last	0.77(27)	0.61(20)	0.53(24)	0.82(19)	0.57(12)	0.62(18)

Table 2.4: Ratios  $r_1 - r_3$ . The first three columns consider all second players, while the last three columns consider only those starting with R. Sample size is in parentheses. Significance: \*\* 0.05

All the ratios are different than predicted. The first and fourth columns of Table 2.4 tell us that second players were not able to induce trust, and that first players would have been better off by trusting, on average, more. Indeed, the third and sixth columns of Table 2.4 confirm that first players were too suspicious, and that more than 50% of times that they choose Not Trust, they would have been better off by trusting the second player. Nonetheless, the second and the fifth column show that in the Full-treatment it was easier to maintain a good reputation while exploiting the first player. In fact, in the Full-treatment it was more likely for the second player to be trusted when playing NR. The difference between treatments increases and becomes significant when considering only those subjects starting with Reward.

**Finding 4.** *Values of ratios  $r_1 - r_3$  are different than predicted. Nonetheless, the probability of being trusted when choosing Not Reward ( $r_2$ ) is higher in the Full-treatment, which is in accordance with Prediction 2.*

## Strategies

So far we classified second players according only to the action they played in the first round. We shall now look at another dimension: the frequency  $f$  with which each subject played Reward across the 20 periods, which we shall refer to as the subject's  $f$ -Reward strategy, or  $f$ -strategy. Recall that our theoretical model predicts that, if  $f \neq 0$ , then  $f$  is concentrated at 100% in the Last-treatment, and is between 25% and 100% in the Full-treatment.

Figure 2.2 shows the distribution of strategies in the two treatments. On the x-axis we plot the  $f$ -strategies, and on the y-axis we report the number of subjects who played that strategy.

As we can see, the mode in the Last-treatment is 100%, while in the Full-treatment the distribution is bimodal, with the two modes at 100% and 75%. Moreover, even if there is a great heterogeneity among subjects, nobody chose an  $f$ -strategy with  $f < 0.25$  in the Full-treatment. Thus, it seems that increasing information lead players to choose less often extremely low or extremely high  $f$ -strategies, in favor of intermediate ones.

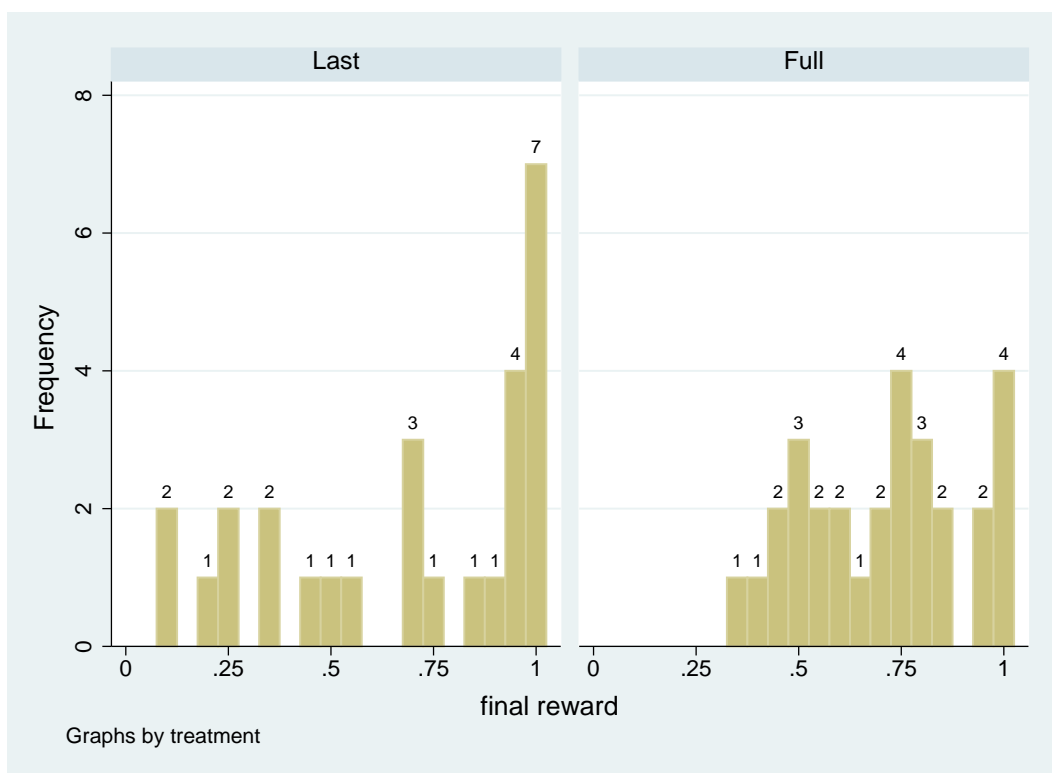


Figure 2.2: Frequency of players' reward at the 20<sup>th</sup> period. Skewness: full = -.033 last = -.555; kurtosis full = 1.9\* last = 1.8\*\*. Epps-Singleton Two Sample test: not equal, p-value < 0.01

To analyze the profitability of the different strategies in different treatments, it is useful to compute the average trust induced by each strategy, i.e. how many times a second player playing a particular  $f$ -strategy was trusted during the game, and averaging over subjects with same strategies. In Figure 2.3 (left) the vertical axes represents the frequency with which second players using different  $f$ -strategies (reported on the x-axis) were trusted by the first players. Coherently with our prediction, we observe that second players who use a  $f$ -strategy with a sufficiently large  $f$  are trusted more in the Full than in the Last-treatment. However, the rate at which trust sharply declines is at around 0.6, while the theoretical model puts it at 0.25.



Figure 2.3: Average trust and average payoff for each  $f$ -strategy

In Figure 2.3 (right) we plot the payoff obtained by second subjects in the two treatments, depending on the  $f$ -strategy they choose. Coherently with our model, the highest payoff accrues to those players in the Full-treatment which used an  $f$ -strategy with a sufficiently large  $f$ . Paradoxically enough, in the Last-treatment the payoff earned by the second players does not depend upon the  $f$ -strategy they choose, as if first players were using an Equalizer strategy, i.e. a strategy that is able to fix the opponent's payoff at a certain level, independently from the strategy chosen. We can summarize those observations in the following:

**Finding 5.** *The frequency of subjects choosing an  $f$ -strategy with extremely high ( $f > 0.8$ ) or extremely low ( $f < 0.3$ ) values of  $f$  is higher in the Last-treatment, in line with our Prediction 1. Trust is higher in the Full-treatment for  $f$ -strategies with  $f \geq 0.6$ , and second players' payoffs are higher than 4 only in the Full-treatment, in line with our Prediction 2.*

### Choices over time

In this section we try to classify all subjects on the basis of the pattern of their choices across periods. Figures 2.A.3- 2.A.8 in Appendix 2.A show each subject's pattern of choices over the first 20 periods. We classify as *Always R* subjects who choose Reward 20 times over 20 periods (see Figure 2.A.3); *Almost R* are subjects who play a mixed strategy in the first part of the

game, but then choose *Always R* in the second part of the game (10 times over periods 11-20, see Figure 2.A.4). *Mild exploiters* choose always Reward only in the first part of the game, and they try to exploit in the second part (see Figure 2.A.5). *Exploiters* and *Hard Exploiters* mix during all the game with a final frequency of at least 25%. The difference is that a player classified as Exploiter never goes below 25% (see Figure 2.A.6, "mixing", for Exploiters, and Figure 2.A.7, "mixing too much", for Hard exploiters). Finally, *Almost NR* play Reward with a final frequency smaller than 25% (see Figure 2.A.8, notice that nobody chose this strategy in the Full-treatment). Table 2.5 shows average frequencies, as well as average payoff and trust induced by each strategy in both treatments.

	Frequency		Trust induced		Payoff	
	Full	Last	Full	Last	Full	Last
Always R	0.14 (4)	0.26 (7)	0.94	0.86**	3.81	3.59*
Almost R	0.07 (2)	0.15 (4)	0.97	0.84**	4.07	3.61***
Mild exploiters	0.10 (3)	0.07 (2)	0.96	0.90	4.17	3.80**
Exploiters	0.48 (14)	0.22 (6)	0.72	0.64**	3.67	3.46*
Hard Exploiters	0.21 (6)	0.18 (5)	0.43	0.64***	2.73	3.68***
Almost NR	0.00 (0)	0.11 (3)	–	0.57	–	3.50

Table 2.5: Frequency of strategies, trust and payoffs - Averages of trust and payoffs are computed considering periods 1-20. Significance of differences between treatments: \* 0.1; \*\* 0.05; \*\*\* 0.01.

**Finding 6.** *Compared to the Full-treatment, in the Last-treatment more subjects can be classified as Always R or Almost R, and fewer as Mild Exploiters or Exploiters, in line with our Prediction 1. Trust and payoffs are higher in the Full-treatment for Mild Exploiters and Exploiters, in line with our Prediction 2. Contrary to our Prediction 2, Exploiters are not able to gain more than the mutual cooperation payoff.*

This Finding suggests that there is scope for extortion, although it is smaller than the theory predicts. Exploiters receive too little trust for their strategy to be profitable. Indeed, payoffs in both treatments are not consistent with first players choosing best responses to information: first players would be better off by trusting more after observing a sufficiently large frequency of Reward (or a Reward in the last treatment), and by trusting less after observing a low frequency of Reward (or a Not Reward in the Last-treatment). Interestingly enough, the result of the Last-treatment seems to show that subjects mix between Reward and Not Reward even when such a mix is not profitable, as almost 30% of the subjects can be classified as Exploiters or Mild exploiters even in settings in which this strategy is not part of any NE.

## The role of information

In this section we try to find an answer to two questions: do first players correctly respond to information? And: is information reliable? To this aim, we analyze behavior over time as a function of the information available.

Consider Figure 2.4 right. On the x-axis we put the average frequency with which Reward has been played (and was shown to the first player), and on the y-axis the frequency with which Trust and Reward was played when this piece of information was available. For example, the graph shows that in periods where the first player observed a frequency of Reward of 80 per cent, average Trust was above .8 and average Reward was slightly above .6. If all the second players were using an unconditional mixed strategy, the information would be reliable and all the blue points would be aligned on the diagonal. If all first players were playing a best response to the information they received, average Trust would be equal to 0 for a probability of Reward smaller than 0.25, and to 1 otherwise. The graph shows that the second players were using roughly an unconditional mixed strategy. The first players, however, trusted too little when receiving positive information about the second player's past choices. In particular, the probability of playing Trust sharply declines for a frequency of Reward below 60 per cent.




Figure 2.4: Average Trust and Reward for each level of information.

In Figure 2.4 (left) we plot the same graph, with data coming from the Last-treatment. Just like in the previous graph, the x-axis reports the frequency with which Reward was played by second players, but keep in mind that in the Last-treatment this information was not shown to first players. Our model predicts that in this case there should only be observations for a frequencies of Reward equal to zero or one. The reason is that there is no equilibrium in which the second player mixes. The graph clearly shows that this is not the case. In fact, the second player seems to play different unconditional mixed strategies. At the same time, the theoretical model predicts that (if first players play their equilibrium strategies) all red points should be aligned on the main diagonal. To see this, consider that in equilibrium first players should play Trust when observing a Reward, and Not Trust when observing a Not Reward. It follows that, as long as first players stick to their equilibrium strategy, the frequency of Trust should match the frequency of Reward. This turns out to be true only for high frequencies of Reward, as in this case the red points lie almost perfectly on the main diagonal. However, when considering lower frequencies of Reward, average Trust never goes below 40%. Indeed, first players were too willing to Trust: they trusted with strictly positive probability even when they interacted (without knowing) with individuals who had never played Reward in the past.

In other words, deviations from the equilibrium predictions go in the direction of too much trustfulness rather than too little, tempting the second players to choose exploitative strategies also in the Last-treatment.

**Finding 7.** *Second players where using approximately unconditional mixed strategies in both treatments, although they are part of an equilibrium only in the Full-treatment. Compared to the predictions of the theoretical model, the first players were Trusting too little in the Full-treatment and too much in the Last-Treatment.*

## 2.6 Conclusions

An interesting feature of our data is that there are subjects who alternate between Reward and Not Reward in both treatments. This may be interpreted as an attempt made by second players to exploit first players, which fails when only the last action is observed, so that no reputations in mixed strategies are possible. If there is confusion on the part of the subjects, it seems that it induces them to try too hard (rather than too little), to obtain more than what mutual cooperation could guarantee them. Not surprisingly, the second player was able



to reach a payoff larger than the mutual cooperation payoff only in the Full-treatment, and only provided that the frequency of Not Reward exceeded 50 per cent.

Several theoretical models have been developed where individuals know, directly or indirectly, the opponent's past behavior, and, in those cases, cooperation is usually attained by using a strategy that is able to punish defectors and thus sustain a fully cooperative equilibrium. Usually, but not always: the Folk-theorem for repeated games proves that there are many more equilibria besides the fully cooperative ones, including those in which one player obtains a payoff larger than the mutual cooperation payoff. In Section 2.3 we showed how such equilibria could be created in simple settings in which players are constrained to use simple memory-one strategies, provided that mixtures are observable. It is easy to dismiss these asymmetric equilibria as mere mathematical curiosities. It may be difficult to imagine real settings in which individuals try to form a reputation for playing a mixed strategy. This is probably the reason why in the large literature on information in the Trust Game, deviations from the fully cooperative equilibria are never explained in terms of other equilibria of the repeated game.

We tested this conclusion with an experimental design and found it wanting. Subjects indeed tend to alternate between Reward and Not Reward more often when they know that the *frequencies* with which they chose Reward is going to be disclosed to the next opponent, compared to the case in which only their last choice is observed. However, we found a substantial fraction of subjects who were willing to randomize between Reward and Not Reward even in the Last-treatment. If anything, we found that subjects try too much to extort a larger payoff from their partners.

At the same time, our data show that in the attempt to exploit the other player's trust, some subjects played Reward with a too small frequency and ended up with a very low payoff, for themselves and for the opponents with whom they interacted. This result contains what we think is an interesting message for those who investigate the solutions to social dilemmas based on repeated interactions and reputations. Any such solution contains an element of conflict in it, because parties are likely to have conflicting interests on the several NE that emerge when a social dilemma is played repeatedly. Failures in securing an efficient equilibrium are usually attributed to player's impatience, or lack of information like in game with imperfect monitoring. Our results shows that sometimes efficient equilibria fail to be reached because some players try to secure for themselves a larger share of the benefits of cooperation. This is an important topic that surely deserves further investigation.

# Appendix

## 2.A More Figures




Figure 2.A.1: Box plot of the final frequency of Reward for each subject

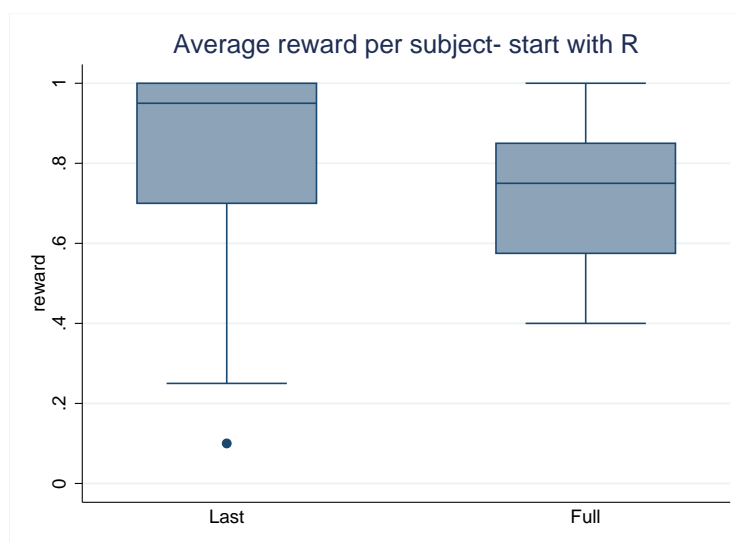


Figure 2.A.2: Box plot of the final frequency of Reward for each subject starting with R

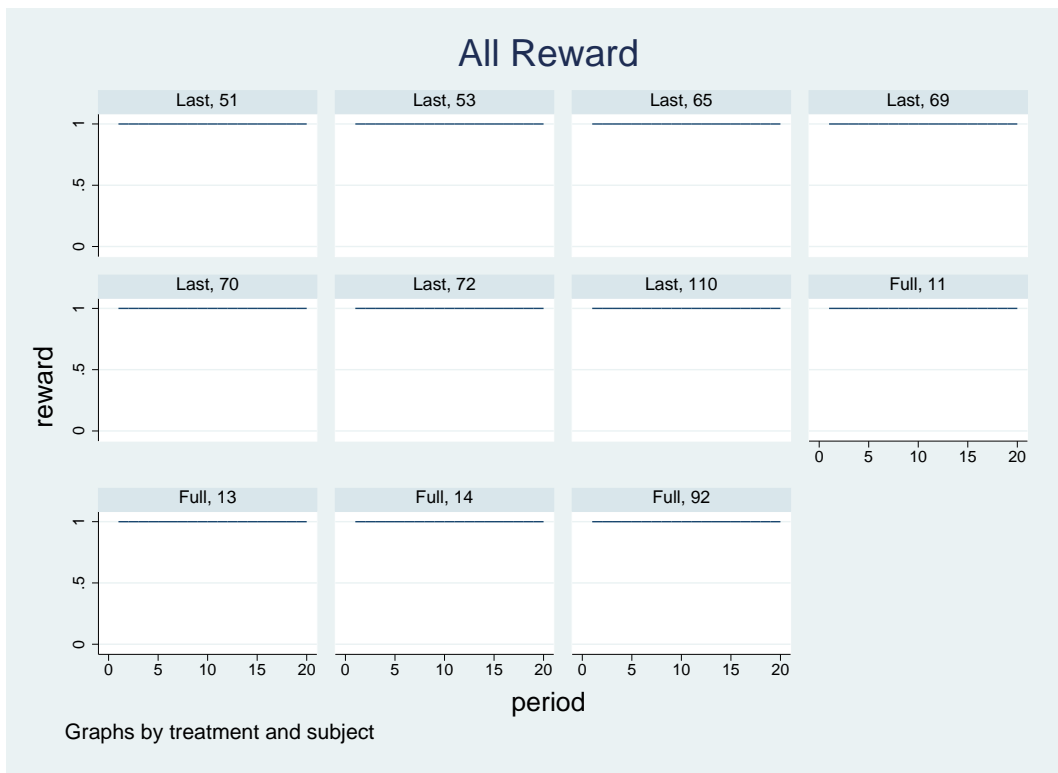


Figure 2.A.3: Always reward:  $f = 1$

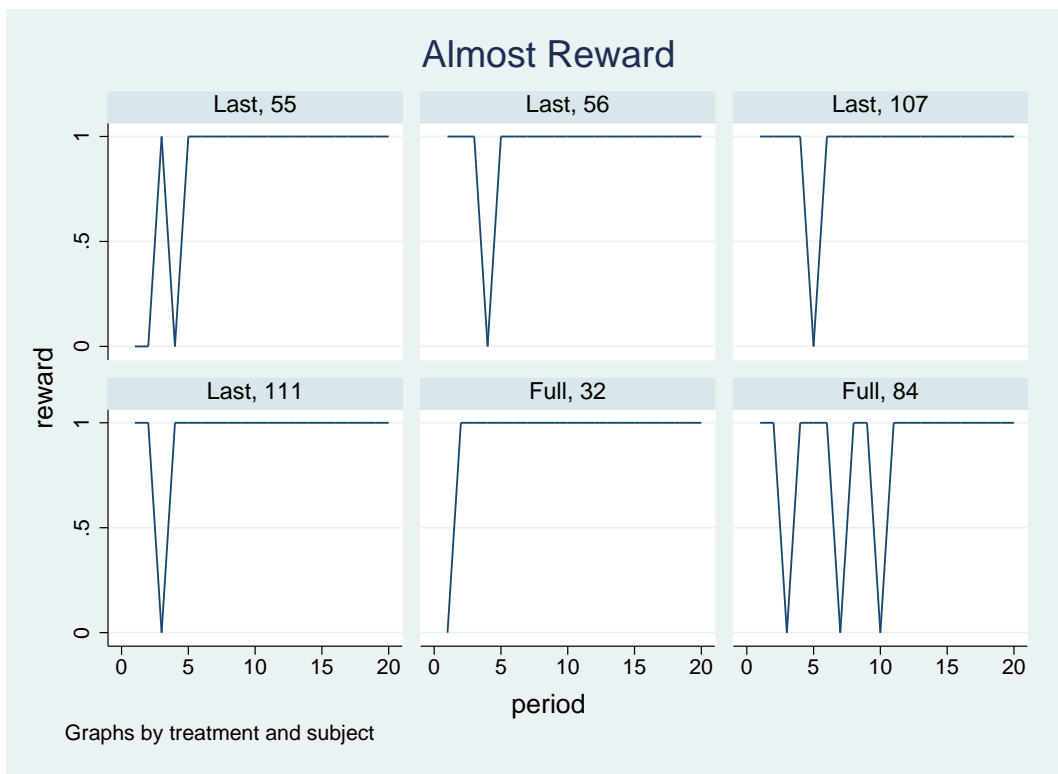


Figure 2.A.4: Almost reward:  $f = 1$  in the second half of the game.

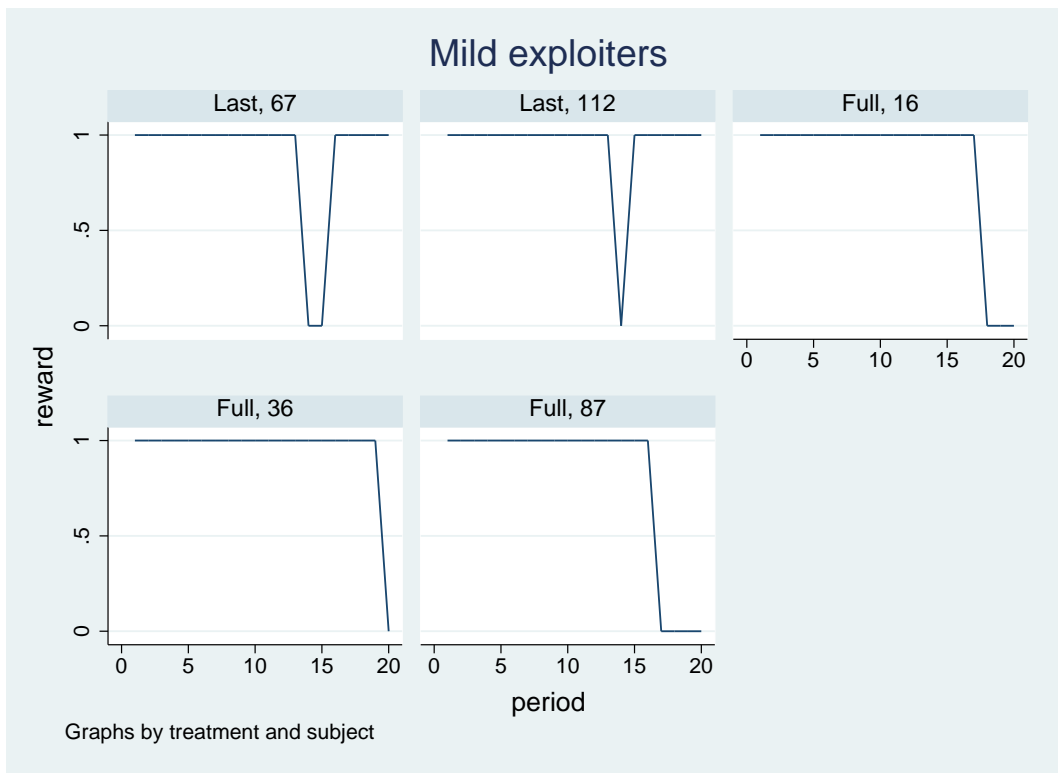


Figure 2.A.5: Mild exploiters:  $f = 1$  in the first part of the game, and  $f < 1$  in the second part of the game

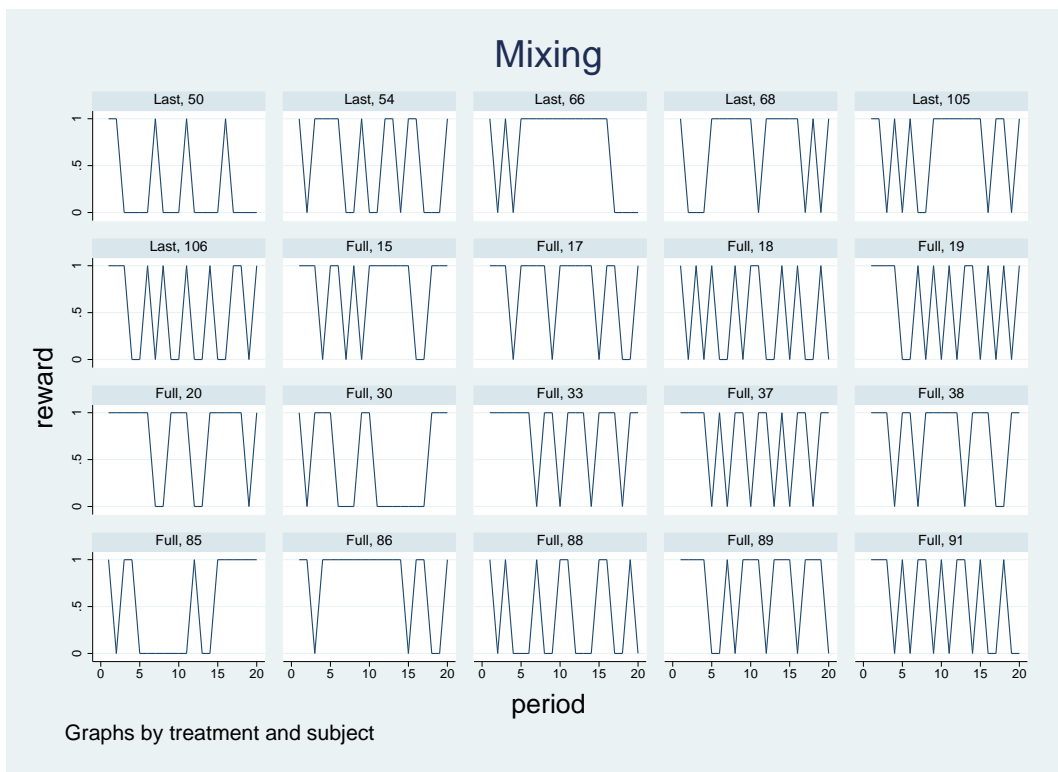


Figure 2.A.6: Exploiters:  $0.25 \leq \min(f) < 1$

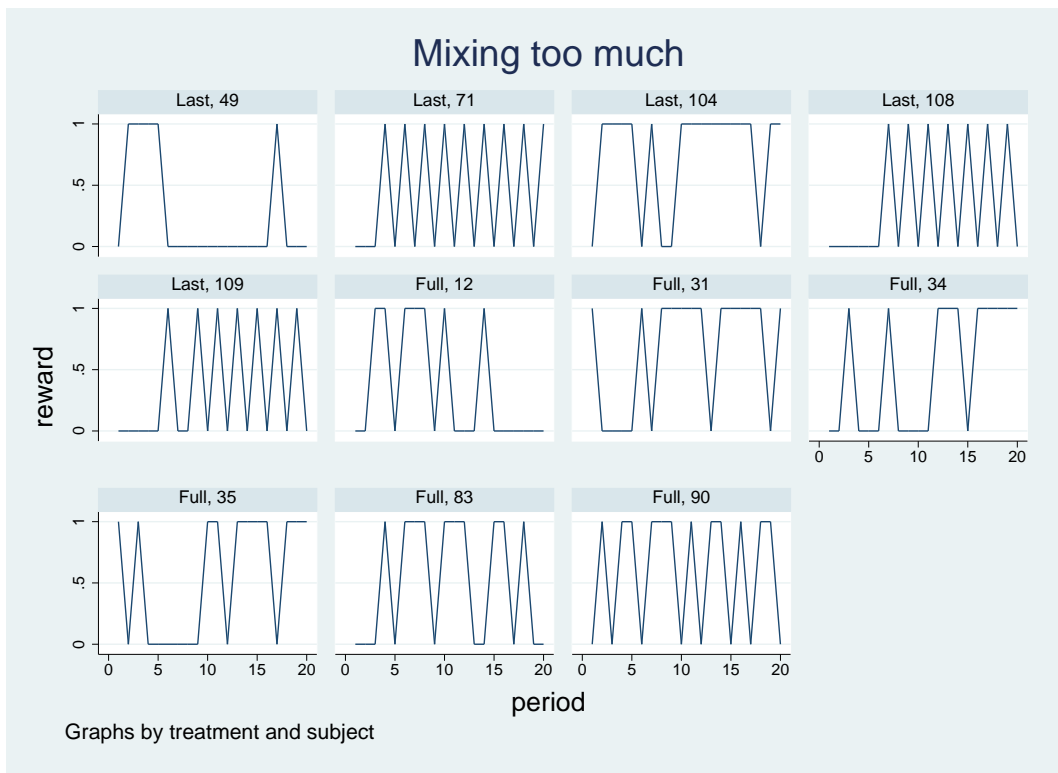


Figure 2.A.7: Hard exploiters:  $\min(f) < 0.25$ .




Figure 2.A.8: Almost not reward:  $f < 0.25$

## 2.B Instructions

### Introduction

Good Afternoon and thank you for participating in this experiment. It won't be difficult, there won't be any tricky questions, but you have to follow the instructions carefully. Your answers will be anonymous, and it will not be possible to know the identity of who gave the single answer. During the experiment you are not allowed to talk with other participants, and if something is not clear in the instructions, just raise your hand and ask the experimenter for further explanations. If something is not clear, again just raise your hand and ask for explanations to the experimenter.

**Participants and roles** At the beginning of the experiment half of participants will take the role of "Partecipante UNO" (participant one), and half the role of "Partecipante DUE" (participant two). Each participant will keep his role until the end of the experiment. During the experiment, each Partecipante UNO will be matched with a Partecipante DUE, and both of them will have to make some choices that will allow them to gain some tokens.

**Structure of the experiment** The experiment consists in two parts, A and B. You will receive more detailed instructions at the beginning of each part. The number of tokens that you will get at the end of the experiment will be determined by the choices that you and the participants with whom you are matched will do. In the two parts of the experiment, you will face the same problem, but you have to make different choices. Specifically:

- in the first part you can make a new choice every time you meet a new participant;
- in the second part you can choose only once, and the same choice will be automatically implemented every time you will meet a new participant

**Payments** You will receive 3 euro for participating in the experiment. Moreover, during the part A of the experiment you will have the chance to receive some tokens, that will be later converted in Euro, at the ratio: 1 token=0,10 euro. During the part B of the experiment, you will have the chance to receive 10 euro, according to the procedure that we will explain later.

## Instructions for the part A of the experiment

**Participants and choices** The part A of the experiment consists in an indefinite number of rounds. At the beginning of each round, each Partecipante UNO will be matched with a Partecipante DUE. Both participants have to make some choices that will determine the number of tokens they will receive in each round. Specifically:

- the Partecipante UNO has to choose between Dentro (IN) and Fuori (OUT);
- the Partecipante DUE has to choose between Alto (HIGH) and Basso (LOW).

To understand the relation between choices and tokens, during the experiment we will give you a more intuitive representation of payments: a table in which your own payment will always be the first value of each box (in red).

		Partecipante DUE	
		ALTO	BASSO
Partecipante UNO	DENTRO	4, 4	0, 6
	FUORI	1, 1	1, 1

		Partecipante UNO	
		DENTRO	FUORI
Partecipante DUE	ALTO	4, 4	1, 1
	BASSO	6, 0	1, 1

The table on the left will be shown to Partecipante UNO, while the one on the right will be shown to Partecipante DUE. In each box there are the number of tokens for Partecipante UNO (first value, in red, in the left table), and the number of tokens for Partecipante DUE (second value, in black, in the left table). Thus:

If UNO chooses Dentro and DUE chooses Alto, both participants get 4 tokens;

If UNO chooses Fuori and DUE chooses Alto, both participants get 1 token;

If UNO chooses Dentro and DUE chooses Basso, UNO gets 0 tokens, and DUE gets 6 tokens;

If UNO chooses Fuori and DUE chooses Basso, both participants get 1 token.

To make your choice you only have to press the button showing your preferred option.


**Matching** When all participants have made their choices, you will know how many tokens you earned for that round. Then, a new round will start and each Partecipante DUE will be matched with a new Partecipante UNO. Is it possible to meet the same participant again, but you have to wait at least for 10 rounds. Thus, if your role is DUE, and in the present round you meet a particular UNO, you are sure not to meet the same Partecipante UNO again for at least 10 rounds (the same reasoning applies if your role is UNO).

During the experiment it is not possible to identify a single participant, but every reference will be done to the roles e.g. “you are matched with a new Partecipante UNO (or DUE)”.

**How long is the experiment?** The experiment consists in an indefinite number of rounds.

You will participate to 20 rounds for sure, after which the experiment will continue with a probability of 50%. The end of the experiment will be thus determined by the result of a random extraction at the end of each round after the 19th. One of the experimenter will put in a box two pieces of papers, with numbers 1 and 2. At the end of each round one number is chosen at random. If the number is 2, the experiment is over. Otherwise, the number is put back in the box and the experiment goes on.

**Information** At the beginning of each round the Partecipante UNO will receive some information on the past choices of his current opponent. Specifically, the Partecipante UNO will be informed over the percentages with which the Partecipante DUE chose the options Alto and Basso in previous rounds (of course this information will be available only from the second round). For example, the percentages 50%, 50%, imply that the current Partecipante DUE chose Alto and Basso the same number of times in the past. The percentages 33%, 66% imply instead that the current Partecipante DUE chose, on average, Alto twice every three times, and Basso once every three times. To better understand this information, on the monitor of Partecipante UNO will appear a bar graph, indicating the percentages with which the Partecipante DUE chose Alto and Basso in the past:



In this example, the bar graph shows the percentages (66.7%, 33.3%), thus the Partecipante DUE chose, on average, Alto twice over three times, and Basso once over three times.

The Partecipante DUE doesn't get any information about the past behavior of the Partecipante UNO with whom he is matched. On his monitor will appear a synthesis of his own previous



choices, represented by the percentages with which he chose Alto and Basso in the previous rounds. Finally, every participant can see a bar showing the average number of tokens earned from the beginning of the experiment until the current round.

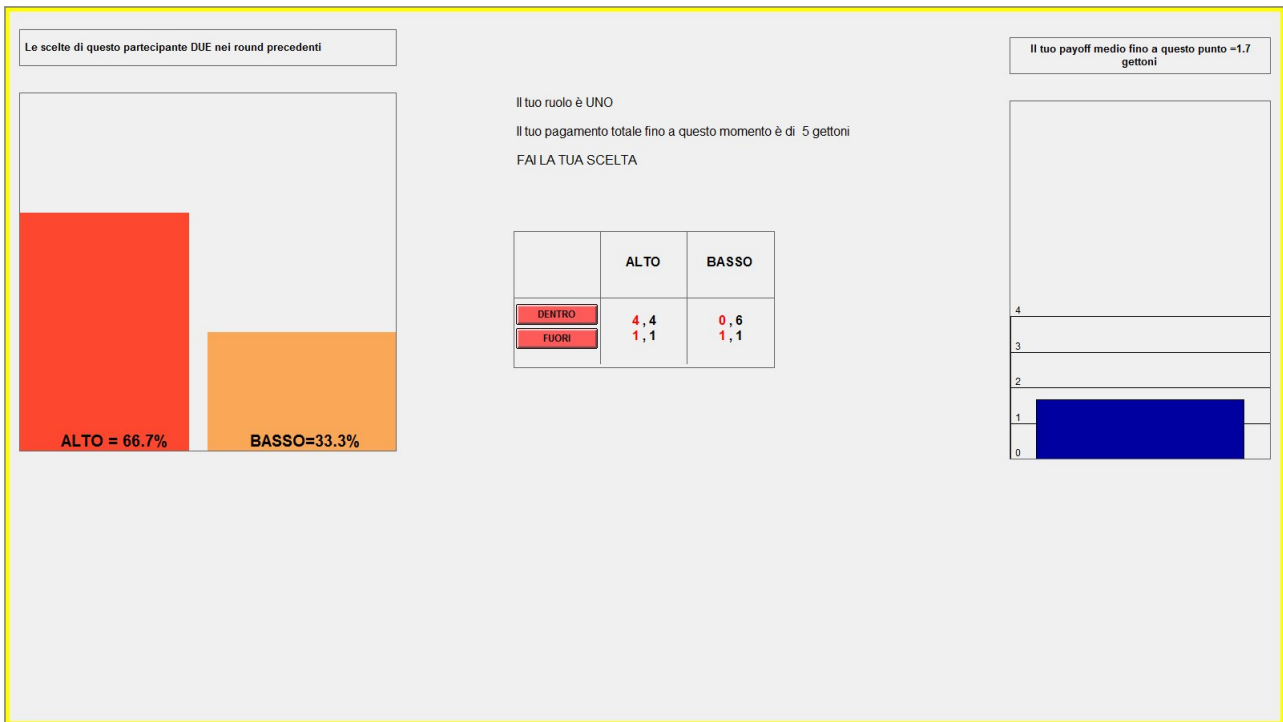


Figure 2.B.2: Example of Partecipante Uno's choice window

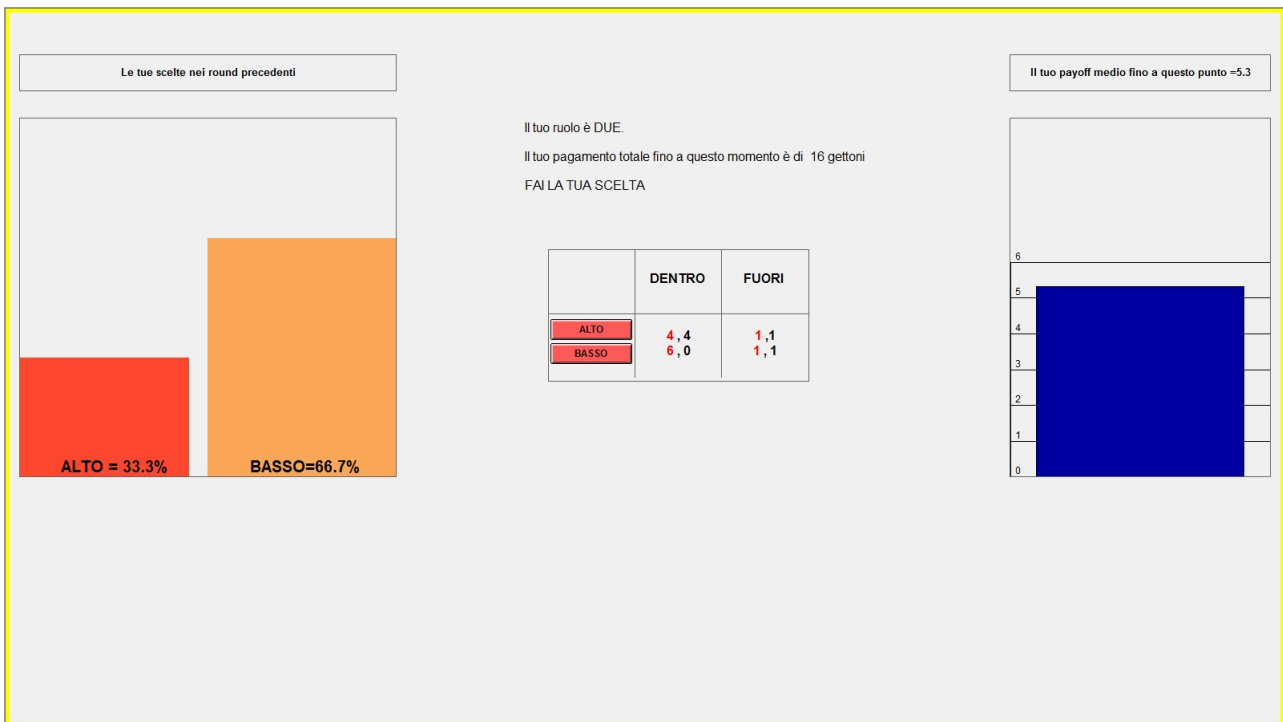


Figure 2.B.3: Example of Partecipante Due's choice window

**Payments** The total payment is equal to the sum of tokens earned in each round. At the end of the experiment, the total of tokens will be converted in Euros and you will get the corresponding amount of money.

## Control Questions

Now we ask you to answer to some control questions.

When all participants will have correctly answered to all questions, the experiment will begin.

1. In one round Partecipante UNO chooses Dentro and Partecipante DUE chooses Alto:
  - How many tokens does Partecipante UNO get in this round?
  - How many tokens does Partecipante DUE get in this round?
2. In one round Partecipante UNO chooses Fuori and Partecipante DUE chooses Alto:
  - How many tokens does Partecipante UNO get in this round?
  - How many tokens does Partecipante DUE get in this round?
3. In every round you will be matched with the same participant. True or False ?
4. At the beginning of every round, Partecipante DUE is informed over Partecipante UNO's previous choices. True or False ?
5. At the beginning of every round, Partecipante UNO is informed over Partecipante DUE's previous choices. True or False ?
6. You are at the first round, which is the probability the experiment will go on?
  - 100%   • 30%   • 50%   • 70%   • 0%
7. You are at the 28th round, which is the probability the experiment will go on?
  - 100%   • 30%   • 50%   • 70%   • 0%

# Chapter 3

## You go first: Coordination in a simulated Crowdfunding Experiment

with Luigi Mittone (University of Trento),

Azzurra Morreale and Mikael Collan (Lappeenranta University of Technology)

In this chapter we present the results from a simulated crowdfunding experiment, designed to assess the impact of *cheap* (i.e. no payoff-relevant) information on coordination. We consider a dynamic coordination game with multiple options. We run the experiment in Italy (with a relatively homogeneous sample), and in Finland (with a more heterogeneous sample). We first show that, with learning, there is a clear tendency for subjects to adopt the weakly dominant strategy “You go first: I’ll wait and see what the majority does”. Our main result is that only in the presence of *cheap* information, the same project is chosen across sessions, independently from the degree of homogeneity of the sample. This suggests an important role for *cheap* information in making a project salient. Our second result is that the presence of *cheap* information makes coordination easier in the very first rounds, by increasing contributions on the same project at the beginning of the game. However, in the presence of uncertainty, disclosing *cheap* information is detrimental for coordination, decreasing contributions both at the beginning and at the end of the game. Applied to the case of crowdfunding, the present results suggest that, when multiple projects are available, it is important for designers to disclose information that is not directly payoff-relevant, but that is able to make a project focal.

## 3.1 Introduction

Crowdfunding (CF) is the practice of gathering resources by collecting small amounts of money from a large number of people, that is by asking funding to “the crowd”. Historical crowdfunding campaigns include the funding of the pedestal for the Statue of Liberty, organized by Joseph Pulitzer<sup>1</sup>, and the “march of dimes” campaign, lead by Franklin D. Roosevelt to fund the fight against polio. Both are successful examples of crowdfunding efforts to mobilize massive crowds in the financing of large common-good projects in the time before the Internet. Since the days of these two examples, crowdfunding has gained momentum, and during the last ten years, fueled by the opportunities brought about by the Internet (including web-based crowdfunding platforms such as Kickstarter.com), the number of crowdfunded investments and the volumes of funding have dramatically increased and are in the billion dollar category <sup>2</sup>.

The rapid growth and increasing importance of crowdfunding makes it an interesting object of study from many points of view that include finance, information systems, and decision-making among others. In particular, existing empirical studies have investigated the determinants of successful fundraising (Mollick, 2014; Colombo et al., 2015), with some of them focusing on backers’ contributions patterns to projects over time (Agrawal et al., 2014; Cholakova and Clarysse, 2015). While the extant literature has enhanced our understanding of the underlying dynamics behind the behavior of the backers of successful projects<sup>3</sup> (Roma et al., 2017), there are still open questions regarding what motivates individuals in funding projects.

For instance, one highlighted aspect of the investment behavior is herding, i.e. the tendency of people to behave as “the majority does”. People appear to be more likely to fund projects that have already received many contributions from others. However, whether the herding behavior is “rational” (i.e. when the accumulated capital is interpreted as a signal of project’s quality) or “irrational” (i.e. when backers blindly follow the behavior of the crowd) is still an open issue in literature (Moritz et al., 2015). Moreover, it is not clear whether backers’ behavior is primarily driven by projects’ characteristics, and to which extent it can be considered an instance of herding behavior (Brüntje and Gajda, 2015). Indeed, being able to attract funding

---

<sup>1</sup>He managed to collect around 150000 dollars from more 120000 than people.

<sup>2</sup>Since its launch in 2009, Kickstarter.com alone managed to collect 3.5 billions of dollars from more than 14 million backers (i.e. the contributors of the crowdfunding campaigns), funding more than 140 thousands projects.

<sup>3</sup>Fundings are usually high during the beginning and towards the end of the campaign, while they slow down during the middle of the campaign (Kuppuswamy and Bayus, 2018)

in the very beginning of the campaign is crucial for a project to succeed, as those projects that receive the higher contributions in the early stage of the campaign are also more likely to reach the threshold, with some of them reaching the threshold in a couple of days, suggesting an important role for herding behavior (Agrawal et al., 2014).

On the one hand, this might lead to think that early investing is a winning strategy, since it can serve as a signal to induce other individuals to choose the same project, thus starting the herding effect. On the other hand, if the backer is budget constrained, he might prefer to wait, not to have his resources blocked in a unsuccessful project, and to later choose a project with a higher chance of being funded. This free riding behavior might cause a coordination failure among backers: if, as it seems to be, early funders generate a valuable (but noisy) signal for later ones over a project's chance to succeed, all investors might have an incentive to wait and see what others do (Agrawal et al., 2014).

Thus, in crowdfunding there are two different sources of risk: a project is subject to the risk of not being funded (arising from the coordination problem), and to the risk of being a failure (arising from the asymmetry of information that characterizes crowdfunding campaigns). It is then important to study the motives behind backers' behavior in the presence of both *market* uncertainty and *strategic* uncertainty. Nonetheless, the evidence on what drives backers' decisions is still scarce, as it is not possible to know from the web platform's data which information was available by the time of the investment's decision.

We design an experiment to try to overcome this difficulty. Our setting allows us to analyze whether successful projects are funded by backers who are mostly subject to herding effect, and to what extent those projects have some salient characteristics that might prompt backers' coordination. In the experiment, subjects play a dynamic coordination game: they can choose to invest, over several rounds, in one of a number of projects that have to reach a minimum funding threshold before a given deadline. The decision to depart from simultaneous moves games, and to give subjects the possibility to choose in a dynamic contest, allowed us to better simulate a real crowdfunding process, and to control whether some form of herding behavior was taking place. Moreover, it allowed us to study which factors, if any, might affect the speed of coordination, i.e. the time at which a project reaches the threshold.

Specifically, in our baseline subjects make their choices in front of identical projects, i.e. projects that have the same return in case of success, and no further information was given.

In the first treatment we add non payoff-relevant information to the description of the projects. We disclose information about the designer of the project, the institutional sponsor, and the rating of the different projects. On the one hand, the presence of information might induce some belief coordination over the project to be funded, speeding up the coordination process. On the other hand, information might increase the noise and hinder the coordination process, if subjects interpret different pieces of information as salient.

In the second treatment, we add payoff-relevant information in the form of uncertainty<sup>4</sup> over the crowdfunding campaigns future prospects, which in turn may affect backer's propensity to finance projects. Indeed, Ahlers et al. (2015) show that the higher the level of uncertainty about the start-up's future prospects, the lower is the success in equity crowdfunding. Therefore, risk preferences might have an impact on subjects' contributions. To analyze this possibility, we modify the baseline by introducing uncertainty over project returns: each project is characterized by a certain level of risk, but all projects have the same expected value.

In the final treatment we consider the most complete setting, by including both *cheap* information and payoff uncertainty. This treatment allow us to test the combined impact of both market uncertainty and project's information on subject's behavior.

For each treatment, we run three market sessions with different groups of projects, to investigate whether some form of learning takes place. Moreover, we collect data from two countries: Italy (University of Trento) and Finland (Lappeenranta University of Technology). As we explain in Section 3.4, the extent to which providing information helps or hinders the coordination process may depend on the degree of homogeneity of the sample. Thus, we run the experiment in two different settings: in Italy we recruit Italian students, while in Finland we recruit Exchange students, i.e. students from all around the world. This allowed us to compare the impact of information in two samples that may well be characterized by different levels of homogeneity.

Overall, subjects succeed in funding a project in all treatments<sup>5</sup>. However, learning causes the number of first movers to decrease, as subjects tend to wait and see what the majority is doing, supporting that there is a difference in behavior between seasoned and new investors.

Introducing information, either non payoff-relevant, or in the form of uncertainty over projects' returns, have effects on project selection and on participants' coordination.

---

<sup>4</sup>For the ease of exposition, in the following we will refer to no-payoff-relevant information as *cheap information* (or simply as *information*) and to payoff-relevant information as *uncertainty*.

<sup>5</sup>This is not surprising, given the parameters we used in the experiment, see Appendix 3.A.

Specifically, introducing cheap information increases the coordination at the very beginning of the game (i.e. the number of subjects choosing the same project in the first round), even if it does not affect the overall coordination (i.e. the number of subjects choosing the winning project), relatively to the baseline. More importantly, when provided with such information, subjects choose the same projects across market sessions. Taken together, these considerations suggest the importance of *cheap* information at the very beginning of the game, determining which project is going to win. Afterwards, behavior appears more driven by herding effect.

Introducing uncertainty over the project's returns seems to increase the coordination at the very beginning of the game, but it reduces the overall coordination, relatively to the baseline. Indeed, in the presence of uncertainty, risk preferences play a role as subjects have a stronger incentive to signal a specific risky project. This is to say that a risk taker player can try to push the other players to bet on the riskiest project, vice versa a risk averse player can implement a similar strategy but supporting the less risky project. In addition, in the presence of uncertainty, a higher number of participants preferred not to invest.

Finally, we show that the combined effect of cheap information and uncertainty is even more detrimental for coordination, as both coordination in the first round and the overall coordination are lower, relatively to the baseline.

Thus, we not only provide further evidence that, with learning, subjects tend to move towards weakly and risk dominant strategies, but we also provide new evidence over which factors might affect the timing of subjects' decisions.

As such, the results presented are very relevant for crowdfunding investors, companies seeking funding via crowdfunding platforms, companies that provide and manage crowdfunding services, and important to the disciplines of finance and decision making generally. The reported findings are new and, to the best of our knowledge, this is the first time similar results are reported in the context of crowdfunding.

## **3.2 Literature Review**

Previous studies on crowdfunding mainly used data from web platforms to analyze backers' behavior, and there seems to be a general agreement that the behavior of experienced and unexperienced backers differ. Experienced backers tend to focus on project's characteristics

in order to decide which project to fund, even if it is still not clear whether they are able to choose successful projects, or if a project became successful because they choose it (Baum and Silverman, 2004; Mollick, 2014; Kim and Viswanathan, 2018).

Indeed, one highlighted aspect of the investment behavior is herding, i.e. the tendency of people to behave as “the majority does”: people appear to be more likely to fund projects that have already received contributions from others (Agrawal et al., 2011; Burtch, 2011). Evidence of herding behavior using field data was found, among others, by Herzenstein et al. (2011), Zhang and Liu (2012), Burtch et al. (2013), and Ward and Ramachandran (2010).

A rationale for herding was given by Mäschle (2012), who noted that backers might use a free-rider strategy and wait to invest only on those projects that already have enough funding, in order to avoid time-costing searching and evaluation activities. He noticed that this behavior can reduce individual costs of information production, but raises the risk of irrational exuberance. Moritz et al. (2015) interviewed several backers in equity crowdfunding and found that not only herding behavior was present, but also that investors were aware of it. The study also confirmed the importance of the very start of the campaign. During this phase a project may collect funds mainly from family and friends, and having little funding after a week may be a signal that not even the closer acquaintance of the designer trust him with the project, discouraging further investors (Belleflamme et al., 2010).

Even if there are several types of crowdfunding (CF)<sup>6</sup>, they all share the same basic characteristics, which resemble a coordination game: a project must reach a funding threshold before a deadline, otherwise the eventual backers are refunded. Thus, several papers have analyzed the factors that might influence contributions and coordination in laboratory experiments. For instance, Lin et al. (2013) investigate how information about pledgers’ friends – which indirectly is a signal of credit quality - affects contributions. In a field experiment, Burtch et al. (2015) highlight that contribution behavior is influenced by the lack of anonymity for the contributors. Finally, Hashim et al. (2017) analyze the effect of information about previous contributions on coordination. Our paper relates to this stream of research by examining how providing information about the projects/funder, and the presence of stochastic outcomes (implied by the

---

<sup>6</sup>Donation CF (e.e. [makeawish.it](http://makeawish.it)), which involves most commonly charities and humanitarian projects; Reward based CF (e.g. [kickstarter.com](http://kickstarter.com)), where backers of a project receive a non-monetary reward, which usually depends on the amount backed; Lending CF (e.g. [prosper.com](http://prosper.com)), where people lend money to other people or business in need, with the promise of having back the amount plus the interest rate; Security CF (e.g. [wefunder.com](http://wefunder.com)) where people are directly involved in the company, buying its debt or equity (Purcell, 2014; Steinberg et al., 2012)



uncertainty over the campaign future prospects), affect coordination.

Moreover, since Bagnoli and Lipman (1989)'s work, a vast body of experimental literature has examined the effects that introducing a threshold has on contributions in coordination and public good games (Corazzini et al., 2015). The majority of these works has mainly focused on the effects of introducing alternative refund levels (the amount returned to subjects in case the threshold is not reached) on total contributions (see e.g. Isaac et al., 1989; Cadsby and Maynes, 1999; Coats et al., 2009; Cartwright and Stepanova, 2015; Cason and Zubrickas, 2017). However, while the afore-mentioned works consider a setting where a single project can be funded, we consider multiple projects and focus on the effects that introducing information has on contributions and coordination. Although this issue is relevant in the real world, scant attention has been devoted to analyzing subjects' behavior in the setting of multiple mutually exclusive projects.

One notable exception is represented by the work of Corazzini et al. (2015), which considers the effect of increasing the number of projects on the level of coordination. Perhaps not surprisingly, the authors found that coordination becomes more difficult when subjects face multiple options, even if this effect is mitigated when one of the options is Pareto dominant. In a similar vein, Wash and Solomon (2014) implement an experiment where subjects can simultaneously decide whether to fund one or more projects, each with a different funding threshold.

We significantly depart from those works in that we incorporate the important time dimension: in our setting subjects have more than one period to decide, knowing the choices made by the other subjects in the previous periods. Previous attempts to introduce time in coordination games were made, for example, by Gächter et al. (2010) and Coats et al. (2009). They allowed subjects to choose sequentially (with random positions) whether to fund a project. This increased the success rate in large groups, even if it had an opposite effect when the game was played in couples. Sequential decisions are easier to analyze, but in the context of crowdfunding it is important for subjects to decide not only *if*, but also *when* to invest in a project.

Closer to our setting is Solomon et al. (2015), where subjects had a window of 60 seconds to choose whether to fund a project, knowing others' decisions in real time.

The main differences with our setting is that Solomon et al. (2015) considered a public good game, where subjects have different endowments for different projects, thus relaxing the coordination problem. This means that projects in their setting are not in competition to each

other for contributions, i.e. all the projects can be funded. Moreover, in their setting the same project has different returns for different subjects, which does not allow to analyze whether timing decisions are dependent on subjects' risk preferences or projects' characteristics.

In our experiment, only one project can be funded (as in Corazzini et al., 2015), and subjects have more than one round to choose (as in Solomon et al., 2015). Contrary to those studies, projects in our setting have the same (expected) return, and we provide subjects with projects' information that can be either payoff-relevant (uncertainty over the campaign future prospects) or non payoff-relevant (e.g. information regarding the designer). The presence of such information can make a project more salient compared to the other projects. In other words, it becomes a focal point as in Schelling (1960), thus promoting coordination (Mehta et al., 1994; Crawford et al., 2008; Isoni et al., 2014; Parravano and Poulsen, 2015).

Finally, we analyze how such information influences convergence to a risk dominant strategy (Harsanyi, Selten, et al., 1988), which, in our setting, corresponds to a "you go first, I'll wait and see what the majority does" strategy. There is a huge experimental literature showing convergence not only to payoff dominant outcomes but also to risk dominant ones (e.g. Van Huyck et al., 1990; R. W. Cooper et al., 1990; Heinemann et al., 2004; Cachon and Camerer, 1996; Bornstein et al., 2002; Blume and Ortmann, 2007). However, differently from our experiment, the majority of these works analyzes setting characterized by Pareto ranked equilibria. We contribute to this literature by examining how providing information and/or introducing uncertainty affects convergence on risk dominant strategy in a dynamic setting characterized by multiple symmetric and asymmetric equilibria.

### 3.3 Experimental Design

The experimental design is based on the well known coordination game (R. W. Cooper et al., 1990; Fehr and Gächter, 2000; Morris and Shin, 2002; Heinemann et al., 2004) with a threshold (Heinemann et al., 2009). Subjects have to choose one project among those available, and if enough subjects choose the same option, the project is funded and can generate profit for the subjects who chose it.

The novelty in the design, which allows us to better simulate the herding effect in a crowdfunding market, is the time dimension: subjects in the experiment have more than one period to decide,

knowing what their opponents did in the previous periods. They can invest in a project only once, and once they do, they can not change their decision in subsequent periods. The choice to depart from standard simultaneous games is driven by the need to make explicit the trade off between waiting for the project selected by the majority, and choosing in the early rounds to make the favorite project more likely to be selected.

We gave to each subject one share to invest, and the parameters were such that only one project could reach the threshold, thus giving an effective incentive to the subjects to coordinate.

The experiment has a baseline and three different treatments (see Table 3.1), in a 2x2 design. In the baseline, no information is given about projects, which are referred to by names of different colors, and all having the same, sure, return, in case the threshold for the project is reached.

In the first treatment (Cheap-Info-treatment, or Info-treatment), we keep sure the return, and we add information about the projects. In the second treatment (Uncertain-treatment), we give no information, but the return of the projects is uncertain: projects have the same expected value, but different levels of risk and return. Finally, in the last treatment (Combined-treatment) we introduce both cheap information and uncertainty over outcomes.

		Uncertainty	
		No	Yes
Information	No	Baseline	Uncertain
	Yes	Cheap Info	Combined

Table 3.1: Baseline and Treatments

In order to define the risk and the characteristics of each project, we run a “designing” experiment, where we asked participants to choose the characteristics of the projects that we later use in the coordination experiment. Then, we run main experiment, involving subjects different from the ones participating to the “designing” experiment, where participants were randomly assigned to one of the four afore-mentioned treatments. In both cases, payoffs are expressed in EMU, Experimental Monetary Unit, which were converted 50:1 into Euro at the end of the experiment (i.e. 1 EMU = €0.02 ).

## Design of the projects

Each subject in the “designing” experiment<sup>7</sup> had to design a project by selecting one option for each of the following characteristics:

- Investor’s profit: this is the share of profits the designer is willing to give to the investor in the case of success. The designer could choose a value in the range 10% - 60%.
- Risk: in the treatments with uncertainty, if the project reaches the threshold, its ability to generate profits depends on the success of a lottery, which is the combination of a probability of success and a profit in the case of success. If the lottery is successful, the project generates a profit for both the designer and the backers of that project. Otherwise, both the designer and the backers get zero. The designer had to choose a lottery among several available, all having the same expected value (1000 EMU). He could choose a value for the risk of the lottery, i.e. the probability of success, in the range 10% – 90%<sup>8</sup>.
- Sponsor: in two of our treatments, we introduce cheap information which might help subjects to coordinate. The first piece of information is the sponsor of the project: the designer had to choose one sponsor among several proposed institutions<sup>9</sup>. Indeed, there is evidence in the literature of the important role played by “superior principals endorsements” in driving backers to support certain projects (Moritz et al., 2015).

At the end of this phase, we run a beauty contest with all the projects. Each designer had to vote up to 3 projects that, in his opinion, would receive the most votes among the designer as the projects that were more likely to reach the threshold. Thus, we were able to rank the projects from the most to the least successful. Such ranking was the second piece of information given to the subjects in the treatments with information. The final pieces of information were about the designer. We disclosed information about the gender, the education (i.e. the designer’s degree program) and the experience (i.e. in how many experiments the designer previously participated). Indeed, there is evidence in the literature that also designer characteristics are important for backers to decide whether to invest in a project (Vismara, 2016).

---

<sup>7</sup>The designing experiment was run in July, 2017, at the University of Trento (CEEL)

<sup>8</sup>Notice that to a risk of 10% corresponds the highest profit of 10000 EMU, while to a risk of 90% corresponds the lowest profit of 1111 EMU.

<sup>9</sup>The proposed institutions were: Bank, District, Insurance Company, Church and University.

Table 3.B.1 in Appendix 3.B gives some examples of chosen projects. Subjects in this phase were aware of the aim of the experiment, and the designer of successful projects were later called to receive the corresponding payments.

## Main Experiment

In all treatments of the main experiment, we had groups of 24 subjects, and we repeated the game three times (we called them three different *markets* or *market sessions*), to analyze whether, and how, subjects' behavior changes with learning. For each market session, we randomly selected one set of projects (already designed in the *designing* experiment) among those having the same expected value (i.e. the same share of the profits) for the investor. We kept the same set of projects for each market session across treatments (see Table 3.B.1).

## Baseline

In the baseline, we didn't give any additional information beside the name (i.e. a color) and the value of each project. At the beginning of each of the three market sessions, we gave one share (valued 150 EMU) to each subject. Subjects then had ten rounds to decide if, and on which project, to invest. The threshold was set at half the number of subjects plus one, and it was the same among all market sessions and all treatments. If a project reaches the threshold, every subject who chose that project would receive the corresponding value. Otherwise, eventual backers would receive 50 EMU as a partial refund. The decision of introducing a partial refund was motivated by the necessity of combining the common use in crowdfunding campaigns of refunding the investment if a project does not reach the threshold, which would imply a total refunding, with the fact that, once invested, the money is not in the direct disposal of the investor until the campaign is over. The partial refunding would thus mimic the cost of losing the availability of money for that period of time. This method was suggested by Corazzini et al. (2015) as "ideal to capture the refunding delay in crowdfunding" (p.19).

In every period after the first one, subjects were informed over the number of participants that, in all previous rounds, chose each project<sup>10</sup>, and they were allowed to choose a project even if it reached the threshold. Moreover, in every round subjects had to indicate how many

---

<sup>10</sup>This information was provided with a bar graph, indicating the level of funding for each project. The bar was red until the threshold was reached. At that point, the bar turned green and stop moving.

participants they believe would chose each project in the current round. This information was useful to understand whether subjects' behavior was consistent with their (stated) beliefs.

## Treatments

In the first (Info) treatment, we modify the baseline experiment by providing information about the projects and the designers. We disclose information about the sponsor and the rating of the different projects, and over the education, sex and experience of the designer of each project. Those pieces of information were summarized on a table (see Appendix 3.C), and subjects could check only one characteristic at the time. By clicking on the name of one characteristic, the values corresponding to each project was shown in the table, and those values remain visible until the subject would click on a different characteristic. This allowed us to control which piece of information the subjects choose to check, and how long they stayed on each characteristic, giving us a partial clue on the relative importance of the different characteristics.

In the second (Uncertain) treatment we modify the baseline by introducing uncertainty over projects' returns. As already mentioned, in each market the expected value of the projects is the same, but each project is characterized by a different probability of generating profits, in the case the threshold was reached. That is, we add market uncertainty to the strategic uncertainty characterizing the coordination game. Each project was linked to lottery, chosen by the designer in the designing phase, and for each project subjects could see the relative probability of generating profits, and the value of its profits. If a project reached the threshold, then we run the corresponding lottery. Only if the lottery was successful the subjects who chose that project could gain the relative payoff. Otherwise, they would get a payoff of zero.

Finally, in the last (Combined) treatment, we modify the baseline by merging the Info-treatment and the Uncertain-treatment. That is, we disclose to participants information about the projects and the designers, and at the same time the projects payoffs' are uncertain.

This experimental design allows us to study the role of *cheap* information in settings characterized by either certain or uncertain payoffs, as well as the role of uncertainty in settings with and without *cheap* information.

## Experimental procedure

For each treatment, we run two experimental sessions, one in Italy (CEEL - University of Trento) and one in Finland (LUT - Lappeenranta University of Technology).

The experiments were run from December 2017 until March 2018. A total of 192 subjects participated, 48 per treatment (24 in Italy and 24 in Finland), and participants were not allowed to participate more than once. The experiment was conducted using the o-tree software (D. L. Chen et al., 2016, see Appendix 3.C for the instructions).

After arriving to the laboratory, participants were randomly assigned to a computer. Instructions were given in Italian in the sessions in Italy, and in English in the sessions in Finland. Participants had to answer a series of control questions before being allowed to start the experiment. The software informed participants when a new market session was starting, and we didn't give any feedback between one market session and the next one<sup>11</sup>. In all treatments, after the final market, we elicited subjects risk preferences through the Bomb Risk Elicitation Task (BRET, Crosetto and Filippin, 2013)<sup>12</sup>.

After the BRET, one market session was randomly extracted by the computer, and the total payment of each subject was computed as the profit in the chosen market session, plus the (eventual) prize for the belief elicitation, plus the profit in the BRET, plus the participation fee. Indeed, we incentivize the belief elicitation in the following way: at the end of the experiment, we randomly selected one period of the game, and we computed the set of beliefs that were the most accurate i.e. closest to the actual number of subjects who invested in that period. A prize of 150 EMU was given to the participant whose beliefs were the most accurate.

The average payment was equal to €11.50, with a maximum of €25 and a minimum of €3 (including the participation fee of €3), for an experiment that took approximately 50-60 minutes.

At the end of the experiment, we asked subjects to answer a series of debriefing questions, and to fill out a short demographic questionnaire.

---

<sup>11</sup>Clearly this was meaningful only in the uncertain and in the Combined-treatments, since subjects didn't know whether the project they chose was successful (i.e. if it won the lottery) or not.

<sup>12</sup>In the BRET, every subject was shown 100 boxes; every subject knows that 99 boxes contain 2 EMU each, while the remaining one contains a bomb. Participants were asked to collect as many boxes as they wanted, knowing that if the box containing the bomb was collected, it exploded, causing earnings equal to 0 for this part of the experiment.

## 3.4 Predictions and Hypotheses

### 3.4.1 Theoretical Framework

Our model is an extension of Heinemann et al. (2009). They considered a group of  $n$  players choosing between a safe option (S), giving a payoff of  $y > 0$ , and a “risky” option (A), giving a payoff of  $r > y$  if at least  $k \leq n$  players choose the same option, and 0 otherwise.

In our case, the original setting is complicated by two aspects: there are several “risky” options available, and players have more than one period to make their decisions.

Consider first the case where players can choose simultaneously among one safe option S, and one out of a set of  $J$  risky options, all giving the same return  $r$  if at least  $k$  players choose the same option. Let  $x_j$  be the total number of players choosing the risky option  $j$ . The payoff from choosing a risky option  $j$  is then:

$$\begin{cases} r & \text{if } x_j \geq k \\ 0 & \text{if } x_j < k \end{cases}$$

and the expected payoff from choosing option  $j$  is  $\pi(j) = rq_j$ , where  $q_j$  is the probability that option  $j$  reaches the threshold, that is the probability that at least  $k - 1$  over  $n - 1$  players make the same choice. The (expected) payoff from choosing the safe option is  $\pi(S) = y$ .

With one safe option and  $J$  risky projects, the game has  $J + 1$  pure Nash equilibria, where all players choose the same option (either safe or risky).

With only one risky option (call it option A), Heinemann et al. (2009) show that, if players are risk neutral, there is one mixed equilibrium in which everybody choose the risky option with the same probability ( $p_A$ ) equal to  $p^*$ . In this case, the number of players choosing the same option follows a binomial distribution, and the probability that option A reaches the threshold, from the prospective of a player who is going to choose it, is:

$$q_A = 1 - \text{Bin}(k - 2, n - 1, p_A)$$

where  $\text{Bin}$  is the cumulative binomial distribution. To be an equilibrium, players should be indifferent between the safe option S and the risky option A, that is:

$$\pi(S) = \pi(A)$$



which implies that  $p^*$  must solve:

$$y = rq_A \rightarrow y = r(1 - \text{Bin}(k - 2, n - 1, p^*))$$

Consider now the case in which there are  $J > 1$  risky options. There are then  $J$  mixed equilibria where players mix between one risky option (with probability  $p^*$ ) and the safe one (with probability  $1 - p^*$ ). With multiple projects, there are also other two types of equilibria: in the first one, the safe option is never chosen, so that players mix among  $\hat{J} \subseteq J$  projects with a probability equal to  $1/\hat{J}$ ; in the second type, also the safe option might be chosen, so that each risky option is played with a probability equal to  $p^*$ , and the safe option with a probability equal to  $1 - \hat{J}p^*$ . For those two types of equilibria to exist, the return from the risky project should be high enough (see Appendix 3.A).

Consider now the case of two periods. In the first period, players can choose one of the risky options, or they can wait. In the second period, if a player already chose one risky option, his action set is empty. Otherwise, he can choose one of the risky options, or the safe one. A strategy in this case specifies an action to take in period 1, and an action to take in period 2, given the number of players who choose each option in period 1. Clearly, if already in period 1 one option reached the threshold, it is optimal for all players who waited in the first period, to choose that option in the second period. If the threshold is not reached in the first period, the coordination problem moves to the second period, and the simultaneous coordination game is played by all players who have chosen to wait in the first period.

The information available in period  $t$  is summarized by a vector  $\mathbf{x}^{t-1}$ , whose general component  $x_j^{t-1}$  is the total number of players choosing option  $j$  up to period  $t - 1$ . The final number of players choosing each option is then given by the vector  $\mathbf{x}^T$ .

With two periods, we denote by  $\Pi(\mathbf{x}^1)$  the equilibrium payoff when the threshold is not reached in the first period, and the information available in the second period is  $\mathbf{x}^1$ . That is,  $\Pi(\mathbf{x}^1)$  is the equilibrium payoff resulting from the simultaneous coordination game played among  $n - \sum_j x_j^1$  players, where each option has a threshold  $k_j = k - x_j^1$ . Notice that, if there is an option  $j$  s.t.  $x_j^1 \geq k - 1$ , then it is (strictly) dominant to choose option  $j$  in the second period.

This implies that the payoff from waiting in the first period is:

$$\begin{cases} r & \text{if } \exists j : x_j^1 \geq k - 1 \\ \Pi(\mathbf{x}^1) & \text{otherwise} \end{cases}$$

The expected payoff to choose one risky option in the first period is instead:

$$\begin{cases} r & \text{if } x_j^2 \geq k \\ 0 & \text{if } x_j^2 < k \end{cases}$$

Since  $\Pi(\mathbf{x}^1) \geq 0 \quad \forall \mathbf{x}^1$ , waiting in the first period is a weakly dominant action, as it gives to players the opportunity to observe the choices of the opponents. Clearly, if everybody waits, the final result is just to postpone the full coordination problem in the second period. Moreover, waiting and then choosing the safe option if the threshold is not met in the first period is the risk dominant strategy, which is also the strategy that subjects in experiments learn to play when the game is repeated (Van Huyck et al., 1990).

In this setting there are 2 types of (pure and symmetric) equilibria: full coordination on one of the (risky) options in the first period and full coordination on one of the (risky or safe) options in the second period. There are also  $J(n-1)$  asymmetric equilibria, where  $x_j^1$  players choose the risky option  $j$  in the first period (with  $n > x_j^1 > 0$ ), and the remaining  $n - x_j^1$  players choose the same option in the next period. Finally, there is a mixed equilibrium in which everybody wait in the first period and then play one of the simultaneous mixed equilibria.

Consider now the more general case of  $T$  periods and  $J$  risky options. The number of (pure and symmetric) equilibria is  $JT + 1$ . They are all weak, as any strategy profile that reaches full coordination on one option by the end of the game is an equilibrium giving the same payoff. Moreover, every strategy profile that also prescribes to choose one option whenever already  $k - 1$  players choose it, can form a subgame perfect equilibrium. Still, the weakly dominant strategy is to wait until  $k - 1$  players already choose one option, and the risk dominant strategy is to choose the safe option if the threshold was not previously reached.

In our setting, we have  $J = 3$  projects and  $T = 10$  periods. It follows from the previous reasoning that there are 31 pure and symmetric equilibria (and many asymmetric ones), as every strategy profile that implies full coordination will form an equilibrium. For instance, one equilibrium is that players fund one the projects in any round over the 10 available rounds. Another equilibrium corresponds to the case where no contributions are provided to any of the projects. Even if this equilibrium is risk dominant, it is dominated by any equilibrium in which the threshold is reached. In fact, contributing to one of the available projects can give a potentially higher payoff than not contributing. However, in the presence of strategic uncertainty (uncertainty regarding whether the other players will choose to contribute) and market uncertainty

(uncertainty over the future prospects of a project), choosing of not contributing represents the safest alternative. Also, in the presence of strategic uncertainty making “late” contributions, i.e. in the last rounds of the game, could be safer than making “early” contributions, i.e. in the first rounds of the game. When deciding between making early or late contributions people face a trade-off. On the one hand, players would wait to see whether other people contribute to any of the projects enough to reach the threshold. On the other hand, it might also be important to contribute immediately, because an “early” contribution may be used as a signal to others, i.e. to encourage them to contribute. Thus, if players are using a weakly dominant action, we should observe coordination in the last rounds of the game. Otherwise, our theory implies no particular pattern, neither in the choice of the option, nor in the timing of coordination. This implies that we should observe no differences between the baseline and the treatment with non payoff-relevant information. Moreover, if players are risk neutral, we should neither observe differences between the baseline and the treatment with uncertain payoffs.

The previous reasoning allows us to make the following theoretical predictions:

THP1 (**Coordination**): In all pure equilibria there is full coordination.

THP2 (**Weak dominance**): Late coordination if players use weakly dominant strategies.

THP3 (**Information**): Introducing *cheap* information does not change players’ behavior.

THP4 (**Uncertainty**): If players are risk neutral, introducing uncertainty does not change players’ behavior.

### 3.4.2 Herding and Behavioral Predictions

According to our theoretical predictions, introducing information should not affect coordination across treatments. Following a huge behavioral literature that speculates human decision-making to be less than fully rational, we expect that providing payoff-relevant or non payoff-relevant information might affect subjects’ choices, specially in the presence of herding behavior.

Recall from our theoretical framework that  $\mathbf{x}^{t-1}$  is the information available in period  $t$ , i.e. is the vector collecting the number of players choosing each option up to period  $t - 1$ . The expected payoff for a player who waited until period  $t$ , and chooses option  $j$  in period  $t$ , is  $rq_j$ , where  $q_j$  is the probability that option  $j$  reaches the threshold, i.e. the probability that at least  $k - 1$  players choose the same project, given that already  $x_j^{t-1}$  players chose that project, and

$n - \sum_j x_j^{t-1}$  players waited. Consider the case of one risky project. We have:

$$q_j^t = 1 - \text{Bin}(k - 2 - x_j^{t-1}, n - x_j^{t-1} - 1, p_j^t)$$

Given  $0 < p_j^t < 1$ , this probability is increasing in  $x_j^{t-1}$ , with  $q_j^t = 1$  if  $x_j^{t-1} \geq k - 1$ . If players are subject to herding behavior, they might overestimate the probability that one project reaches the threshold, given that it was already chosen by some players. We model herding behavior as a non-rational updating of the probability  $q_j^t$ : a player subject to herding behavior (call him H-player) will always choose the project that has the highest number of collected shares, as long as those are higher than  $\hat{x}$ , for some  $1 < \hat{x} < k - 1$ :

$$q_{j'}^t = \begin{cases} 1 & \text{if } x_{j'}^{t-1} = \max\{x_j^{t-1}\} \text{ and } x_{j'}^{t-1} \geq \hat{x} \\ 0 & \text{otherwise} \end{cases}$$

As soon as at least  $\hat{x}$  players choose the project  $j$ , an H-player will also choose that project. The highest the herding effect, the lowest the value of  $\hat{x}$ . Moreover, an H-player will never choose first, so that if all players are H-players, coordination can only happen on the safe option. Recall that, in a population of NH-players (i.e. players that are not subject to herding behavior) the weakly dominant strategy is to wait. In the presence of H players, this might not be true anymore: players that do not follow the herd might have an incentive to early invest and “start the herd”. Moreover, if there is uncertainty over projects’ payoff, and if NH-players are not risk neutral, they might have also the incentive to signal their favorite project to orient the herd in their favorite direction. Then, the coordination game is moved in the very first rounds, and is played by those that we call *first movers*, i.e. NH players.

In this setting, providing information has a non trivial effect. On the one hand, it might help those who do not follow the herd to better coordinate in the very first rounds: information can help first movers, if it is able to make one option more salient, thus facilitating early coordination. On the other hand, information might increase the noise, and, consequently, hinder the coordination process, as backers can be influenced by their own preference towards specific characteristics of projects that serve as a signal of quality of the projects themselves. More precisely, whether such information is able to facilitate coordination depends on its ability to make an option salient for the highest number of players, which might be more likely to happen in a homogeneous population. Thus, we might expect information to help coordination more in homogeneous settings, rather than in heterogeneous ones.

Introducing payoff-relevant information, in the form of uncertainty, might also affect behavior.

Recall that the risk in this setting is both strategic and payoff-related. A risk seeking player may have a strong incentive to signal his favorite (riskier) option, while a risk averse player may have an incentive to signal his favorite (less risky) option, even if the latter tendency is milder, as choosing at the beginning of the game has a higher strategic risk. Thus, if players have different risk preferences, the presence of uncertainty might hinder coordination. If instead they have the same (not risk neutral) preferences, uncertainty might help them in coordinating. Finally, introducing both payoff-relevant information (i.e. uncertainty) and non payoff-relevant information might hinder the coordination process, as it might be even more difficult for both first movers and followers to interpret an option as salient. In fact, in this case players face a trade-off that might hinder the coordination process, since a project being salient for some informational attributes may not be the favorite from the point of view of its riskiness.

The previous reasoning allows us to make the following behavioral predictions:

**BHP1 (Information):** Introducing cheap information might help the coordination process in homogeneous settings, i.e. in settings where different players are able to focus on the same piece of information. If information is effective in making one option salient, then the option with the same characteristics should be chosen across sessions.

**BHP2(Uncertainty):** Introducing uncertainty might hinder the coordination process if there is a lot of variance in players' risk preferences, as players might be tempted to signal different projects from the very beginning of the game. If introducing uncertainty is effective in making one option salient, the same risky option should be chosen across sessions.

**BHP3(Combination):** The combined effect of uncertainty and information might hinder coordination, especially in a more heterogeneous settings.

## 3.5 Results

In this research, we are mainly interested in analyzing the impact of cheap information (i.e. non payoff-relevant information about the project and the designer), as well as uncertainty, on the fundraising success. We first compute several success indicators and look at whether and how they differ across treatments. Specifically, given that in “the context of crowdfunding platforms, funding success is a multifaceted concept” (Ahlers et al., 2015, p.961), in addition to the traditional success indicators (whether a project reaches the threshold and how many subjects

invest on that project), we also consider the speed of coordination as an important success measure. We then analyze the effects of treatment manipulation on participants' behavior at the very beginning of the game (first round of the first session) and over the game.

## Samples

Two of our behavioral predictions strictly depend on the composition of the players' sample, since the ability of the available information to make an option salient depends on whether subjects in the sample are able to see the same piece of information as focal. This process should be easier if the sample is homogeneous in terms of subjects' background and risk preferences. Therefore, we need to try to assess whether the participants in our samples were more or less homogeneous in terms of the two aforementioned dimensions. All subjects in the experiment were university students. Recall that, in Italy, we recruit mostly Italian students, while in Finland we recruit Exchange students, i.e. students from all around the world. In Table 3.1 we report the main demographic variables, separately for the Italian and the Finnish sample, together with our measure for risk aversion, i.e. the number of boxes collected during the BRET, with a higher number indicating a lower risk aversion. Table 3.1 also reports the standard errors of the aforementioned variables for the two samples, which can be interpreted as a rough measure of homogeneity.

	Age	Experience	Faculty	Payoff	Boxes
ITA	22 (2.53)	10 (9.08)	2 (1.78)	450 (408)	43 (17.63)
FIN	27 (4.47)	2 (2.62)	3 (2.95)	396 (354)	41 (20.48)

Table 3.1: Sample demographics (standard errors in parentheses)

The ANOVA analysis, when applied to our measure of risk aversion, allow us to reject the null hypothesis that the Italian and the Finnish samples are equal (p-value < 0.1). The same holds for all the other variables (p-value < 0.01). Moreover, the Finnish sample has always a higher standard error than the Italian one, except for the variable Experience. The variable Experience measures the number of experiments a subject participated in the past. The majority of students in the Finnish sample participated only in a few experiments, and thus it is not surprising that it exhibits a lower standard deviation. For all the other variables (age, faculty, and risk aversion), the Finnish sample has a higher standard deviation than the Italian one.

We might then speculate that subjects in Italy were more homogeneous in terms of background and risk preferences, compared to subjects in Finland. Thus, in the Italian sample, we expect a higher coordination rate in the treatments with either *cheap* information or uncertainty.

### 3.5.1 Winning Projects

Our first theoretical prediction implies that we should have no reason to believe that the same option is chosen across treatments. Our first and second behavioral predictions instead imply that projects with the same piece of information or the same risk level should be chosen across sessions, if introducing information or uncertainty is effective in making one project focal. Thus, we start by analyzing the total number of shares that each project managed to collect in each session of every treatment, summarized in Table 3.2. Table 3.B.1 in Appendix 3.B gives an overview of projects' characteristics and risk levels.

Projects	Baseline		Information		Uncertainty		Combined		
	ITA	FIN	ITA	FIN	ITA	FIN	ITA	FIN	
Yellow <sub>1</sub>	2	1	23	15	22	18	23	6	First mkt session
Red <sub>1</sub>	1	3	1	2	1	0	1	14	
Blue <sub>1</sub>	21	20	0	4	1	5	0	2	
Yellow <sub>2</sub>	23	2	1	2	19	2	5	22	Second mkt session
Red <sub>2</sub>	0	20	23	19	2	0	18	1	
Blue <sub>2</sub>	0	2	0	0	2	21	0	1	
Yellow <sub>3</sub>	0	0	0	0	0	3	0	0	Third mkt session
Red <sub>3</sub>	0	1	1	0	17	0	1	17	
Blue <sub>3</sub>	0	0	0	1	1	0	0	0	
Green <sub>3</sub>	0	21	23	17	0	1	14	2	
Purple <sub>3</sub>	22	2	0	1	0	0	1	2	
Orange <sub>3</sub>	1	0	0	0	2	18	2	1	

Table 3.2: Number of subjects who choose each project by the end of each session.

In all market sessions subjects succeeded in funding a project (recall that the limit to fund a project was 13). As we can see in Table 3.2, the lowest number of participants who choose the winning project was 14, in the Combined-treatment. Indeed, on average, a lower number of participants chooses the winning project in the Combined-treatment, compared to all the other treatments. Notice that the Combined-treatment is the one where, according to our

behavioral prediction BHP3, coordination should be more difficult. Moreover, the highest number of subjects who choose the winning project was 23, in all the Italian sessions of the Info-treatment (and in one market session of the baseline). Notice that the Info-treatment is the one where, according to our behavioral prediction BHP1, coordination should be easier. The same conclusions do not apply to the Finnish sample, because a lower number of participants choose the winning project in the Info-treatment, compared to the baseline. This result might be explained by the lower level of homogeneity of the Finnish sample.

However, and more interestingly, the only treatment where the same projects were chosen across samples in all market sessions is the Info-treatment, implying that information succeeded in making one project with specific characteristics salient. Despite information either about the project (e.g. institutional funder) or about the designer (e.g. his experience) is non payoff-relevant, our results show that it helped in focusing on the funding of a specific project, relatively to the case where the same information is not provided (Baseline). Moreover, subjects in Italy chose the same projects also in the Combined-treatment.

**Finding 8.** *In the Info-treatment, the two national samples chose the same projects in all sessions. The Italian sample chose the same projects also when uncertainty was added.*

### 3.5.2 Coordination

Our goal is to assess the impact of introducing cheap information and uncertainty on subjects' coordination. We first look at the difference of contributions across treatments. Table 3.3 reports the percentage of participants that invested in the successful project in each treatment. For each treatment, we have 48 independent observations and 144 decisions (subjects made decisions over 3 market sessions).

		Uncertainty	
		No	Yes
Information	No	88,19% (Baseline)	79,86% (Uncertain)
	Yes	83,33% (Info)	75% (Combined)

Table 3.3: Coordination measure: Percentage of subjects who chose the successful project across treatments (N=144 for each treatment)



The percentage of subjects choosing the successful project in the baseline treatment is significantly higher than it is in the Uncertainty-treatment condition and in the Combined treatment condition (two-sided Fisher’s exact  $p$  equal to 0.076 and 0.006 respectively). However, this is not the case when comparing the same percentage to the Information treatment condition (two-sided Fisher’s exact  $p = 0.312$ ). It is worth remembering that two of our theoretical predictions state that introducing either cheap information (THP3) or uncertainty (THP4) (conditioned to a wide prevalence of risk neutral players in the sample) should not affect coordination.

**Finding 9.** *On average, compared to the baseline, introducing uncertainty or a combination of cheap information and uncertainty reduces coordination. The introduction of cheap information alone does not affect coordination.*

It follows that we can reject THP4, while THP3 seemed to be confirmed. On the other hand, it is interesting to notice that the combination of cheap information and uncertainty produces the worst impact on coordination. It seemed that, within the experimental scenario here adopted, the more pieces of information are provided the poorer is coordination. To investigate this result in more detail, we also ran logit regressions (Table 3.4). We use as a dependent variable the probability of choosing the successful project option in each choice. The independent variables are the four decision scenarios: Baseline, Information, Uncertainty, Combined, the gender (variable “Women”) and whether a participant played the game in Italy or in Finland (variable “Italy”). The Baseline scenario, Men, and Finland are our omitted groups.

	Model 1	Model 2
Information	-0.402 (0.342)	-0.458 (0.346)
Uncertainty	-0.633 (0.331)	-0.651 (0.334)
Combined	-0.912*** (0.322)	-0.961*** (0.327)
Women		-0.35 (0.223)
Italy		0.662*** (0.226)
Constant	2.01*** (0.258)	1.91*** (0.297)
Observations	576	573

Table 3.4: Determinants of choosing the successful project (Pooled Italian and Finnish samples, logit model. Standard errors in parentheses. Dependent variable takes value 1 if the subject chooses to invest in the successful project and zero otherwise. \* $p < 0.1$ , \*\* $p \leq 0.05$ , \*\*\* $p \leq 0.01$ )

Looking to Table 3.4 (Model 1) one can immediately notice that the results from the Fisher test statistics are confirmed: the coefficients of treatments Uncertainty and Combined are statistically significant and negatively correlated with our dependent variable. This means that in these treatments the participants coordinated themselves in a less efficient way than they did in the baseline treatment. On the other hand, the coefficient of the treatment Information is not significant, that is a confirmation of the almost insignificant effect generated by cheap information on the overall rate of coordination. For what it regards gender (variable “Women”) and the country (variable “Italy”, Model 2) we can say that there are no gender effects, while there is a strong Nationality effect. Given the impact of the national samples on the coordination measure, we decided to run the non-parametric tests (i.e. Fisher tests) and the regressions separating the Italian and the Finnish groups. Table 3.5 shows the percentages of subjects choosing the successful project in the Italian sample.

		Uncertainty	
		No	Yes
Information	No	91,67% (Baseline)	80,56% (Uncertain)
	Yes	95,83% (Info)	76,38% (Combined)

Table 3.5: Coordination measure: Percentage of subjects who chose the successful project across treatments (Italy, N=72 for each treatment)

Looking at the results we have a confirmation of what we already obtained from the pooled data-base. More precisely, the percentage of Italian subjects choosing the successful project in the baseline treatment is significantly higher than it is in the Uncertainty treatment condition and in the Combined treatment condition (two-sided Fisher’s exact  $p$  equal to 0.09 and 0.021 respectively), while this is not the case when comparing the same percentage to the Information treatment condition (two-sided Fisher’s exact  $p = 0.494$ ). It is also worth noticing that contrary to the result reported from the pooled sample, this time the highest coordination has been reached in the treatment with cheap information and not in the baseline treatment, even if this difference is not significant. Looking to the results from the logit regression (Table 3.6, Italy) we have, like we already had in the pooled sample, a confirmation of the Fisher statistics. Indeed, only the coefficient of the Uncertain-treatment and Combined-treatment are significantly different from zero, and negative, meaning that they are negatively correlated with

the probability of choosing the successful project.

	Model 1 (Italy)	Model 1 (Finland)
Information	0.737 (0.727)	-0.826** (0.418)
Uncertainty	-0.976* (0.520)	-0.378 (0.437)
Combined	-1.224** (0.508)	-0.687* (0.422)
Constant	2.398*** (0.426)	1.713*** (0.328)
Observations	288	288

Table 3.6: Determinants of choosing the successful project (by country, logit model. Standard errors in parentheses. Dependent variable takes value 1 if the subject chooses to invest in the successful project and zero otherwise. \* $p < 0.1$ , \*\* $p \leq 0.05$ , \*\*\* $p \leq 0.01$ )

Considering now the Finnish sample (see Table 3.7), the percentage of subjects choosing the successful project in the baseline treatment is significantly higher than it is in the Information treatment condition (two-sided Fisher's exact  $p = 0.07$ ; Pearson test= 0.045), while is borderline significantly higher in the Combined treatment condition (two-sided Fisher's exact  $p = 0.15$ , one side = 0.075, Pearson test=0.101). However, this is not the case when comparing the same percentage to the Uncertainty treatment condition (two-sided Fisher's exact  $p = 0.516$ ).

		Uncertainty	
		No	Yes
Information	No	84,72% (Baseline)	79,17% (Uncertain)
	Yes	70,83% (Info)	73,61% (Combined)

Table 3.7: Percentage of subjects who chose the successful project across treatments (Finland, N=72 for each treatment)

Running the logit regression (see again Table 3.6, Finland), we have a confirmation of these results, with only the Information treatment significantly (negatively) correlated with the depended variable and a weak (negative) correlation with the Combined treatment.

Recalling our first Behavioral prediction, we expect that cheap information fostered coordination in more homogeneous groups (BHP1) because it should trigger the participants to identify

a common focal piece of information (and therefore the same project). Looking comparatively to the results just shown we should discharge BHP1 because in the more homogeneous sample, which is the Italian one, adding cheap information does not improve coordination in a significant way. This result is a bit delicate because the data suggest that some kind of “roof effect” is taking place in our sample; in fact, in the Baseline treatment the percentage of participants coordinating on the successful project was very high (91,67%) and therefore to obtain an even small increase in this value (in the Information treatment this value raises at 95,83%) could be seen as a sort of confirmation of BHP1. On the other hand, the results from the Finnish sample, the less homogeneous one, open an intriguing issue because it seemed that adding cheap information not only does not foster coordination in inhomogeneous groups, like is predicted by BHP1, but it is even detrimental. In this sense one cannot definitely exclude that cheap information has not any impact on coordination, what reasonably matter is the degree of homogeneity of the players that determines their efficiency in using the information signals as focal points. This is to say it in other words that introducing cheap information seemed not enough to trigger coordination on the basis of the identification of a common focal point in homogeneous groups but, at the same time, it weakens the participants’ efficiency in interpreting the information signals in inhomogeneous groups. Furthermore, it is worth underlining that the participants were aware of their degree of cultural proximity because they knew to be part of a group of persons coming from the same social-cultural environment (the Italian sample) or from different socio-cultural milieus (the Finnish sample, which was made by students coming from different countries). This awareness could have had some effect on the participants’ trust in their own capability to correctly interpret the cheap signals as focal points. Our second behavioral prediction states that the introduction of uncertainty should reduce coordination (BHP2) in presence of high heterogeneity in risk attitudes. This prediction seemed confirmed by the results from the pooled sample but it is falsified looking to the results from the Finnish sample which is the most inhomogeneous one also for what it regards the risk attitudes. Finally, our third behavioral prediction says that combining information and uncertainty should hinder coordination in particular when the sample is inhomogeneous (BHP3). Looking to both the Italian and the Finnish samples this final prediction seemed confirmed, at least broadly, even if, and contrary to the hypothesis, is more established in the Italian (more homogeneous sample) than in the Finnish (less homogeneous) sample. These considerations are summarized as it follows:

**Finding 10.** *In the Italian sample, introducing respectively uncertainty or a combination of uncertainty and cheap information was detrimental in fostering coordination. In the Finnish sample, the introduction of cheap information damages coordination while uncertainty does not affect coordination in a significant way.*

Trying to wrap up the results just discussed one can conclude that all the factors that we introduced in the experimental design (cheap information, uncertainty, group heterogeneity) do play a role in influencing coordination, we will see in the concluding remarks how these results can take to some normative considerations.

### 3.5.3 Choices over time

Now we examine the effects of cheap information, the role of experience, as well as the impact produced by uncertainty, on the speed of coordination. This new performance indicator (the speed of coordination) uses the number of rounds that a project needs to reach the funding threshold. For each market session, Figure 3.1 shows the average percentage of shares cumulated in each period, as well as the average number of rounds needed for the winning project to reach the threshold, considering the whole sample and pooling over treatments. It is possible to observe that, on average, subjects need more rounds to reach the threshold in the third market session. In fact, while in the first market session, the project is funded in five rounds; seven and eight rounds are needed in the second and in the third session respectively. This is a confirmation of previous findings that subjects may learn to play the weakly dominant action, and wait more as they become more experienced.




Figure 3.1: Funding over time: average proportion of shares collected up to each period, by market session (MS)

We now look at the impact of cheap information and uncertainty on the speed of coordination. Focusing on the first market session, Figure 3.2 shows, for each treatment, the average percentage of shares cumulated in each round, as well as the average number of rounds needed for the winning project to reach the threshold.




Figure 3.2: Funding in the first MS: average proportion of shares collected up to each period in the first market session, by treatment.

It is possible to observe that introducing only cheap information increases the speed of coordination. On average, the winning project reaches the threshold within a lower number of rounds, relatively to the baseline. Notice that the Info-treatment is the one where, according to our behavioral prediction BHP1, coordination would be easier. Conversely, introducing market uncertainty decreases the speed of coordination, as, on average, more rounds are needed in order to reach the threshold. Finally, coordination was the slowest in the Combined-treatment. In fact, four and eight rounds are needed, on average, for a winning project to reach the threshold in the baseline and in the Combined-treatment, respectively.

It is important to highlight that the same considerations arise when looking at the two national samples separately (see Figure 3.3). Introducing only cheap information increases the speed of coordination, even if it reduces the average level of coordination, i.e. the total number of subjects choosing the winning project, in the Finnish sample. Moreover, introducing uncertainty decreases the speed of coordination (to the extreme in the case of the Combined-treatment in the Finnish sample, where the threshold was reached only in the last period). As a side note, one can notice that, regardless of the treatment, more shares were collected in Italy.

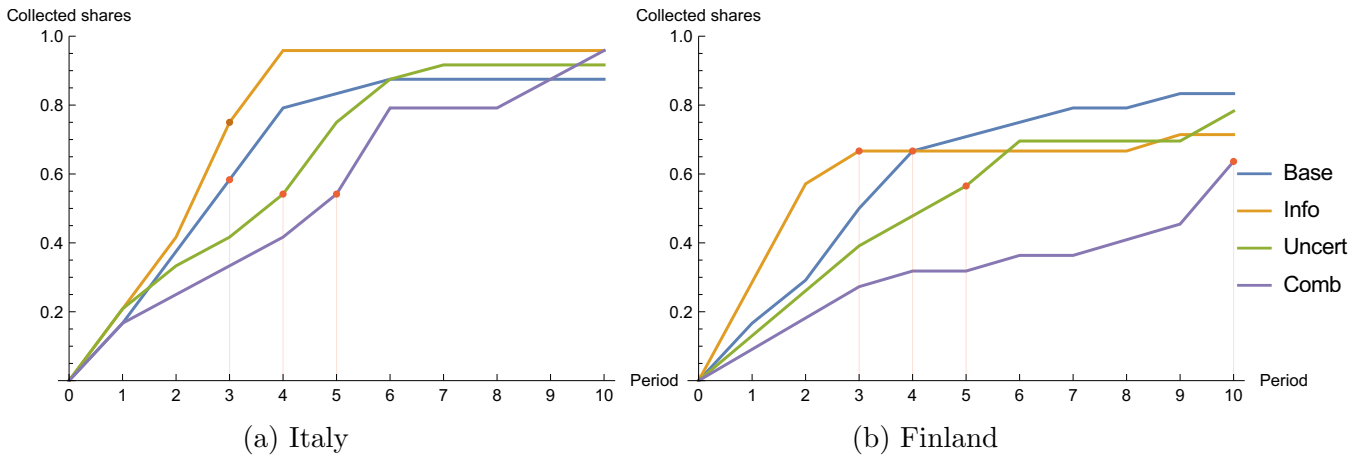


Figure 3.3: Funding in the first MS: average proportion of shares collected up to each period in the first market session, by treatment. Figure 3.3a considers only the Italian sample, Figure 3.3b considers only the Finnish sample.

We summarize these considerations in the following:

**Finding 11.** *In the first market session, the speed of coordination is higher in the presence of information, and it is lower with uncertain payoffs, compared to the baseline.*

*The combined effect slows down the coordination process, especially in the Finnish sample.*

Finally, Figure 3.4 shows, for each treatment, the average percentage of shares cumulated in each round, as well as the average number of rounds needed for the winning project to reach the threshold, considering all the three market sessions. On average, coordination was slower in all treatments, compared to the baseline. Notice that, in the baseline, the only information available to the players is the number of shares collected. It follows that, for a project, it is enough to collect one share more than the others, to become focal. Conversely, in the presence of either cheap information or uncertainty, the participants must, in some way, balance the salience implied by the number of previously collected shares, with the salience induced by the available information (either cheap or payoff-relevant). Thus, a project could be less likely to be considered “focal”, even if it has a higher number of collected shares. Furthermore, in such a case, risk and preference attitudes about the information signal (either cheap or pay relevant) might play a role. This consideration is especially true when the introduction of uncertainty is combined with information about the project and about the designer, as it strongly decreases the speed of coordination.




Figure 3.4: Funding over time: average proportion of shares collected up to each period, considering all market sessions, by treatment.

### First movers

Given the importance of first movers in our coordination game, we now explore more in details their behavior. According to our theoretical model, if participants choose weakly dominant or risk dominant options, we should find no share allocations in the very first period of the game. Table 3.8a shows the (total) number of first movers, i.e. subjects choosing in the very first round for each market session. Our prediction is not confirmed. Nonetheless, there is a clear pattern across sessions in every treatment, as the number of first movers decreases significantly. Indeed, all the differences between the number of first movers at the first round across market sessions are statistically significant accordingly with a Pearson test (see Table 3.8a). It is reasonable to interpret this pattern as the result of some kind of learning process that triggered the participants towards the weakly dominant and the risk dominant strategies.

	MS1	MS2	MS3
Baseline	13	5**	4**
Info	17	3***	5***
Uncertain	13	6*	6*
Combined	13	4**	1***

(a)

	MS1	MS3	TOT
Baseline	0.27	0.09**	0.15
Info	0.38	0.12***	0.19
Uncertain	0.28	0.14*	0.19
Combined	0.28	0.03**	0.14

(b)

Table 3.8: First movers- Table (a) shows the total number of subjects choosing in the first period in each Market session, by Treatment. Table (b) shows % of subjects choosing in the first period when considering only those who chose a project. Stars indicate statistical significance when compared with the first market session (Pearson test, \* $p < 0.1$ , \*\* $p \leq 0.05$ , \*\*\* $p \leq 0.01$ ).

These results are also confirmed when looking at the proportion of early bidders among all



subjects who choose a project, shown in Table 3.8b. Indeed, when considering only subjects who made a decision, the result is even stronger: more subjects chose in the first period in the Info treatment compared to the baseline. Conversely, in the first market session there is no difference when comparing the number of subjects who chose in the first period in the other treatments with those who chose in the baseline.

**Finding 12.** *With learning, there is a shift towards more weakly dominant and risk dominant strategies, i.e. more subjects choose to wait at the beginning of the game. This holds for both samples and in every treatment. Introducing cheap information increases the number of subjects choosing in the first period, while introducing “too much” information (i.e. both cheap information and uncertainty) decreases the number of subjects choosing in the first period.*

Looking now at the number of participants who invested across rounds on the successful project one can notice (see Table 3.9 ) a general pattern similar to the one already commented with regard to the number of first round investors. More precisely one can notice that in the second and third market sessions there is a strong convergence of the choices towards a coordination which takes place from the third to the fifth round. This is to say that the market sessions do matter in determining the winning strategy of the players: the more experienced they are, the more they wait before deciding to invest.

		Rounds considered:						All
		(1)	(1-2)	(3-5)	(6-8)	(9-10)	(10)	
MS1	Baseline	8	16	21	3	1	0	41
	Information	11	22	15	0	1	0	38
	Uncertainty	8	14	17	8	3	3	42
	Combined	6	10	10	8	9	6	37
MS2	Baseline	4	6	26	8	3	1	43
	Information	3	6	17	17	2	1	42
	Uncertainty	5	6	13	15	6	2	40
	Combined	3	3	11	18	8	4	40
MS3	Baseline	2	4	17	17	5	1	43
	Information	3	4	10	17	9	3	40
	Uncertainty	4	9	13	1	12	9	35
	Combined	0	2	3	15	13	10	33

Table 3.9: Contributions over time - Winning projects

To check for the statistical significance of the comment just suggested with regard to the data

shown in Table 3.9, we have run a Logit regression (Table 3.10).

Model 3	
Information	-0.2467 (0.372)
Uncertainty	-0.492(0.365)
Combined	-1.406*** (0.367)
Market session	-0.534*** (0.153)
Period	0.541*** (0.61)
Constant	1.26*** (0.413)
Observations	576

Table 3.10: Determinants of choosing the successful project (Pooled Italian and Finnish samples, logit model. Standard errors in parentheses. Dependent variable takes value 1 if the subject chooses to invest in the successful project and 0 otherwise. \* $p < 0.1$ , \*\*  $p \leq 0.05$ , \*\*\*  $p \leq 0.01$ )

The results of the regression show that the variable market session is significant and negatively correlated with the choice to invest in the winning project. Indeed, average contributions to the winning project were lower in the last market sessions. Moreover, the variable “period” is significant and positively correlated with the choice to invest in the winning project. Thus, the “wait and see” strategy was not only present, but also profitable.

Considering again first movers’ behavior, one could speculate that if information is effective in making one project salient (either payoff-relevant information or cheap information), it should be easier for subjects to coordinate on the winning project from the very first round. Therefore, we should find a higher number of early choices (both in general and on the winning project) in the info-treatment and in the uncertain-treatment, compared to the baseline. In the combined-treatment, instead, we should find a lower number of subjects choosing in the first period (both in general and on the winning project), since it is more difficult both to interpret a project as salient, and to signal a favorite one. Accordingly, by looking at Table 3.9, the average number of subjects choosing the winning project in the first period is higher in the Info-treatment (in the first market session) and in the Uncertain-treatment (in the others market sessions).

### 3.5.4 The role of *cheap* information

The result that, with cheap information, the same project was chosen 3 out of 4 times in all market sessions, do suggest that information played a role on subjects choices. To better clarify this point, we now examine the *clicking* behavior of subjects. Recall that in the info-treatment and in the combined-treatment, subject could click on the piece of information they wanted to see, and we collected data on both the number of clicks and the time they spent on each characteristic. Clearly, we cannot control for attention level, or check whether a subject was indeed looking at the value, but we can say for sure that, if a subject didn't click, then information didn't have any role on his choice.

Overall, we observed the same behavior in both treatments. In the first market session, subjects who clicked on the information did choose the same project, and they were able from the first round to create a signal for subsequent players. Considering all market session, the positive effect of information (i.e. the increase in the likelihood to choose the winning project after having clicked on projects' characteristics) is still present in all cases except for the combined-treatment in Finland (see Figure 3.5).




Figure 3.5: Clicks and choices: click and win (lose) means that a subject clicked at least once on at least one of the characteristics in the first period, and he chose the winning (a losing) project in the first period; no click and win (lose) means that a subject didn't click on any characteristics in the first period, and he chose the winning (a losing) project in the first period.

In Figure 3.5, we grouped first movers, i.e. the firsts participants who choose a project, according to whether they clicked on any characteristics in the period they chose, and to whether they chose a winning or losing project. On average, the proportion of first movers who clicked on some characteristics is higher in Italy than in Finland. We might speculate this is the reason why more shares were collected in Italy. Moreover, the proportion of participants who clicked and then choose a losing project is close to zero in all cases except for the Combined-treatment in Finland (Purple area in Figure 3.5, respectively). Finally, the proportion of subjects choosing

the winning project is higher among those who clicked, than among those who didn't clicked (Yellow and Orange in Figure 3.5, respectively).

Pooling the two national samples, and considering both treatments, we notice that, on average, 88% of subjects who clicked on the characteristics choose the winning project, while only the 12% of subjects who didn't click chose the winning project. Thus, we might speculate that subjects who clicked on some information were indeed more likely to choose the winning project, suggesting an important role for information in helping the coordination process.

Notice that, alternatively, one could interpret the clicking behavior as a proxy of the degree of attention that subjects were putting in the task. In this case, a lower number of clicks would imply less attention, and subjects who put more attention (i.e. clicked more) also managed to have better results, i.e. they were able to reach a higher level of coordination.

It is natural to ask whether the projects that were consistently chosen among the different treatments shared some characteristics that might attract participants' attention. From Table 3.B.1 in the Appendix 3.B we can see that the winning projects in the market sessions with information did share some basic characteristics: they were sponsored by the Bank, and they were designed by an Experienced Male studying Economics.

Given the abundance of information available to subjects, it is difficult to say if those characteristics were indeed the main driver of subjects' behavior. We can get further insights by considering the type of characteristics that was more clicked, and by analyzing participants' answers to the debriefing questionnaire.

Regarding the first point, we found that the most clicked characteristics were the sponsor and the rating of the project, while the least clicked was the gender of the designer. Thus, we might speculate that the sponsor could be one of the important drivers of subjects' behavior, as it was both one of the most clicked characteristics, and winning projects all share the same sponsor (i.e. the Bank). Moreover, even if the designer of the winning project was always a male, very few subjects did click on this characteristics. It is thus unluckily that the gender had any influence on subjects' decisions.

Moving to the second point, Table 3.11 shows the proportion of subjects who considered each characteristic as *important*, according to their answers to the questionnaire, for each treatment. We can see that in the info-treatment the variable "Information" was considered important by almost half of subjects, while in combined-treatment the important variable for most of the

subjects was the level of risk of the project. Notice that Table 3.11 considers all subjects, not just those choosing as first movers. The evidence from Table 3.11 points to the direction that subjects did indeed considered the information provided as a tool to decide.

	<b>Baseline</b>		<b>Information</b>		<b>Uncertainty</b>		<b>Combined</b>	
	ITA	FIN	ITA	FIN	ITA	FIN	ITA	FIN
Collected shares	0.92	0.78	0.50	0.33	0.54	0.30	0.21	0.18
Color	0.04	0.09		0.04		0.04		
Info			0.46	0.38			0.17	0.30
Risk				0.08	0.42	0.58	0.58	0.48
Other	0.04	0.13	0.04	0.17	0.04	0.08	0.04	0.04

Table 3.11: Debriefing question: Decisions. % answers to the question: Which variable was most important to take your decision?

Finally, in Table 3.12 we split those subjects who considered the information as important, according to which characteristic, among those available, was more important for their decision. Overall, we can say that Education of the designer was not the driver of subjects decisions, as only the 10% of subjects listed it as important. Notice that rows do not add to 1, as subjects could list more than one characteristic.

	Rating	Sponsor	Education	Experience
Info (20)	0.3	0.45	0.1	0.5
Combined (11)	0.45	0.36	0.09	0.55

Table 3.12: Debriefing question: Information. % of subjects who listed each characteristics as important to make the decision

Thus, recalling that the projects that were chosen more often in the Info-treatment and in the Combined-treatment were all sponsored by the Bank, and they were designed by an Experience Male studying Economics, we can speculate and say that the main driver of subjects' decision were the Sponsor and the Experience of the Designer. Indeed, we can disregard the Gender of the Designer, since it was almost never clicked, and his Education, from what we noticed in Table 3.12. Indeed, the Sponsor of the project was the most clicked characteristics, and, together with the Experience of the Designer, was among the most relevant variables to decide, as stated by subjects.

### 3.5.5 Strategies and choices

We now check whether choices were rational according to two very board criteria. First notice that, in all treatments, if a project already reached the threshold, choosing a different project can be regarded as irrational, since keeping the share would give a strictly higher payoff. When considering all the investing choices that could have been made by the subjects, we have a total of 24 (participants) x 3 (market sessions) x 4 (treatments) x 2 (nations), for a total of 576 choices. Over all, only 5 choices, i.e. less than 1%, can be classified as irrational according to this criterion. Second, notice that, in both treatments with certain payoff, if a project already reached the threshold, keeping the share and not investing can be regarded as irrational, since it is strictly dominated by investing in the “winning” project. Over all, 16 choices (3 in the baseline and 13 in the Info-treatment) can be classified as irrational according to this second criteria, which is around the 5% of choices in those treatments.

In one of the debriefing questions, we asked participants whether they would have preferred to use a different strategy. The first row of Table 3.13 shows the number of subjects who answered affirmatively. The remaining rows show the proportion of subjects who stated that they would have preferred to choose one of the following strategies: to decide later (Wait), to keep the share and not invest (Keep) or to decide sooner (Signal).

	<b>Baseline</b>		<b>Information</b>		<b>Uncertainty</b>		<b>Combined</b>	
	ITA	FIN	ITA	FIN	ITA	FIN	ITA	FIN
Change?	4	7	0	9	12	10	7	8
Wait	100%	72%	-	22%	17%	10%	43%	62.5%
Keep	0	0	-	0	25%	50%	0	0
Signal	0	14%	-	33%	33%	30%	43%	12.5%
Other	0	14%	-	45%	25%	10%	14%	25%

Table 3.13: Preference for different strategies

It seems that subjects understood in which situation it was more convenient to wait rather than to early bid in order to signal a project. As Table 3.13 shows, the majority of subjects in the baseline and in the combined-treatment reported that they would have preferred to wait. On the contrary, the majority of subjects in the uncertain-treatment reported that they would have preferred to invest before or to keep the share.

## 3.6 General discussion and Conclusions

This research belongs under the umbrella of behavioral economics and focuses on creating a better understanding of how individuals, and crowds, make investment decisions under uncertainty, within the context of crowdfunding, and what drives their decision-making. More precisely, we investigate the effect that information about the project and the other funders (the existence of an institutional funder) and market uncertainty have on the success of projects in raising funds. Furthermore, we study the speed at which a project reaches the funding threshold and the role of information in the coordination of investor behavior. The research set-up used is experimental and included four treatments, which were all done twice, once in Italy and once in Finland, to reach proper validation of the results. Consider first our theoretical predictions. Recall that the first one implies full coordination in all pure equilibria, and the second one implies that, if subjects choose weakly dominant strategies, they should wait until enough subjects choose a project before taking a decision.

Regarding the first prediction, we can reject that subjects behave according to a pure Nash equilibrium in all cases except for the Italian sample in the Info-treatment. Regarding the second prediction, our results show a clear tendency of subjects to chose the weakly dominant strategy, that is to wait. Indeed, this is the most chosen strategy, and the portion of subjects adopting this strategy is increasing over time, in line with previous experimental results.

Our last theoretical predictions imply that we have no reason to expect that one particular project is chosen across the different market sessions, but our results do suggest that both information and uncertainty played a role in determining subjects' behavior.

Moving to the behavioral predictions, recall that the first and second ones imply that, in homogeneous samples, introducing information, either in the form of not payoff-relevant project's characteristics, or in the form of uncertainty over project's returns, might help coordination.

We found that, when only *cheap* information is present, coordination is faster and the same project is chosen across sessions. Moreover, in the more homogeneous sample (i.e. the Italian one), coordination was higher, in the sense that more subjects choose the same project.

When only uncertainty is present, coordination is slower and lower, compared to the baseline. This result might confirm the intuition that crowd's risk preferences are relevant for the success of a crowdfunding campaign.

Our last prediction implies that the presence of both information and uncertainty might hinder coordination, as it increases the noise, making more difficult to interpret one option as salient. Indeed, in the Combined-treatment coordination is slower and lower, in the sense that fewer subjects choose the same option. Moreover, in this treatment there is the highest number of subjects choosing the waiting strategy, and the highest number of subjects choosing the risk dominant strategy, that is to keep the share.

Overall, our results provide strong evidence for the presence of the herding effect, and, perhaps more importantly, subjects were aware of it. Indeed, first movers did try to signal to the crowd their favorite project, even in the case where the information provided was no payoff-relevant, and there was no reason to prefer one project over the other. We found evidence that the project which managed to collect the highest number of shares in the very first round was the one reaching the threshold in all but one session. Interestingly, introducing cheap information, was more effective in coordinating first movers actions, compared to both the baseline and the treatment with uncertain payoffs. This might suggest that, when designing a crowdfunding campaign, more emphasis should be given to the characteristics of the projects in terms of its quality, rather than on its capability to generate a monetary return. Overall, our results suggest that the crowd is mainly driven by others behavior. Since we found evidence that the behavior of the first movers is crucial to determine which project would succeed, understanding their motives can be highly important to design a successful crowdfunding campaign. In a research in progress, Bretschneider et al. (2014) suggest several possible factors that may influence backers: intrinsic motivation ( “fun to make investments”, curiosity, altruism, reciprocity), projects characteristics (personal preferences for the object of the project, return), salience (“direct and indirect identification”, recognition), and others’ behavior. In our setting, giving non payoff-relevant information might refer to project salience, while introducing uncertainty refers to project’s returns. Our results suggest that first movers are indeed driven by both project’s characteristics and salience. Nonetheless, when comparing the first movers coordination rates between the treatments with information and with uncertainty, we found that coordination is in a sense easier with homogeneous samples and cheap information. Thus, we might speculate that, to increase the chances of begin chosen by first movers, designers should push on project’s salience, rather than to focus on projects returns or its relative risk.



# Appendix

## 3.A Notes on Coordination

**Two players** Consider the case of 2 players, who have two periods to choose between a safe option (N), giving a payoff of 0, and two “risky” options A and B. In  $t_1$  the action set is  $A_i(t_1) = \{A, B, W\}$ , where  $W$  is the option to wait for the next period. In  $t_2$ , if  $a_i(t = 1) = W$ , the action set is  $A_i(t_2) = \{A, B, N\}$ , otherwise the action set is empty. The payoff of player  $i$  from choosing a risky option  $a_i(t) = j \in \{A, B\}$  in period  $t$  is:

$$\pi_i(a_i(t) = j) = \begin{cases} R_{ijt} > 0 & \text{if } a_1 = a_2 \\ -c < 0 & \text{if } a_1 \neq a_2 \end{cases} \quad (3.1)$$

As we saw in Section 3.4, to wait in  $t_1$  is weakly dominant, and to wait and choose the safe option if the opponent didn't choose in  $t_1$  is risk dominant.

If the returns from the risky options are the same, that is  $R_{ijt} = R \quad \forall i, j, t$ , there are 3 sets of (pure) Nash equilibria: one set where both players choose today, one set where they both choose tomorrow, and one set where one player chooses today and the other chooses tomorrow. We can try to select among those equilibria by introducing some form of asymmetry. Consider first the case where the returns from the risky options are different for the two players: for one player we have  $R_A > R_B$ , and for the other player  $R_B > R_A$ . We can then state the following:

**Remark 10.** *The set of equilibria where both players choose tomorrow is not subgame perfect.*

Previously, we assumed that players' cost to make the decision is the same in  $t_1$  and  $t_2$ . Clearly, with equal returns from projects, if choosing today has a higher cost than choosing tomorrow, then in all subgame perfect equilibria players chooses in  $t_2$ , while if choosing today is cheaper than choosing tomorrow, in all subgame perfect equilibria players chooses in  $t_1$ . Nonetheless, if players have different costs to make decisions, and different returns from the risky options,

the only subgame perfect equilibrium gives to the player who has the smaller cost of investing in  $t_1$  is preferred option:

**Remark 11.** *The only subgame perfect equilibrium gives to the player that has the smaller cost of investing in  $t_1$  his preferred option.*

**N players** Consider the case where  $N$  players has to choose between a safe option and a risky one, giving a positive return only if  $k \leq N$  players chose it. Players have  $T$  periods to decide. In each period, they can choose between waiting or investing. Once they chose investing, they don't make any other choice. In this simple setting, we can state the following:

**Remark 12.** *If  $T \geq k$ , choosing the safe option is not a s.p.e..*

We now consider the simultaneous case where  $N = 10$  players can either coordinate on one of  $J = 2$  risky options, A and B, or choose the safe option, S. The safe option has a return of  $y > 0$ , while each risky option has a return of  $r > y$  if at least  $k = 6$  players choose it, and 0 otherwise. As we saw in Section 3.4, the game has  $J + 1 = 3$  pure NE, where everybody choose the same option. The game has also  $J = 2$  mixed equilibria where players mix between the safe option and one of the risky options. Whether the game has also other mixed equilibria crucially depends on the returns of the risky options.

Consider first an equilibrium where the safe option has probability zero. Since players are mixing between the risky options, and since the risky options all have the same return, in equilibrium players will choose each option with the same probability,  $p_j$ . Thus, if in an equilibrium players mix among  $\hat{J}$  options, they will choose each option with probability  $p_j = 1/\hat{J}$ . In our example, if players mix between the two risky options A and B, they will set  $p_A = p_B = 1/2$ . Then, the expected payoff of choosing option  $j$  is  $r/2$ , meaning that this is an equilibrium if and only if  $r \geq 2y$ . Notice that, if in equilibrium more than 2 risky options have positive probability, the resulting distribution of choices will follow a multinomial distribution. If  $q_j(1/\hat{J})$  is the probability that project  $j$  reaches the threshold, given that each player chooses option  $j$  with probability  $1/\hat{J}$ , player  $i$  will mix if and only if:

$$r q_j(1/\hat{J}) \geq y$$

. Thus, there can be a mixed equilibrium where players mix among  $\hat{J}$  risky options only if  $r \geq \frac{y}{q_j(1/\hat{J})}$ .

### 3.B Projects

Projects	Uncertain Ret.		Cert. Ret.	Project's charact.		Designer's characteristics			Market Session
	Profit	Success	Profit	Sponsor	Rating	Experience	Education	Sex	
Yellow <sub>1</sub>	572	0.7	400	Bank	4	15	Economics	M	MS1
Red <sub>1</sub>	1000	0.4	400	Insurance	4	8	Sociology	F	
Blue <sub>1</sub>	1000	0.4	400	University	2	5	Law	F	
Yellow <sub>2</sub>	1000	0.5	500	University	9	6	Law	F	MS2
Red <sub>2</sub>	1250	0.4	500	Bank	2	20	Economics	M	
Blue <sub>2</sub>	834	0.6	500	Insurance	5	2	Engineering	M	
Yellow <sub>3</sub>	1000	0.5	500	University	9	6	Law	F	MS3
Red <sub>3</sub>	1000	0.5	500	Curia	6	8	Law	F	
Blue <sub>3</sub>	834	0.6	500	University	3	3	Engineering	M	
Green <sub>3</sub>	1250	0.4	500	Bank	2	20	Economics	M	
Purple <sub>3</sub>	834	0.6	500	Insurance	5	2	Engineering	M	
Orange <sub>3</sub>	715	0.7	500	District	9	5	Economics	M	

Table 3.B.1: Randomly selected projects of each session (in the third market session we used the same projects as in the second plus 3 randomly selected projects). Projects that were chosen in the Info-treatments are highlighted.

### 3.C Instructions

Welcome! We thank you for accepting to participate in this experiment, which will allow you to earn some money. You have already earned 2.5 Euro to arrive on time. During the experiment you will not be able to communicate with the other participants. If something in the instructions is not clear to you, please raise your hand and ask for information directly from the people conducting the experiment. The payment you will receive at the end of the experiment will be determined by the choices you make and by the choices that will be made by other participants, according to the methods that will be explained below. During the

experiment, the amounts are expressed in EMU (Experimental Monetary Units), where 50 EMU = 1 Euro, or 1 EMU = 0.02 Euro. This experiment is divided into two parts. The instructions for the first part of the experiment are given below. You will be given instructions on the second part later. During the experiment your choices will remain anonymous.

### **First part - Instructions**

The aim of this experiment is to simulate a project market. In a previous experiment, we asked other participants, called designers, to design a set of projects. Each designer had the opportunity to choose the characteristics of a project, in order to make it attractive to the largest possible number of lenders. You and the other participants of today can finance one of these projects, which will be indicated by different colors (blue, pink, white, ect). In order to be financed, a project must reach a minimum funding threshold, which is the same for all the projects. If a project exceeds the minimum funding threshold, it might generate profits for all the lenders who have invested in that project. The profits of a financed project will depend on the result of a lottery, according to the procedures that will be described later. Each participant will receive one share at the beginning of the experiment, valued at 150 EMU, and can choose to hold it, or use it to invest in only one project. The experiment consists of 10 rounds. In each round after the first, you can see the level of funding that each project has gotten until that round, and if you have not already invested your share, you can decide whether to invest it in a project or to wait, but you cannot get back a share you have already invested. Since the minimum funding threshold for a project is 13 shares, at least 13 lenders will be needed to reach the threshold for the realization of the project, and only one project can be financed. If a project fails to reach the minimum funding threshold, any lenders of that project will be partially reimbursed of the invested portion, according to the procedures that will be described later.

### **The projects**

The Projects are characterized by the following elements:

- Value of the project
- Probability of success

- Type of institution that sponsors the project
- Reputation of the project
- The designers are instead characterized by the following characteristics:
- Education
- Experience

These characteristics will be summarized in a table like the following:

Project	Value	Prob.succ.	Sponsor	Rating	Sex	Education	Experience
Yellow	1000	0.5	University	8	F	Law	6

Figure 3.C.1: Example of a project

To view each characteristic, you will have to click on the corresponding name, and the respective values for each project will appear. You will be able to view one characteristic at a time, but you can click on each name as often as you wish. Moreover, by passing the mouse over the name of each characteristic, a window will be opened with a brief explanation of the variable. Below you will find the explanation of each characteristic of the projects:

**Value of the project** As already mentioned, if a project reaches the funding threshold, your profit will depend on the result of a lottery. A lottery is made of a given probability of success and a corresponding profit. The first characteristic that you will find in the table is the value of the project, that indicates the profit generated by the project in the case of success, that is, if the project reaches the minimum funding threshold and the lottery is successful.

**Probability of success** The second characteristic is the risk of the project, that is the probability with which the funded project is able to generate profits. This variable is related to the value of the project, since the higher is the value of the project, the lower is the probability to obtain the corresponding profit. For example, to a probability of 10% corresponds a profit of 10000 EMU, while to a probability of 90% corresponds a profit of 1111 EMU.

**Type of institution that sponsors the project** The third characteristic of the project is the institutional sponsor. The designers of the projects have in fact chosen the sponsor from a list of various proposed institutions that they have thought is the sponsor that will make the

project most attractive for lenders.

**Reputation of the project** The fourth characteristic of the project indicates the reputation of the project among the designers. After having defined the various projects, the designers were asked to vote for the projects based on their perception of which project will be the most likely to reach the funding threshold. The third characteristic indicates how many votes each project received during this evaluation.

Below we provide the explanation of each characteristic of the designer:

**Designer education** This characteristic indicates the education of the designer of the project, that is, in which degree programme is the designer enrolled in.

**Designer experience** This characteristic indicates the experience of the designer accumulated by having participated in experiments in the past. The experience is indicated by the number of experiments in which the designer took part until now.

**Designer Gender**

**Your choice**

In each round, you can decide whether to invest your share in a project or to wait, by selecting the appropriate option: Once your share has been invested, it is not possible to get it back or to change project in the following rounds. In each round you will also have the following pieces of information:

**Collected shares**

It indicates the total shares that the project collected until that round. Both the total number of shares collected and the percentage of funding will be indicated. Furthermore, bar charts are used to make the comparison between the funding levels of the different projects more intuitively understandable. The bar graph will be red if the corresponding project has not yet reached the minimum funding threshold, and, once the threshold is reached, it will turn green. It is possible to participate in the financing of a project, even when this has already reached the minimum funding threshold. When making your choice, before clicking on a characteristic, you will visualize the following window:

## Project Financing

This is round number 1 of 10. You can still invest your share.  
Still 24 participants have to decide.

Project	Profit		Collected shares	Forecast
Yellow	400	0		0% <input type="text"/>
Red	400	0		0% <input type="text"/>
Blue	400	0		0% <input type="text"/>

Make your choice :

Yellow  Red  Blue  Not Invest in this round

Figure 3.C.2: Choice Window: Baseline

## Project Financing

This is round number 1 of 10. You can still invest your share.  
Still 24 participants have to decide.

Project	Profit	Sponsor	Rating	Sex	Education	Experience	Collected shares	Forecast
Yellow	400						0	0% <input type="text"/>
Red	400						0	0% <input type="text"/>
Blue	400						0	0% <input type="text"/>

Make your choice :

Yellow  Red  Blue  Not Invest in this round

Figure 3.C.3: Choice Window: Info-treatment

## Project Financing

This is round number 1 of 10. You can still invest your share.  
Still 24 participants have to decide.

Project	Value	Prob.succ.	Collected shares	Forecast
Yellow	1429	0,7	0	0%
Red	2500	0,4	0	0%
Blue	2500	0,4	0	0%

Make your choice :

Yellow
  Red
  Blue
  Not Invest in this round

Figure 3.C.4: Choice Window: Uncertain-treatment

## Project Financing

This is round number 1 of 10. You can still invest your share.  
Still 24 participants have to decide.

Project	Value	Prob.succ.	Sponsor	Rating	Sex	Education	Experience	Collected shares	Forecast
Yellow	1429	0,7						0	0%
Red	2500	0,4						0	0%
Blue	2500	0,4						0	0%

Make your choice :

Yellow
  Red
  Blue
  Not Invest in this round

Figure 3.C.5: Choice Window: Combined-treatment



In each round you will also be informed of how many participants have not yet chosen to invest. In the last column of the table, called forecasts, you will have to indicate how many participants, among those who have not yet chosen to invest, are choosing each project during the current round. .

### **Experimental Procedure:**

As anticipated, the experiment consists of 10 rounds. In each round you can indicate whether to invest your share in a project, or to wait. You can invest in a project even if this has already reached the funding threshold. You can decide not to invest in any project and to keep your share. In each round following the first, you will have information on the level of financing of the projects. Remember that only one project can be financed, because the sum of shares available to the lenders, is such that only one project at most can be financed. It is important to note that once you have invested your share in a project, it will no longer be possible to change your decision, i.e. you will not be able to remove your share and transfer it to another project, nor will you be able to get back your share. Remember that, in each round of the experiment, you will also need to indicate the number of subjects you believe will invest in each project during that round.

### **Profits calculation**

After the last round, profits will be calculated. If a project has reached the minimum funding threshold, it can generate profits, according to the result of the corresponding lottery. In particular:

- If you have invested in a project, that reaches the funding threshold, and won the lottery, you will receive the value of the project. For example, if the project value is 450 EMU, you will receive 450 EMU.
- If you have invested in a project, that reaches the funding threshold, and didn't win the lottery, you will lose your share and your profit will be 0 EMU.
- If you have invested in a project, and the project does not reach the funding threshold, you will receive 50 EMU, as a partial refund of the amount invested.
- If you have not invested in any project, you will keep your initial share, i.e. 150 EMU.

## Sessions

During the experiment you will participate in three sessions, i.e. three markets of projects. In every session, the above-explained procedure is repeated. This means that each session will last 10 rounds, and in each session you will receive a share that you can use to finance a project from those available in the session. Within each session, projects will have the same value and the same funding threshold. The value of projects can change between one session and another.

## Payments

Each participant will receive 2.5 Euro for participating in the experiment. For this part of the experiment, the gain will be calculated as follows: at the end of the last session, one of the three sessions (markets) will be randomly drawn, and your earnings will be equal to your profit in the selected session. Within the selected session, one round will be randomly drawn and the number of participants who have invested in each project during that round will be calculated. The participant who indicated the number of subjects that is closest to the true number, will receive 100 EMU. If there are any equal-merit, the winner will be randomly picked among them.

## Control questions

1. There are 10 rounds to finance a project within each session. True or false?
2. In the randomly selected session the White project was financed, worth 450 EMU, and it didn't win the lottery. You have invested in the White project. Your profit is of 0 EMU. True or false?
3. Suppose that you invest your share on a project in the third round. At the seventh round you decide to remove your share from that project, and invest it on another project. Is it possible?
4. Your profit is given by the sum of the profits obtained in each session. True or false?
5. In the randomly selected session the White project was financed, worth 450 EMU, and it won the lottery. You have invested in the White project. What is your profit?
6. In the randomly selected session the White project was financed, worth 450 EMU. You have invested in the Black project. What is your profit?

7. In the randomly selected session the White project was financed, worth 450 EMU. You have decided to keep your share and to not invest it. What is your profit?
8. At each round you will have to indicate how many players have already invested in each project. True or false?

### **Second part - Instructions**

On your screen you will see 100 boxes. In one of them, there is a bomb; each of the other 99 boxes contains 2 EMU. You do not know where the bomb is, but you know that it might be in any of the 100 boxes with the same probability. Your task is to select all the boxes you want. You will earn 2 EMU for each box that you collect without the bomb. To select a box, you can simply click on it. Selection does not imply that it will immediately open; instead, you will discover the actual contents (EMU or the bomb) of your boxes only at the end of the experiment. If you select the box containing the bomb, then everything you have collected is destroyed and you will earn 0 EMU for this second phase of the experiment. After collecting all the boxes you like, select the STOP button. This ends the second phase of the experiment. The earnings (in EMU) you acquire in this second phase of the experiment will be added to the earnings (in EMU) you obtained during the first phase of the experiment.



# Conclusions

In this thesis we tried to explore the relation between information and equilibrium selection. In Chapter One we developed a theoretical model which predicts a higher level of cooperation when less information is available. We then applied the model to an asymmetric situation where one long-run player interacts with a sequence of short-run players. We call the information available to the short-run player about the long-run player previous behavior as the *reputation* of the long-run player. We show that giving more information to the short-run player is equivalent to expand the set of possible reputations, and consequently of equilibria, so that also extortionate outcomes become equilibria.

In Chapter Two we tested the model developed in Chapter One with a laboratory experiment. We show that subjects in the experiment are indeed able to recognize the different settings triggered by different amount of information. They tend to use more a fully cooperative strategy when little information is available, and to use more an extortionate strategy when the theory prescribes that is optimal to do so, i.e. when more information is available. At the same time, the data reveal that first players tend to resist to extortion. They are not willing to tolerate low levels of cooperation by the second player, even when this would be profitable as long as monetary payoffs are concerned.

There is a clear analogy between our result and the large literature on ultimatum bargaining (see D. Cooper and Kagel, 2016 for a review). It is one of the best established results in the experimental literature that in the ultimatum game offers are made that are above the minimum predicted by the subgame perfect equilibrium, and that "unfair" offers are usually rejected. A wealth of "social preferences" models have been proposed to accommodate this pattern of behavior. The repeated TG resembles ultimatum bargaining because, by choosing the frequency of rewards the second player chooses the "fairness" of the final distribution of (expected) payoffs. The first player can either accept it (by playing Trust) or reject (by playing

Not Trust). This implies that part of the observed behavior can be explained, at least partially, by subjects' preferences involving social motives like altruism or inequality aversion. How much of the behavior in experimental setting involving repeated games can be explained by models of social preferences is a relatively unexplored question, that surely deserves more scrutiny.

Chapter Three reports the results from an experiment designed to test whether non payoff-relevant information is able to help subjects solve a coordination problem. Our main finding is that information is able to make an option focal, even if it does not increase coordination. In the presence of uncertain payoffs, introducing information is instead detrimental for coordination. Those results, when applied to the context of crowdfunding, suggest that project's designers should push more on the inner characteristics of the project, and less on its ability to generate profits. Those results might be corroborated by repeating the experiment with different types of information, to get further insights into the relevance of the disclosed information.

# Bibliography

- Adami, C. and Hintze, A. (2013). “Evolutionary instability of zero-determinant strategies demonstrates that winning is not everything”. In: *Nature Communications* 4, pp. 1–7.
- Agrawal, A., Catalini, C., and Goldfarb, A. (2011). *The geography of crowdfunding*. Tech. rep. National bureau of economic research.
- (2014). “Some simple economics of crowdfunding”. In: *Innovation Policy and the Economy* 14.1, pp. 63–97.
- Ahlers, G. K. et al. (2015). “Signaling in equity crowdfunding”. In: *Entrepreneurship Theory and Practice* 39.4, pp. 955–980.
- Aumann, R. J. (1981). “Survey of repeated games”. In: *Essays in game theory and mathematical economics in honor of Oskar Morgenstern*.
- Axelrod, R. and Hamilton, W. D. (1981). “The evolution of cooperation.” In: *Science* 211.4489, pp. 1390–1396.
- Axelrod, R. (1980a). “Effective Choice in the Prisoner’s Dilemma”. In: *The Journal of Conflict Resolution* 24.1, pp. 3–25.
- (1980b). “More Effective Choice in the Prisoner’s Dilemma”. In: *The Journal of Conflict Resolution* 24.1, pp. 379–403.
- (1984). *The Evolution of Cooperation*. Basic Books.
- Baek, S. K. et al. (2016). “Comparing reactive and memory-one strategies of direct reciprocity”. In: *Scientific Reports* 6.25676, pp. 1–20.
- Bagnoli, M. and Lipman, B. L. (1989). “Provision of public goods: Fully implementing the core through private contributions”. In: *The Review of Economic Studies* 56.4, pp. 583–601.
- Baklanov, A. (2018). “Nash equilibria in reactive strategies”. working paper.
- Barlo, M., Carmona, G., and Sabourian, H. (2009). “Repeated games with one-memory”. In: *Journal of Economic Theory* 144.1, pp. 312–336.

- Baum, J. A. and Silverman, B. S. (2004). “Picking winners or building them? Alliance, intellectual, and human capital as selection criteria in venture financing and performance of biotechnology startups”. In: *Journal of Business Venturing* 19.3, pp. 411–436.
- Belleflamme, P., Lambert, T., and Schwienbacher, A. (2010). “Crowdfunding: An industrial organization perspective”. In: *Prepared for the workshop Digital Business Models: Understanding Strategies’, held in Paris on June*, pp. 25–26.
- Bendor, J. and Swistak, P. (1995). “Types of evolutionary stability and the problem of cooperation.” In: *Proceedings of the National Academy of Sciences* 92.8, pp. 3596–3600.
- Bergstrom, C. T. and Lachmann, M. (2003). “The Red King effect: when the slowest runner wins the coevolutionary race”. In: *Proceedings of the National Academy of Sciences* 100.2, pp. 593–598.
- Bigoni, M., Camera, G., and Casari, M. (2014). *Money is More than a Memory*. Working Papers 14-17. Chapman University, Economic Science Institute.
- Blume, A. and Ortmann, A. (2007). “The effects of costless pre-play communication: Experimental evidence from games with Pareto-ranked equilibria”. In: *Journal of Economic Theory* 132.1, pp. 274–290.
- Boerlijst, M. C., Nowak, M., and Sigmund, K. (1997). “Equal pay for all prisoners”. In: *The American Mathematical Monthly* 104.4, pp. 303–305.
- Bohnet, I. et al. (2005). “Learning trust”. In: *Journal of the European Economic Association*, pp. 322–329.
- Bolton, G., Katok, E., and Ockenfels, A. (2005). “Cooperation among strangers with limited information about reputation”. In: *Journal of Public Economics* 89.8, pp. 1457–1468.
- Bolton, G., Ockenfels, A., and Ebeling, F. (2011). “Information value and externalities in reputation building”. In: *International Journal of Industrial Organization* 29.1, pp. 23–33.
- Bornhorst, F. et al. (2004). *How do People Play a Repeated Trust Game? Experimental Evidence*. Sonderforschungsbereich 504 Publications 04-43. Sonderforschungsbereich 504, Universiteit Mannheim.
- Bornstein, G., Gneezy, U., and Nagel, R. (2002). “The effect of intergroup competition on group coordination: An experimental study”. In: *Games and Economic Behavior* 41.1, pp. 1–25.
- Bracht, J. and Feltovich, N. (2009). “Whatever you say, your reputation precedes you: Observation and cheap talk in the trust game”. In: *Journal of Public Economics* 93.9, pp. 1036–1044.



- Bretschneider, U., Knaub, K., and Wieck, E. (2014). “Motivations for crowdfunding: What drives the crowd to invest in start-ups?” Working paper.
- Brüntje, D. and Gajda, O. (2015). *Crowdfunding in Europe: State of the art in theory and practice*. Springer.
- Burch, G. (2011). “Herding Behavior as a Network Externality.” In: *International Conference on Information Systems 2011*. Vol. 2.
- Burch, G., Ghose, A., and Wattal, S. (2013). “An empirical examination of the antecedents and consequences of contribution patterns in crowd-funded markets”. In: *Information Systems Research* 24.3, pp. 499–519.
- (2015). “The hidden cost of accommodating crowdfunder privacy preferences: a randomized field experiment.” In: *Management Science* 5.61, pp. 949–962.
- Cachon, G. P. and Camerer, C. F. (1996). “Loss-avoidance and forward induction in experimental coordination games”. In: *The Quarterly Journal of Economics* 111.1, pp. 165–194.
- Cadsby, C. B. and Maynes, E. (1999). “Voluntary provision of threshold public goods with continuous contributions: experimental evidence”. In: *Journal of Public Economics* 71.1, pp. 53–73.
- Cartwright, E. and Stepanova, A. (2015). “The consequences of a refund in threshold public good games”. In: *Economics Letters* 134, pp. 29–33.
- Cason, T. N. and Zubrickas, R. (2017). “Enhancing fundraising with refund bonuses”. In: *Games and Economic Behavior* 101, pp. 218–233.
- Charness, G., Du, N., and Yang, C.-L. (2011). “Trust and trustworthiness reputations in an investment game”. In: *Games and Economic Behavior* 72.2, pp. 361–375.
- Chatterjee, K., Zufferey, D., and Nowak, M. (2012). “Evolutionary game dynamics in populations with different learners”. In: *Journal of Theoretical Biology* 301, pp. 161–173.
- Chen, D. L., Schonger, M., and Wickens, C. (2016). “oTree - An open-source platform for laboratory, online, and field experiments”. In: *Journal of Behavioral and Experimental Finance* 9, pp. 88–97.
- Chen, J. and Zinger, A. (2014). “The robustness of zero-determinant strategies in iterated prisoner’s dilemma games”. In: *Journal of Theoretical Biology* 357, pp. 46–54.
- Cholakova, M. and Clarysse, B. (2015). “Does the possibility to make equity investments in crowdfunding projects crowd out reward-based investments?” In: *Entrepreneurship Theory and Practice* 39.1, pp. 145–172.

- Coats, J. C., Gronberg, T. J., and Grosskopf, B. (2009). “Simultaneous versus sequential public good provision and the role of refund: an experimental study”. In: *Journal of Public Economics* 93.1-2, pp. 326–335.
- Colombo, M. G., Franzoni, C., and Rossi-Lamastra, C. (2015). “Internal social capital and the attraction of early contributions in crowdfunding”. In: *Entrepreneurship Theory and Practice* 39.1, pp. 75–100.
- Compte, O. and Postlewaite, A. (2015). “Plausible cooperation”. In: *Games and Economic Behavior* 91, pp. 45–59.
- Cooper, D. and Kagel, J. (2016). “Other regarding preferences: A selective survey of experimental results”. In: 2.
- Cooper, R. W. et al. (1990). “Selection criteria in coordination games: Some experimental results”. In: *The American Economic Review* 80.1, pp. 218–233.
- Corazzini, L., Cotton, C., and Valbonesi, P. (2015). “Donor coordination in project funding: Evidence from a threshold public goods experiment”. In: *Journal of Public Economics* 128, pp. 16–29.
- Crawford, V. P., Gneezy, U., and Rottenstreich, Y. (2008). “The power of focal points is limited: even minute payoff asymmetry may yield large coordination failures”. In: *The American Economic Review* 98.4, pp. 1443–58.
- Crosetto, P. and Filippin, A. (2013). “The ”bomb” risk elicitation task”. In: *Journal of Risk and Uncertainty* 47.1, pp. 31–65.
- Dal Bo, P. and Frechette, G. R. (2018). “On the determinants of cooperation in infinitely repeated games: A survey”. In: *Journal of Economic Literature* 56.1, pp. 60–114.
- Duffy, J., Xie, H., and Lee, Y.-J. (2013). “Social norms, information, and trust among strangers: Theory and evidence”. In: *Economic Theory* 52.2, pp. 669–708.
- Dutta, P. K. and Siconolfi, P. (2010). “Mixed strategy equilibria in repeated games with one-period memory”. In: *International Journal of Economic Theory* 6.1, pp. 167–187.
- Ely, J. C. and Valimaki, J. (2002). “A robust folk theorem for the prisoner’s dilemma”. In: *Journal of Economic Theory* 102.1, pp. 84–105.
- Ely, J. C., Hörner, J., and Olszewski, W. (2005). “Belief-free equilibria in repeated games”. In: *Econometrica* 73.2, pp. 377–415.
- Ely, J. C. and Välimäki, J. (2003). “Bad reputation”. In: *The Quarterly Journal of Economics* 118.3, pp. 785–814.

- Fehr, E. and Gächter, S. (2000). “Cooperation and punishment in public goods experiments”. In: *The American Economic Review* 90.4, pp. 980–994.
- Fischbacher, U. (2007). “z-Tree: Zurich toolbox for ready-made economic experiments”. In: *Experimental economics* 10.2, pp. 171–178.
- Fudenberg, D. and Levine, D. K. (1989). “Reputation and Equilibrium Selection in Games with a Patient Player”. In: *Econometrica* 57.4, pp. 759–778.
- (1992). “Maintaining a reputation when strategies are imperfectly observed”. In: *The Review of Economic Studies* 59.3, pp. 561–579.
- Gächter, S. et al. (2010). “Sequential vs. simultaneous contributions to public goods: Experimental evidence”. In: *Journal of Public Economics* 94.7-8, pp. 515–522.
- Garcia, J. and Van Veelen, M. (2016). “In and out of equilibrium I: Evolution of strategies in repeated games with discounting”. In: *Journal of Economic Theory* 161, pp. 161–189.
- Garcia, J., Van Veelen, M., and Traulsen, A. (2014). “Evil green beards: Tag recognition can also be used to withhold cooperation in structured populations”. In: *Journal of Theoretical Biology* 360, pp. 181–186.
- Guth, W., Ockenfels, P., and Wendel, M. (1997). “Cooperation based on trust. An experimental investigation”. In: *Journal of Economic Psychology* 18.1, pp. 15–43.
- Hao, D., Rong, Z., and Zhou, T. (2015). “Extortion under uncertainty: Zero-determinant strategies in noisy games”. In: *Physical Review E* 91.5, p. 052803.
- Harsanyi, J. C., Selten, R., et al. (1988). *A general theory of equilibrium selection in games*. Vol. 1. The MIT Press.
- Hashim, M. J., Kannan, K. N., and Maximiano, S. (2017). “Information feedback, targeting, and coordination: An experimental study”. In: *Information Systems Research* 28.2, pp. 289–308.
- Hayek, F. A. (1988). *The Fatal Conceit: The Errors of Socialism. The Collected Works of Friedrich August Hayek, Vol. I*. Ed. by W. W. B. III. London: Routledge.
- Heinemann, F., Nagel, R., and Ockenfels, P. (2004). “The theory of global games on test: experimental analysis of coordination games with public and private information”. In: *Econometrica* 72.5, pp. 1583–1599.
- (2009). “Measuring strategic uncertainty in coordination games”. In: *The Review of Economic Studies* 76.1, pp. 181–221.

- Herzenstein, M., Sonenshein, S., and Dholakia, U. M. (2011). “Tell me a good story and I may lend you money: The role of narratives in peer-to-peer lending decisions”. In: *Journal of Marketing Research* 48.SPL, S138–S149.
- Hilbe, C., Nowak, M., and Sigmund, K. (2013). “Evolution of extortion in Iterated Prisoner’s Dilemma games”. In: *Proceedings of the National Academy of Sciences* 110.17, pp. 6913–6918.
- Hilbe, C., Nowak, M., and Traulsen, A. (2013). “Adaptive dynamics of extortion and compliance”. In: *PLoS ONE* 8.11, pp. 1–9.
- Hilbe, C., Röhl, T., and Milinski, M. (2014). “Extortion subdues human players but is finally punished in the prisoner’s dilemma”. In: *Nature Communications* 5, pp. 1–6.
- Hilbe, C., Traulsen, A., and Sigmund, K. (2015). “Partners or rivals? Strategies for the iterated prisoner’s dilemma”. In: *Games and Economic Behavior* 92, pp. 41–52.
- Hilbe, C., Wu, B., et al. (2015). “Evolutionary performance of zero-determinant strategies in multiplayer games”. In: *Journal of Theoretical Biology* 374, pp. 115–124.
- Ichinose, G. and Masuda, N. (2018). “Zero-determinant strategies in finitely repeated games”. In: *Journal of Theoretical Biology* 438, pp. 61–77.
- Imhof, L. A., Fudenberg, D., and Nowak, M. A. (2007). “Tit-for-tat or win-stay, lose-shift?” In: *Journal of theoretical biology* 247.3, pp. 574–580.
- Isaac, M., Schmitz, D., and Walker, J. (1989). “The assurance problem in a laboratory market.” In: *Public Choice* 3.62, pp. 217–236.
- Isoni, A. et al. (2014). “Efficiency, equality, and labeling: An experimental investigation of focal points in explicit bargaining”. In: *The American Economic Review* 104.10, pp. 3256–87.
- Kalai, E., Samet, D., and Stanford, W. (1988). “A note on reactive equilibria in the discounted prisoner’s dilemma and associated games”. In: *International Journal of Game Theory* 17.3, pp. 177–186.
- Kanagaretnam, K. et al. (2010). “Trust and reciprocity with transparency and repeated interactions”. In: *Journal of Business Research* 63.3, pp. 241–247.
- Kandori, M. (1992). “Social norms and community enforcement”. In: *The Review of Economic Studies* 59.1, pp. 63–80.
- Keser, C. (2003). “Experimental games for the design of reputation management systems”. In: *IBM Systems Journal* 42.3, pp. 498–506.

- Kim, K. and Viswanathan, S. (2018). “The Experts in the Crowd: The Role of Reputable Investors in a Crowdfunding Market”. To appear in *MIS Quarterly*.
- Kreps, D. M. and Wilson, R. (1982). “Reputation and imperfect information”. In: *Journal of Economic Theory* 27.2, pp. 253–279.
- Kreps, D. M. (1996). “Corporate culture and economic theory”. In: *Firms, Organizations and Contracts*, Oxford University Press, Oxford, pp. 221–275.
- Kuppuswamy, V. and Bayus, B. L. (2018). “Crowdfunding creative ideas: The dynamics of project backers”. In: *The Economics of Crowdfunding*. Springer, pp. 151–182.
- Lee, C., Harper, M., and Fryer, D. (2015). “The art of war: Beyond memory-one strategies in population games”. In: *PLoS ONE* 10.3.
- Li, J. and Kendall, G. (2014). “The effect of memory size on the evolutionary stability of strategies in iterated prisoner’s dilemma”. In: *IEEE Transactions on Evolutionary Computation* 18.6, pp. 819–826.
- Lin, M., Prabhala, N. R., and Viswanathan, S. (2013). “Judging borrowers by the company they keep: Friendship networks and information asymmetry in online peer-to-peer lending”. In: *Management Science* 59.1, pp. 17–35.
- Mailath, G. J. and Samuelson, L. (2001). “Who wants a good reputation?” In: *The Review of Economic Studies* 68.2, pp. 415–441.
- (2006). *Repeated games and reputations: long-run relationships*. Oxford university press.
- Mäschle, O. (2012). *Which information should entrepreneurs on German crowdfunding-platforms disclose?* Tech. rep. Thünen-Series of Applied Economic Theory.
- McAvoy, A. and Hauert, C. (2016). “Autocratic strategies for iterated games with arbitrary action spaces”. In: *Proceedings of the National Academy of Sciences* 113.13, pp. 3573–3578.
- (2017). “Autocratic strategies for alternating games”. In: *Theoretical Population Biology* 113, pp. 13–22.
- Mehta, J., Starmer, C., and Sugden, R. (1994). “The nature of salience: An experimental investigation of pure coordination games”. In: *The American Economic Review* 84.3, pp. 658–673.
- Minozzi, W. (2015). “Anarchy in the lab”. working paper.
- Mollick, E. (2014). “The dynamics of crowdfunding: An exploratory study”. In: *Journal of Business Venturing* 29.1, pp. 1–16.

- Moritz, A., Block, J., and Lutz, E. (2015). “Investor communication in equity-based crowdfunding: a qualitative-empirical study”. In: *Qualitative Research in Financial Markets* 7.3, pp. 309–342.
- Morris, S. and Shin, H. S. (2002). “Social value of public information”. In: *The American Economic Review* 92.5, pp. 1521–1534.
- Nowak, M. (1990). “Stochastic strategies in the prisoner’s dilemma”. In: *Theoretical Population Biology* 38.1, pp. 93–112.
- (2006). “Five rules for the evolution of cooperation”. In: *Science* 314.5805, pp. 1560–1563.
- Nowak, M. and Sigmund, K. (1988). “Game Dynamical Aspects of the Prisoner’s Dilemma”.
- (1989). “Oscillations in the evolution of reciprocity”. In: *Journal of Theoretical Biology* 137, pp. 21–26.
- (1990). “The evolution of stochastic strategies in the prisoner’s dilemma”. In: *Acta Applicandae Mathematicae* 20.3, pp. 247–265.
- (1995). “Invasion Dynamics of the Finitely Repeated Prisoner’s Dilemma”. In: *Games and Economic Behavior* 11.02, pp. 364–390.
- Parravano, M. and Poulsen, O. (2015). “Stake size and the power of focal points in coordination games: Experimental evidence”. In: *Games and Economic Behavior* 94, pp. 191–199.
- Piccione, M. (2002). “The repeated prisoner’s dilemma with imperfect private monitoring”. In: *Journal of Economic Theory* 102.1, pp. 70–83.
- Press, W. H. and Dyson, F. J. (2012). “Iterated Prisoner’s Dilemma contains strategies that dominate any evolutionary opponent”. In: *Proceedings of the National Academy of Sciences* 109.26, pp. 10409–10413.
- Purcell, S. (2014). *The Definitive Guide to Equity & Debt Crowdfunding*. online book.
- Roma, P., Petruzzelli, A. M., and Perrone, G. (2017). “From the crowd to the market: The role of reward-based crowdfunding performance in attracting professional investors”. In: *Research Policy* 46.9, pp. 1606–1628.
- Schelling, T. C. (1960). *The strategy of conflict*. Cambridge, Mass.
- Selten, R. and Hammerstein, P. (1984). “Gaps in Harley’s argument on evolutionarily stable learning rules and in the logic of tit for tat”. In: *Behavioral and Brain Sciences* 7.1, pp. 115–116.
- Simon, H. A. (1971). “Designing organizations for an information rich world”. In: *Computers, communications, and the public interest*. Ed. by M. Greenberger. Baltimore, pp. 37–72.

- Solomon, J., Ma, W., and Wash, R. (2015). “Don’t wait!: How timing affects coordination of crowdfunding donations”. In: *Proceedings of the 18th acm conference on computer supported cooperative work & social computing*, pp. 547–556.
- Steinberg, S. et al. (2012). *The Crowdfunding Bible: How to Raise Money for Any Startup, Video Game Or Project*. online book.
- Stewart, A. J. and Plotkin, J. B. (2012). “Extortion and cooperation in the Prisoner’s Dilemma”. In: *Proceedings of the National Academy of Sciences* 109.26, pp. 10134–10135.
- Szolnoki, A. and Perc, M. (2014a). “Defection and extortion as unexpected catalysts of unconditional cooperation in structured populations”. In: *Scientific Reports* 4, pp. 1–6.
- (2014b). “Evolution of extortion in structured populations”. In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 89.2, pp. 1–5.
- Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1990). “Tacit coordination games, strategic uncertainty, and coordination failure”. In: *The American Economic Review* 80.1, pp. 234–248.
- Van Veelen, M. (2012). “Robustness against indirect invasions”. In: *Games and Economic Behavior* 74.1, pp. 382–393.
- Vismara, S. (2016). “Information cascades among investors in equity crowdfunding”. In: *Entrepreneurship Theory and Practice* November.
- Wang, Z. et al. (2016). “Extortion can outperform generosity in the iterated prisoner’s dilemma”. In: *Nature Communications* 7.12, pp. 1–7.
- Ward, C. and Ramachandran, V. (2010). “Crowdfunding the next hit: Microfunding online experience goods”. In: *Workshop on Computational Social Science and the Wisdom of Crowds at NIPS2010*, pp. 1–5.
- Wash, R. and Solomon, J. (2014). “Coordinating donors on crowdfunding websites”. In: *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. ACM, pp. 38–48.
- Xie, H. and Lee, Y.-J. (2012). “Social norms and trust among strangers”. In: *Games and Economic Behavior* 76.2, pp. 548–555.
- Zagorsky, B. M. et al. (2013). “Forgiver triumphs in alternating prisoner’s dilemma”. In: *PLoS ONE* 8.12, e80814.
- Zhang, J. and Liu, P. (2012). “Rational herding in microloan markets”. In: *Management Science* 58.5, pp. 892–912.