UNIVERSITY OF GHANA

COLLEGE OF HUMANITIES

DEPARTMENT OF ECONOMICS

HORIZON AND COOPERATION:
HOW THE LENGTH OF A GAME AFFECTS PLAYER COOPERATION

BY

GUY AUGUSTO DOS SANTOS

(10602596)

THIS THESIS IS SUBMITTED TO THE UNIVERSITY OF GHANA, LEGON IN PARTIAL FULFILMENT OF THE REQUIREMENT FOR THE AWARD OF MASTER OF PHILOSOPHY IN ECONOMICS

JULY 2019

DECLARATION

I hereby declare that this thesis is the product of my own original research undertaken under supervision and sent forth in fulfilment of the obligations required for the award of Master of Philosophy in Economics at the University of Ghana, Legon. All references used in the work have been fully accredited.

…………………………………... ………………………

GUY AUGUSTO DOS SANTOS DATE

(10602596)

…………………………………... ………………………

DR F. AGYIRE-TETTEY DATE

SUPERVISOR

…………………………………... ………………………

DR. A. BARIMAH DATE

SUPERVISOR

ABSTRACT

The study of conflict and cooperation in competitive situations has contributed greatly to various fields. From modelling to describing behaviour, the application of game theory in these fields has significantly improved their depth and robustness. Despite the field's usefulness there is still a significant disconnect between theoretical predictions and what is evidenced empirically. Illuminating this gap greatly explains the motivation behind this empirical study.

An experiment of unique parameters is conducted in order to test the robustness of the contemporary literature and explore how cooperation evolves with horizon. The study finds that cooperation does not always behave as the conventional literature suggests, even when multiple repetitions allow for subjects to gain experience. Game theoretic literature has found that in finitely repeated games the evolution of cooperation follows a distinct pattern of early cooperation followed by defection near the end of the game (Embrey et al. 2016). However our results show that this pattern does not always hold. We find that in a "stranger" setting, where players are paired only for a single round of play before being re-matched with another, they will disregard the horizon of finitely repeated PD games and treat the supergame as a series of one shot games. This finding was robust even when experience was accounted for. Consistent with recent literature on infinitely repeated games, horizon and experience both have a significant effect on cooperation, with the latter magnifying the former. However contrary to other studies on the subject, we observe that increases in the horizon decreased cooperation. This study suggests that parameters play an important role in how key variables affect cooperation and that they need to be factored in when using the prisoner's dilemma.

DEDICATION

I dedicate this study to God and those He worked through to bring this research to fruition.

ACKNOWLEDGEMENTS

I thank God for successfully bringing this work to completion. A special thanks to those who supported me through this process, especially to my father Tirso Dos Santos, whom with love, wisdom and diligence brought me to where I am. I also thank Dr. F. Agire-Tettey for his patience, openness and guidance throughout this process. And finally to my sister, Yeye, who in her own little way brought this study to life.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

PD	Prisoner's Dilemma
FRPD	Finitely Repeated Prisoner's Dilemma
IFRPD	Infinitely Repeated Prisoner's Dilemma

1. CHAPTER ONE:

INTRODUCTION

The study of conflict and cooperation in competitive situations has contributed greatly to various fields; in so much that it has created its own: Game Theory. From modelling to describing behaviour, the application of game theory in these fields has significantly improved their depth and robustness. In the economic context game theory has been integral to the description and modelling of agents' behaviour as they strategically interact with each other. It is a valuable tool; aiding in the fundamental analysis of industries and economic sectors, explaining and predicting economic behaviour, and even making prescriptions in social dilemmas and public policy[1].

Despite the field's usefulness there is still a significant disconnect between theoretical predictions and what is evidenced empirically. Illuminating this gap greatly explains the motivation behind this empirical study.

A central area of interest in the study of game theory in economics are the factors that affect cooperation between the players and how changes in these factors affect the outcome of the game. Chief among such factors is the length of the game that is its horizon. This study attempts to give a more nuanced view of how a game's length affects player cooperation in the hopes that it can help bridge the gap between theoretical predictions and what is observed empirically.

Time horizon plays an integral role in a game's outcome: its importance cannot be understated. This one variable affects nearly all other factors such as player behaviour and strategy, available equilibria and ultimately the game's outcome, as players can react to each other.

[1]Other examples include, natural resource extraction, the tragedy of the commons, Cournot competition, team collaboration with unobservable effort and public good provision, to name a few.

Games fall into two distinct categories based on their time horizon: finitely and infinitely repeated games. The study of these categories of games has specifically contributed to a wide range of economic approaches and helped describe otherwise puzzling economic phenomena. The economic subjects of auctions, duopolies, oligopolies, and pricing have especially benefited from the insights gleaned from the study of repeated games. For this study we will focus on the Prisoner's Dilemma (PD) game as it best captures the tension and cooperation problems caused by the mismatch between individual incentives and collective goals best (Rapoport & Chammah, 1965).

Finitely repeated games are those where the time horizon is known and observable; players know how many repetitions of the stage game will be played and know when the game will end. Theoretically players will play the Nash equilibrium strategy in all rounds of the game as they use backward induction to maximize their payoff[2].

In finitely repeated games experimental evidence has shown that cooperation rates are positive and generally declining over time, with a collapse of cooperative behaviour at the end of a supergame (Embrey, et al., 2016). This consistent result is opposite of the standard theory that cooperation is a dominated strategy and that the Nash equilibrium of choosing to defect will be played in all rounds.

While infinitely repeated games have unknown time horizon: players do not know how many repetitions of the stage game will be played and thus expect the game to continue indefinitely. A key feature is that players do not know when the last round is (or when the game will end) but may have an idea of how long the game will last (for example through random termination rules).

[2] See Chapter 3 for a more in depth analysis of how the standard theory reaches the Nash equilibrium

Theoretically in these games cooperation is sustainable and players can play a socially optimum strategy for all repetitions of the game.

In effect the players' knowledge and beliefs regarding the number of total repetitions in the game[3] is one of the most important factors in determining their behaviour and subsequently the game's outcomes. It is well understood that through repeated play more equilibria and higher levels of cooperation can be achieved. Nevertheless meaningful yet unanswered questions remain in game theoretic literature regarding a game's length and its effect on cooperation.

1.1. Statement Of The Problem

While there is evidence that longer lasting games increase cooperation in both finite and infinitely repeated games (Dal Bó, 2005; Cooper, et al., 1996), this relationship has not been thoroughly studied throughout game theoretic literature and largely remains inconclusive. Since the first prisoner's dilemma experiment conducted over 60 years ago (Flood, 1952), the question of whether people learn to cooperate or defect in PD games has been difficult to answer. Different studies give contradictory pictures of the evolution of play with increased repetition and experience. Some have suggested that the source of these contradictory results may be the diverse parameters employed. Different payoffs, horizon, design features and analysis all contribute to differing pictures of how player behaviour evolves.

This study serves to shed light on the relationship between a game's horizon and cooperation rates, specifically by studying, in a novel way, the effects between the number of repetitions played of a stage game and the level of cooperation between players. Much like Dal Bó (2005) this study

[3] The game's horizon or length

analyses and compares both finitely and infinitely repeated games, however this study employed a "block method" to induce infinitely repeated games and uses a Probit model to estimate changes in the likelihood of cooperation when game length is changed. The study also tests robustness of the relationship by using unique parameters.

1.2. Research Questions

In order to attempt to understand and explore the nuanced relationship between game length and player cooperation, the following questions will be asked:

I. Does the horizon affect cooperation in both finitely and infinitely repeated prisoner's dilemma games under novel parameters?

II. How does cooperation change when the horizon increases?

III. Is the effect of horizon on cooperation constant across supergames?

IV. How does experience effect the relationship between horizon and cooperation?

1.3. Objectives Of The Study

The goal of this study is to:

I. Test the robustness of recent studies that have found horizon to have a significant effect on cooperation in both finite and infinitely repeated games

II. Test whether the horizon behaves in the way the literature predicts

III. Understand the nuance of the relationship between horizon and cooperation

IV. Understand how experience can affect this relationship

1.4. Significance Of Study

The significance of this research is threefold:

First, it fleshes out a gap in game theoretic literature with robust empirical analysis: adding to the theoretical and empirical understanding concerning time horizon and cooperation in discrete games.

Economic literature has established that cooperative equilibriums can be supported through repeated play. Broadly speaking increasing the probability of future interactions increases the level of cooperation between agents (Dal Bó, 2005). The literature has shown that this is true when moving from one shot games to repeated games (Cooper, et al., 1996) and from shorter repeated games to ones of longer length (Dal Bó & Fréchette, 2016). Despite this relationship's empirical validity, the recent literature has not been robustly tested nor attempted to discover the possible nuances. For example, few attempts in the literature have empirically explored under what conditions and parameters this relationship holds.

This study serves to fill this gap by empirically testing the validity of the relationship by using novel parameters such as: a different method of inducing infinitely repeated games other than the usual random termination rule and employing a matching procedure that randomizes players after each round of play. These unique parameters test the robustness of the findings in recent studies by showing whether or not the relationship between horizon and cooperation still holds under different test conditions. The experiment also focuses specifically on repetition by employing various horizons and allowing for experience, while holding other key variables such as payoff constant, making the results robust in determining the horizons' effect on cooperation. The relationship between the length of a repeated game and cooperation between the agents is explored through mapping the changes in cooperation and varying the number of supergames the players will play to show how experience affects the relationship.

Secondly, it sheds light on the evolution of cooperation in repeated games. Isolating how game length affects cooperation can help further decompose the factors that dictate agents' behaviour.

Finally, the question of how game length affects cooperation between players and understanding the effect's nuances can have substantial real-world implications.

Multiple agents interact, cooperate and compete repeatedly to achieve their desired outcome in numerous arenas of life be it social, economic, political or otherwise. Understanding how different parameters and under what conditions cooperation can be sustained is important to ensure efficient and mutually beneficial outcomes. In many cases these socially optimum and desirable outcomes require agents to overcome the temptation to be completely self-serving. Theory and some empirical evidence suggest that with more repetition agents can overcome the temptation not to cooperate. Understanding under what conditions this happens is imperative to ensure those in charge of setting up the playing fields are aware of how the parameters chosen can affect the likelihood of cooperation and ultimately the potential implications on the outcomes for all stakeholders. It is possible that horizon can be treated as a variable to be optimized and as an additional tool to incentivize cooperation. Considering the effect each additional interaction has on agent behaviour and the game's outcome is integral if value is to be maximised.

1.5. Contribution To Knowledge

The findings of this study improve our knowledge of repeated games and lead to a more nuanced view of a game's time horizon. This study suggests that an experiment's parameter can affect how horizon possibility of optimizing or controlling for a game's length rather than arbitrarily designating one based on convenience. Accounting for the impact of game length could improve future models of player behaviour and the development of better game theoretic models, in the hopes that they become more robust and useful thus increasing their predictive power.

1.6. Definitions Of Terms

Round	Refers to a single repetition of the play
Horizon/Game Length	Total number of rounds or repetitions of play
Period	A set of rounds from which each horizon is taken. For the finite treatment in this study the period and horizon have the same number of rounds, while in the infinite treatment each period is set at ten rounds.
Supergame	A set made of five periods comprising of each horizon
Session	Refers to the entire treatment of three supergames
Finite Horizon:	Players know the number of repetitions of the stage game will be played
Infinite Horizon:	Players do not know how many repetitions of the stage game will be played
Prisoner's Dilemma:	The type of game the subjects play
Stage game:	The game the subjects play. In this study this is the Prisoner's Dilemma game
Payoff:	The total points the players get from playing the game
Players/Subjects:	Participants in the study
Experience:	Skill, expertise or familiarity gained through playing multiple repetitions

Average Cooperation: The measure of cooperation used in the study, refers to the mean of the subjects' choices aggregated over a particular period of play (round, supergame or session).

Strategy: This is a prescription or contingent plan that tells a player what action (C or D) to take in all circumstances they could possibly face in a PD game.

Matching Procedure: Protocol for pairing two different subjects together for a round of play.

Partner Setting: Matching procedure where players are paired for the entire period, horizon or supergame, after which they are matched with different players.

Stranger Setting: Matching procedure where players are randomly matched with other players after every round of play

1.7. Organization

Where Chapter One introduced the study, contextualized the subject matter and provided a list of definitions. The following chapters of this thesis are as follows: Chapter Two summarizes previous experimental research and literature on the topic. Chapter Three describes the experimental design, the methodology and theoretical predictions that drive the analysis of the results while Chapter Four presents, discusses and interprets the results of the experiment. Chapter Five concludes by summarizing the key findings of the study and makes recommendations for future study on the subject.

2. CHAPTER TWO:

LITERATURE REVIEW

2.1. Introduction

This study uses the game theoretical framework based on that initially built by Von Neumann and Morgenstern and since expounded upon by various academics and researchers. This framework models economic behaviour through "games of strategy" to address several key limitations and basic problems of general economic theory, especially of those inherent with the assumption of rational economic behavior. For example, the use of games removes this assumption and introduces the possibility of irrational economic behavior, thus making these theories more mathematically complete and general (Von Neumann & Morgenstern, 1944). This section surveys the economic literature and prevailing theories relevant to this game theoretic study.

The basic theoretical foundation underlying this study is similar to those found in seminal studies on Folk Theorems, such as (Friedman, 1971), (Benoit & Krishna, 1985), (Rubinstein, 1980) and (Aumann & Shapley, 1976), which posit that in infinitely repeated games with sufficiently little discounting, any individually rational outcome can become the Nash equilibrium. This is because under repeated play, players can respond to each other's actions and must take into account their opponent's reactions when employing their strategies.

2.2. Finitely Repeated Games

Experimental evidence in finitely repeated games show similar evolutions in cooperation as the game progresses. There is a preponderance of positive cooperation in the early rounds followed by a significant collapse or “unravelling” of cooperation in the final rounds as subjects begin to defect as the game nears a close (Kreps et al., 1982; Selten & Stoecker, 1986; Embrey et al., 2016). Through repeated play subjects learn to defect quicker in finitely repeated PD game thus in subsequent games the unravelling of cooperation tends to happen earlier as a result. Despite the collapse of cooperative play near the end of the game, subjects end up with payoffs that are strictly bigger than they would have obtained under equilibrium play.

For example Selten and Stoecker (1986) examine player behaviour in a finitely repeated PD with a ten round horizon. Subjects played 25 supergames and were re-paired after each supergame. They observed that player behaviour converges to a particular pattern with experience: in early rounds players jointly cooperate, once defection is initiated by either player, in subsequent rounds there is joint defection. Significantly, the experimenters determine that with experience, the point at which players intend to first deviate moves earlier. In the initial supergames players learn to cooperate and with experience they learn the risks of not defecting first. At this point cooperation begins to unravel (Roth, 1988). These early PD experiments had the impression that the decline of cooperation comes about with experience.

Recent experimental evidence echo Selten and Stoecker’s (1986) findings, showing that cooperation rates are positive and generally declining over time, with a collapse of cooperative behaviour at the end of a supergame (Embrey, et al., 2016).

This consistent result is contrary to the standard theory that cooperation is a dominated strategy and that the uncooperative pareto-dominated Nash equilibrium will be played in all rounds of the finitely repeated PD game.

Two main theories have emerged to explain this phenomenon, both of which focus on behavioural aspects of the players which are discussed below.

2.3. Incomplete Information Theory

The first theory proposed to explain the preponderance of player cooperation observed in finitely repeated PD games was made by Kreps et al. (1982). It maintains the assumption that players are self-interested, however the repeated nature of the game creates incentives for players to cooperate. The central assumption in this theory is that there is uncertainty in what type of opponent is being played against. Players hold a small belief that that their opponent is a cooperative player, which induces the self-interested players to cooperate. However the common knowledge of rationality does not hold under this theoretical structure.

2.4. Reputation Building Theory

The second theory to explain the high level of cooperation in finitely repeated games postulates that not all agents in the game are completely self- interested; that some of the players are altruists and benefit from cooperation in a manner not reflected in the payoff matrix[4] (Andreoni & Miller, 1993) Theories based on this postulate that cooperation is not a dominated strategy as the true payoffs differ from those given in the PD game.

Studies that compare the two theories find that while there is evidence of both altruism and reputation building; both theories alone fail to model the preponderance of cooperative behaviour in finitely repeated games and fail to consistently explain the pattern of play (Cooper, et al., 1996). This has led to the further exploration of cooperation in finite games using hybrid models and other explanatory variables.

[4] In the literature this added benefit from cooperation is referred to as a “warm glow”

2.5. Horizon In Finitely Repeated PD Games

Experiments into finitely repeated PD (FRPD) games show that as horizon increases from a one-shot PD game to a repeated PD game, cooperation increases significantly. However the overall evolution of cooperation in FRPD of differing horizons remains the same: high cooperation in the initial rounds followed by a breakdown of cooperation in the latter rounds, experience causing players learn to defect sooner.

Cooper et al. (1996) conducts a study to separate the altruism and reputation building hypotheses. The study compares two treatments: one with 20 one shot[5] PD games and the other FRPD games with a ten round horizon. The experiment had a single horizon and set of unchanging payoffs. They find higher cooperation rates in the finitely repeated PD treatment than in the one shot treatment. Their results show that cooperation rates start above 50% in the finitely repeated treatment and end below, however cooperation rates are always lower for the one shot game. Noticeably, cooperation is significantly above zero in both treatments. The study omits the analysis of experience and the evolution of player behaviour due to the limited number of repetitions. The experimenters conclude that while there is evidence of both altruism and reputation building; both theories alone fail to model the preponderance of cooperative behaviour in finitely repeated games and fail to consistently explain the pattern of play (Cooper, et al., 1996). This has led to the further exploration of cooperation in finite games using hybrid models and other explanatory variables.

Bereby-Meyer and Roth (2006) compare one shot PD games to play in FRPD games with either stochastic or deterministic payoffs. The one shot treatment involved 200 rounds with random re-matching, while the FRPD game consisted of 20 supergames with a ten round horizon. They find

[5] One shot refers to a game played only for a single round or repetition

higher cooperation in the first round of the repeated games than in the one shot games. With experience subjects in the repeated treatment learn to cooperate more in earlier rounds and less towards the end of the supergame. They report that stochastic payoffs dampen this effect and interpret their results to be consistent with reinforcement learning models: given expected payoffs, adding uncertainty to the link between an action and its consequences, will slow learning.

Many recent papers study heterogeneity in cooperative player behaviour and the role of reputation building in the FRPD. For example Schneider and Weber (2013) allow subjects to choose the game length (or horizon) of each supergame. They observe that a commitment to long-term relationships works as a screening device. Generally it is found that conditionally cooperative players are more likely to commit to longer term relationships relative to uncooperative players. The study reports that longer games facilitate more cooperation and that even higher cooperation can be achieved through endogenously chosen long-term commitments, rather than if the game length is exogenously imposed.

2.6. Infinitely Repeated Games

Initial studies into the relationship between horizon and cooperation showed it to be limited (Roth & Murnighan 1978; Murnighan & Roth 1983; Feinberg & Husted 1993). However in the past decade the question of repetition has gained significant interest in the study of infinitely repeated games. The lack of sharp predictions has caused the standard theory of infinitely repeated games to be heavily criticized and prompted further empirical investigation.

Roth and Murnighan (1978) were first to address whether infinite repetition affects cooperation experimentally. They found that cooperation tends to increase with the probability of continuation, however the effect was not consistently monotonic. This led to their conclusion that the results [of whether higher horizon affects cooperation rates] were "equivocal" (Roth, 1995). Similarly, Feinberg and Husted (1993) when studying collusive equilibria in duopoly markets found that the effect of repetition on cooperation was not large.

Unlike the previous studies based on PD games Palfrey and Rosenthal (1994) employed a public good game for the stage game of their experiment. They concluded that the impact of repetition on cooperation is limited after their results showed that the effect of increased repetition on contribution rates had a small magnitude despite being larger under a probability of continuation of 0.9 than those in one shot games.

More recent papers have provided a more positive answer to the question of repetition's effect on cooperation. This contemporary literature shows that cooperation increases with the probability of future interactions and the effect's magnitude increases with experience. Unlike previous papers where players were limited to playing one supergame, these papers allowed players to participate in several repeated games and gain experience.

For example, Dal Bó (2005) was first to present the idea that "the shadow of future" cast by longer horizons does in fact increase cooperation. The study shows that cooperation is four times higher when going from a one shot prisoner's dilemma game to an infinitely repeated game with a probability of continuation of 0.75. The experiment was conducted using finitely repeated PD experiments as controls for the study. The experiment used two stage-game payoffs for horizons of 1, 2 or 4 rounds. The primary focus of the paper is to compare player behaviour in finitely repeated PD games to player behaviour in randomly terminated or infinitely repeated PD games of the same expected length. The results show the infinite treatment has much higher cooperation rates in the first round than when the treatment where the PD games are finitely repeated. In fact aggregate cooperation rates decline with experience for the finitely repeated games. Consistent with studies FRPD, there is a significant deterioration of cooperation rates in the final rounds when cooperation is observed within a supergame of the finite horizon. Also consistent with previous findings for the finite horizon such as Cooper et al (1996), first round cooperation is higher in games with a longer horizon.

The main difference between the initial studies that found repetition or horizon's effect on cooperation to be limited and these more recent papers is that the latter allow subjects to play several repeated games and thus gain experience. This suggests while in an uncertain horizon having repeated interactions is important for the subjects' ability to support cooperation only once they have gained experience.

2.7. Comparison of the FRPD and the IRPD

Subjects with enough experience, evidence shows, will behave differently under finitely repeated and infinitely repeated games. After comparing finitely and infinitely repeated trust games, Engle-Warnick and Slonim (2004) find that once players have gained significant experience they become more trustworthy and trust others more under infinite repetition than finite.

Dal Bó (2005) also finds that cooperation is much greater under infinitely repeated PD games than finitely repeated ones. However, Lugovskyy, et al. (2017), compared the behaviour of subjects in finitely and infinitely repeated public goods games and did not find contributions to be higher in the infinite treatment. Their results did show however, that the fall in contributions between the first and last round was much greater in the finite case. Similarly, Normann and Wallace (2012) find no clear difference in cooperation between the two treatments.

Engle-Warnick and Slonim (2006a) show that in infinitely repeated trust games, the length of the supergames correlates positively with subsequent choices. A finding echoed by Dal Bó and Fréchette (2011) whom show that the current supergame length affects the probability of the subjects cooperating in the following supergame, despite being paired with completely different opponents.

Several papers including Camera and Casari (2009), Embrey et al. (2013), Sherstyuk et al. (2013), Bernard, et al. (2014), and Fréchette and Yuksel (2014) repeat these results: the longer the supergame, the more likely the subjects are to cooperate. Interestingly the opposite also seems to true; as subjects observed are more likely to defect in shorter supergames.

The process is consistent with subjects using their experience to update their beliefs regarding the proportion of the population that will cooperate and thus increasing their likelihood of cooperation and the value of cooperating.

The results of these contemporary papers suggest that cooperation is affected by: how robust cooperation is to strategic uncertainty, particularly when cooperation is risk dominant, the realized length of previous supergames and the choices of the past subjects with whom the player was paired with (Dal Bó & Fréchette, 2011).

The theory of repeated games has shown that repetition can lead to cooperation by helping people overcome opportunistic behaviour. This idea has been applied in many economic fields and other social sciences, yet only recently has experimental literature largely considered the limits of its validity and under what conditions or parameters it holds. Further study is needed and can help us understand the conditions under which longer horizons can lead to cooperation and ascertain the nuances of how this relationship evolves.

3. CHAPTER THREE:

METHODOLOGY

This chapter explains the methodological approach this study will employ and explores some underlying assumptions and theoretical underpinnings of how horizon and cooperation arise. The chapter is broadly separated into: a general explanation of the prisoner's dilemma game and the dominant standard theory, a discussion regarding the variables and parameters and an explanation of the experiment, data collection and analysis.

3.1. Overview of the Prisoner's Dilemma

An overview of the Prisoner's Dilemma stage game is as follows:

Assuming two players (Player 1 and Player 2) have two actions to choose from: Cooperate (C) or Defect (D) the game is specified by the following payoffs:

- mutually defecting players each receive the punishment payoff (P),
- a cooperating player that matches with a defecting player receives the sucker's payoff (S),
- a defecting player matching with a cooperating player will get the temptation payoff (T), and
- mutually cooperating players each receive the reward payoff (R)

For the game to be considered a PD game the payoffs must be such that:

$$T > R > P > S$$

The tension in the PD game[6] stems from the fact that regardless of what the other player chooses, defecting (D) has a greater payoff than cooperating (C) (i.e. T > R and P > S). Mutual cooperation however yields a higher payoff than mutually defecting (R > P). Therefore rational actors will end up with the less desirable payoff from mutual defection (P), despite there being a pareto-optimal payoff (R) from mutually cooperating.

The prisoners' dilemma used in this study take the standard form shown in Figure 3.1

Figure 3.1 Basic Prisoner's Dilemma Payoff Matrix

		Player 2	
		D	C
Player 1	D	(P; P)	(T; S)
	C	(S; T)	(R; R)

Figure 3.1 Payoff table where Player 1 plays from the perspective of the row and Player 2 from the column
P is the Punishment payoff, S is the sucker's payoff, T is the Temptation payoff and R is the Reward payoff

In this experiment we assume that the payoffs are true payoffs; that is the payoffs reflect accurately the utility the players will receive from each outcome.

Reputation building theories state that payoffs in the stage game are not true reflections of the utility players receive; that there is an additional utility (or "warm glow") players receive from either cooperating or defecting that is unaccounted for in the game's payoff matrix.

[6] For this to be a repeated Prisoner's Dilemma, it must hold that $T > R > P > S$, and $R > \frac{(S+T)}{2}$ so alternating between cooperation and defection yields a less efficient outcome that mutually cooperating

This has brought about the categorization of players as altruists or egoists in cooperation games. The unaccounted utility changes the payoff matrix and is purported to explain the difference between theoretical predictions and what is observed empirically. This theory has seen mixed results in the literature would not specifically change the study's findings has not been taken into account in this study.

To understand the basis of this study, one must understand the theoretical underpinnings of how repetition can change the players' strategies and incentives, and thus final outcomes of the game.

Based on the standard theory in a PD game, playing the uncooperative action "D" strictly dominates playing the cooperative action "C", thus the only Nash equilibrium is where both players choose to play action "D" (DD) and the equilibrium outcome id the payoff (1; 1) as shown in Figure 3.2.

Figure 3.2 PD Payoff Matrix with Figures

		Player 2	
		D	C
Player 1	D	(1; 1)	(4; 0)
	C	(0; 4)	(3; 3)

Figure 3.2 Payoff Matrix where: T = 4, R = 3, P = 1, S = 0

Even though players can achieve higher collective payoffs if both Player 1 and Player 2 choose to cooperate (C), each player's best response, given the other's actions, is to defect (D). Thus defection is the best response strategy for each player, causing the outcome between two rational players to always be an inefficient one. Herein lies the dilemma: the best response strategy create a socially inefficient pareto-dominated outcome.

3.2. Theoretical Underpinnings of Increased Repetition on PD

In a one shot (one round) PD game there is no escape from this dilemma; the game parameters would have to be changed. Any argument for choosing to cooperate would have to be based on the argument that payoffs in the matrix do not represent the players' preferences, or that some facet of the game is not indicative of the actual interaction. When the game is repeated however, things change significantly.

Suppose player 1 has a utility function, u_1 and maximizes the normalized discounted sum of payoffs according to equation 3.1 below.

$$(1-\delta)\sum_{t=0}^{\infty}\delta^t u_1(a_1^t, a_2^t) \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots.. (1)$$

Equation 3.1 shows the normalized discounted sum of payoffs

Where: u_1 is Player 1's utility function; the $\delta \in (0, 1)$ the common discount factor; t is the period, a_1^t is the action Player 1 chooses in period t; and a_1^t is the action[7] Player 2 chooses in period t.

Using the PD payoff matrix found in Figure 3.2, assume that the strategy of Player 1 is to choose C in the first round and every subsequent round if Player 2 chooses C, while choosing D in all other circumstances (known commonly as the "Grimm" Strategy). Suppose finally the period T is reached. The player can either continue to play C ad infinitum which, given Player 2's strategy,

[7] The actions for both Player 1 and Player 2 are either to Cooperate (C) or to Defect (D)

would yield a (discounted, normalized) payoff of 3; or play D after which each player can do no better than play D in all subsequent periods, as the Player 2 will retaliate and defect continuously thus yielding the following payoff[8]:

$$(1+\delta) = [4 + \sum_{t=t+1}^{\infty} \delta^{t-T}\, 0] = 4(1-\delta) \quad (2)$$

Equation 3.2 Payoff given Player 1 defects in period T

Therefore continued cooperation is optimal if:

$$3 \geq 4(1-\delta) \quad (3)$$

Equation 3.3 The condition for continued cooperation between Player 1 and Player 2

or $\delta \geq \frac{1}{4}$

As future play depends on current actions, incentives can be altered to shape these actions, in this case allowing an equilibrium where players choose C. This example of a folk theorem for repeated games shows that through repeated play coordination occur and yield equilibria closer to the socially optimum. A key assumption here is that players are sufficiently patient and the future payoffs are sufficiently important.

Nevertheless, the standard theory concerning finitely repeated games of T repetitions predicts that the players repeat the one-stage Nash equilibrium in all rounds of the game. As the vectors of player strategies grows super exponentially with number of repetitions in a repeated game, a two round finitely repeated PD game of Figure 3.2 will be used an example illustrates this.

[8] These payoffs are the same as those found in the PD Payoff Matrix in figure 3.2

Using backward induction reveals that the strategy players employ in the one shot equilibrium (DD) is repeated again in both repetitions of the two round game[9].

To prove that at every equilibrium of the two round repeated PD game[10] the players play (DD) in both stages, suppose a case where there exists an equilibrium with a positive probability that a player does not play (DD) at some stage, and that the player in this case chooses to play C. Let $t \in \{1, 2\}$ be the last round where there is positive probability that the players will not play (DD).

Case 1 t=1

Consider the strategy of a player who plays D in both rounds. This strategy grants them a higher payoff. Since C is strictly dominated by D, the player's payoff rises if they switch from C to D in the first round: as the player's payoff is either 1 or 4 depending on the second player's choice. The following round both players play D by assumption. Thus, the sum total payoff of the player increases.

Case 2 t = 2

Consider the player strategy where in the first round he plays the original strategy and in the second round he plays D. His playoff in the first stage does not change, but because D strictly dominates C, his second stage payoff increases. Thus, his total payoff increases.

This same rationale applied to T-stage repeated PD shows that, at equilibrium, the players play (DD) in every round.

[9] This is a special case of the general claim which states every strategy where players play an equilibrium of the base game is an equilibrium of the T round game
[10] And by extension any finite T-stage repeated PD game

Due to the proliferation of strategies caused by repeated games, this study will not attempt to compute all equilibria of the repeated games. The main theoretical takeaway regarding finitely repeated games is the theoretical prediction that players will always choose to defect in every round. Rather than strategies this study instead focuses on actions and outcomes. Another issue with focusing on strategy would be the ambiguity as multiple strategies can yield the same actions and outcome, thus experimental data cannot accurately observe them. For example in the case where subjects are paired for entire supergames, if a cooperating subject is paired with a partner who chose to cooperate for every round of a repeated game the equilibrium outcome and actions are consistent with their playing a “Grimm” strategy and also “Always Cooperate” strategy.

In the case of this study as the matching procedure changes players’ partners after every round of play, every round played will have a different partner. Therefore strategies such as the Grimm strategy cannot be correctly employed: this is because the punishment mechanism of the strategy would not be born on the defecting player but all future players.

3.3. Furthering Repetition to the Infinite Case

When the PD is extended even further to the infinite stage, things change considerably. As with the finite case there are multiple equilibrium in the infinitely repeated PD[11] including the finite best response equilibrium where all players persistently defect regardless of past play. However in infinitely repeated games coordination can yield equilibria closer to the socially optimum.

[11] The infinite repeated game has even more equilibria than that of the finite case

3.4. Cooperation

Cooperation is the key variable of study in game theoretic research and so needs explicitly mentioned and analysed in the context of longer horizons. While early PD experiments on the matter found limited and equivocal results when it came to the impact of repetition and higher horizon on cooperation (Feinberg & Husted, 1993) in repeated PD games, recent literature on the determinants of cooperation have shown that higher horizons lead to increases in cooperation between players if subjects are allowed to gain experience (Dal Bó & Fréchette, 2016).

In this study we use two measures of cooperation: the arithmetic mean of the cooperative choice (C) aggregated for all players over a given period of play[12] and the likelihood of a player choosing to cooperative choice (C).

3.5. Experimental Approach

Our methodological approach varies the number of supergames allowing for experience and tests the impact of horizon by using five horizons of increasing length. Aside from the high number of discrete horizons, the parameters of the study are unique from those found in the majority of the literature (a discussion regarding the parameters of this study can be found below). The results from the study form a robustness test for recent literature concerning the horizon and provide insight into the relationship between cooperation and repetition.

[12] Period of play here could be a round, horizon, period, supergame or session

3.6. Overview of Experiment

This experiment was designed to be a reliable test of the hypothesis concerning horizon and cooperation. The experiment is based on a simple PD game where the subjects interact anonymously with each other using computer terminals. Subjects were paired with different opponents after each round of play. This matching procedure aims to capture interactions where subjects have no prior experience nor knowledge about their opponent thus controlling for reputation building and partner specific learning. The subjects were paid an initial participation fee of 10GHC (Ghana Cedi) and given a chance to double the amount through play (subjects chance to double their earnings increased proportionally to their accrued points). The conversion between points and probability of doubling their earnings ensured that players had significant incentives to increase their earnings. Two separate sessions for the finite and the infinite were held on different days. Players were given a practice supergame of two rounds before beginning the actual game session to: familiarize themselves with the controls, control for early learning and ensure that all systems were working as expected. The game session consists of three supergames, each supergame was made up of five periods of the PD stage game. Each period coincided with a horizon of a certain number of rounds. At the end of the session, the points were tallied and the player earnings were disbursed.

3.7. Data Collection, Processing and Selection Criterion

According to the results of previous empirical studies on the matter, personal characteristics such as students' major (Dreber et al. 2014; Sherstyuk et al. 2013) and gender (Sherstyuk et al. 2013;

Proto et al. 2014) had minimal to no effect on cooperation. Therefore, other than age[13], variables such as academics and demographics were not factored into the subject selection criteria.

Subjects were taken from the University of Ghana Legon Campus where the experiment was held. As the experiment would last for up to several hours and subjects needed enough time to complete the experiment in its entirety, university students with no classes on the day of the experiment and be old enough to give consent were chosen. An open call was made for students to be a part of the experiment that fit these criteria. Subjects were collected through advertising and word of mouth at the university. The participants were predominantly male[14] undergraduate students between the ages of 18 and 25 of varying majors.

The experiment was held over two days with the finite treatment being conducted on the first day while the infinite treatment, with new subjects, was conducted on the second day. The experiments were conducted at the same time of day, in the same computer lab, by the same invigilators in order to ensure a consistency in the environment. This is important to help control for any unknown variation due to the way the sessions were conducted. The data collection was done through the z-Tree (Fischbacher 2007) computer program and the data analysis was done through Stata 13 data processing software.

[13] Age is considered as subjects must be of age to give consent, however there is little empirical evidence that age is a significant factor in affecting cooperation

[14] Over 85% of the subjects in both the finite and infinite treatments were male

3.8. Discussion of Parameters

This study relies heavily on the belief that the parameters of repeated games play a key role in cooperative outcomes in PD games. As such it is important to address them explicitly. Below is a brief discussion of key experiment parameters beginning with the payoffs.

3.8.1. PD Payoffs

The payoffs for the study were kept constant throughout all PD stage games for both treatments. This was in order to limit the variables subjects had to consider and adjust to when making their decisions. The experiment was conducted such that the only dynamic parametric variable was in fact the change in horizon.

The payoff parameters of the experiment were restricted to:

$$R > \frac{S+T}{2} > P \quad \text{..(4)}$$

Where the first inequality,

$$R > \frac{S+T}{2} \quad \text{..(5)}$$

ensures that mutual cooperation is more efficient than that of an asymmetric outcome. This guarantees that alternating between cooperation and defection is not an efficient strategy for subjects to employ.

The second inequality,

$$\frac{(S+T)}{2} > P \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (6)$$

infers that choosing to cooperate always improves efficiency, a condition relatively overlooked in PD literature save for Friedman and Sinervo (2016).

Friedman and Oprea (2012) study 4 different stage game payoffs with a fixed 8 round horizon. The study found that cooperation rates increase with experience when payoffs keep the temptation to defect low and create high efficiency gains from choosing to cooperate. However when payoffs are not conducive to cooperation then cooperation rates decrease. The payoff matrix used for this study is the same found in Figure 3.1.1 where: T = 4, R = 3, P = 1, S = 0

The payoff adheres to both inequalities:

$$3 > \frac{0+4}{2} > 1 \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (7)$$

3.8.2. Matching Procedure

Subjects were paired with different opponents after each round of play. This parameter was chosen with the specific aim to capture interactions where players have no prior experience or knowledge about their opponent: controlling for reputation building and learning opponent behaviour. This means players face a high degree of uncertainty about their partners' "type". Like Gilboa and Schmeidler (1989) we assume that subjects resolve this uncertainty by using the most pessimistic among multiple priors. As the subjects are randomized after each round of play cooperation becomes more difficult to sustain, as it is harder to build reputations, punish a specific defecting partner or threaten to punish a potentially defecting partner. Indeed authors such as Andreoni and

Miller (1993) and Cooper et al. (1996), found that more cooperation arises in "Partner" settings; where players are matched with the same opponent for the entire supergame. Boone, et al. (1996) however found there to be no difference, while others found even more cooperation in "Stranger" settings (Andreoni, 1988).

Logic suggests, at least in infinitely repeated games, that there should be more cooperation in the former. However, this higher cooperation found in the "Partner" setting can be offset by path dependence and stronger negative reciprocity (Mengel & Peeters, 2011).

A general explanation of how a "Partner" setting can affect cooperation in repeated PD games is as follows:

Suppose two people are aware they will be playing a repeated PD game, many times in succession. These two people are essentially engaged in a continuing strategic relationship. It seems likely cooperation may arise from this situation. We can expect that in nearly every round of play each player will choose the cooperate action, despite the action being strongly dominated when the game is played just once. Indeed, in many recent laboratory experiments when subjects play the repeated PD game, this cooperative pattern of play is usually observed. This cooperation result is also observed in real-life continuing strategic relationships such as those found in oligopolies.

Why would we expect to see cooperative play? When the PD game is being played repeatedly between a matched pair of agents, cooperation can be construed as a signal that they will continue to cooperate in future rounds. It establishes a "cooperative" reputation, and in order not to lose this cooperation the other player is induced to cooperate as well. The same argument applies to the other player: they may cooperate for the same reason— in order to establish their own cooperative reputation, hoping it will induce future cooperation.

Additionally, a defection by a player can be punished by the other choosing to defect in future rounds. This will unequivocally harm the initially defecting player, as they lose out on their opponents' cooperation in future rounds. Punishing defection and establishing a reputation creates a cooperative pattern of behavior that could only exist in a repeated setting with the same players. Specifically, it is the continuing nature of the players' strategic relationship makes punishment and reputations possible.

Therefore, a simple change in the matching procedure can have significant implications on the subjects' strategies and their ability to cooperate. When a continued strategic relationship is no longer possible, cooperation can suffer as the partner effects of reputation building and punishment disappear. Many experiments in the literature focus on the "Partner" setting, however this study will concentrate on the lesser studied "Stranger" setting.

The "Stranger" setting employed in this experiment through its matching procedure[15], has the added benefit of isolating the effect caused by changing the horizon, as cooperation rates are much less affected by these partnership effects.

3.8.2.1. Contagion and Punishment Discussion

Research on random matching games similar to this study is sparse. However a study by Ellison (1994) uses a similar matching procedure where: subjects could only observe the outcomes of the games in which they participated and they could neither communicate nor recognize the identity of any past opponents. The principal conclusion of the study was that cooperation is possible in equilibrium and is relatively robust, even in games with anonymous random matching.

[15] Players change their opponent after each repetition

Kandori (1992) finds that cooperation is possible if players employ strategies with "contagious" punishments under special payoffs. Ellison (1994) extends this finding, showing that cooperation under a "Stranger" setting is possible for general payoffs with sufficiently patient players.

The contagion effect stems from a breakdown of cooperation after a single defection, where a player initially defects in one period, causing his opponent to defect in the next[16] thus infecting the next player who then defects from then on. The contagious process shows that a defecting player in the first period will cause two players to defect in the second period and four to defect in the third, and so on. After one player deviates the collapse of social cooperation punishes all players. Unlike the "Grimm" strategies employed in Partner matching protocols where defecting players are immediately punished, the contagious punishment takes time to spread though the population and reach the defecting player.

Randomizations play two key roles in these types of games, first they provide a coordinating device so all players in the study can simultaneously return to cooperation in the end of a punishment phase. This is important as all players only just prefer to cooperate when all others are cooperating; no player would return to cooperation if the probability of all other players returning to cooperation is not close to one. Coordination allows the construction of a stable universal equilibrium. The second is to adjust for the expected duration of the punishment[17]. Thus players can employ punishments that can deter defection but are not so severe that players would not be willing to carry them out. For example in the first period of our study the severity of punishments are relatively low, as there are only two repetitions of the PD game played, however the severity of the punishment increases with each subsequent period.

[16] This player will choose to defect in the next round of play and all future rounds of play in the period.
[17] The expected duration of the punishment becomes the severity of the punishment. Earlier defections cause longer durations thus there is a higher punishment severity for earlier defections and for longer games.

Ellison (1994) finds contagious punishments to be a fairly powerful tool for enforcing cooperation, sustaining cooperative equilibria despite large populations of somewhat patient players with infrequent individual interactions.

3.8.3. Information Availability

At the bottom of players' screen a play history summary is displayed. This history includes: past player choices, previous game outcomes and a running point score for the period of play. This running history resets at the end of each period. In the end of the supergame the point total the player achieved is displayed to them. This helps subjects learn and strategize while also reminding subjects about game length. In the finite treatment the subjects are shown the number of rounds of play left before the end of the period.

3.8.4. Horizon

3.8.4.1. Finite Treatment

In the finite treatment, three supergames of finitely repeated games with varying horizons were played. The number of repetitions in each horizon was 2, 4, 6, 8, and 10. These horizons also matched the realized number of rounds of the infinitely repeated games. This enables direct comparison between both treatments.

3.8.4.2. Infinite Treatment and the Block Method

The infinite treatment employed a "Block" method for inducing an infinite horizon, this method was modified from the block random termination design found in Fréchette and Yuksel (2014).

A popular method for inducing an infinite time horizon in a laboratory setting is by using a random termination rule. This method works as follows: after every repetition of the stage game there is a fixed probability (d) that the game will continue for an additional repetition and a probability (1-d) that the supergame will end. These probabilities are known to all players. The Block Random Termination method of inducing infinite time horizons created by Fréchette and Yuksel (2014), plays as the standard random termination, however it is conducted in blocks of a pre-announced length (number of rounds). Within a block, players do not get feedback about whether or not the period has continued until that round. Once the end of a block is reached, the players are told whether the supergame ended within that block and if so, in which particular round. If the supergame has not yet ended, they continue play into a new block. Subjects are paid only for the rounds up to the end of the supergame and all play and outcomes for subsequent rounds in the block are void.

While the two methods are theoretically equivalent their results show that random termination shows slightly higher levels of cooperation and the block random termination method was shown to be less affected by experience. Nevertheless both methods generate sharp and congruent comparative statics with only differences in their magnitudes, for example both show that when a player's opponent in the previous round cooperates, one is more likely to cooperate in their next round. The researchers find that both methods generate behavior consistent with theory and that the block random termination method seems to the most appropriate method if one needs to observe more rounds but requires the strategy to be close to what they would be under random termination.

In this study we modify the Block Random Termination method by using fixed lengths for each horizon that would not necessitate play into another block. We fix the horizons to 2, 4, 6, 8, and 10 rounds and use a block size of ten rounds. Thus subjects play a period of ten rounds. Each period corresponded to a particular horizon and the number of rounds that counted towards point allocation was determined by the period's horizon. For example the first period of the infinite treatment is associated with the first horizon, thus only the first two rounds of the stage game will count towards allocating points to the player and for analysis. The second period is associated with the second horizon and only the first four rounds of the stage game will count, so on until the final fifth period and horizon where all the ten rounds count. Subjects were verbally told and made to understand that although the number of rounds for each period of play was fixed at ten rounds, the number of rounds that would be counted from each period would not only vary. They were also made to understand that for each successive period the probability of the total number of rounds to be counted increases- with the first period having the lowest probability and the fifth having the highest. In this manner subjects understood that the horizon increased with every period[18]. At the end of the session the correct number of rounds would be picked from each period played and the results of those particular chosen rounds would be used to allot points to the players, while the rest would be discarded and not count towards the players' final scores.

The motivation behind the use of this approach to induce the infinite horizon is to ensure direct comparability between the two time horizons whilst controlling for the subjects' belief of future interaction. This modification also avoids the "restart effect" found in Fréchette and Yuksel (2014)

[18] Players were made to understand that horizons reset after each supergame

where cooperation rates spiked upwards at the beginning of new blocks, only to decrease back to levels expected from the continuation of the previous rounds of play.

In order to provide a test for whether the block method does in fact induce an infinite horizon in a laboratory setting we use three variables that have historically had statistically significant effects on cooperation rates: horizon, experience and the last opponent's choice. If the results gained from the block method show that these factors are significant then it suggests that the infinite horizon was successfully induced.

A major criticism of this method is that players know when the last round of the period will be even if they do not know expressly when the last round that will count for the game will be. This limitation however is not one unique to this method but shared among all infinitely induced PD experiments as subjects understand that lab experiments cannot last forever.

The assigned probabilities for the number of rounds that would count towards the horizon works as a random termination rule assuring that the last round is unknown to the player. A limitation of this design however is that subjects could easily over or underestimate the probability of continuation.

It is important to be aware that the horizons for both finite and infinite treatments in the study increase in increments of two rounds- thus a unite increase or change in horizon in the experiment results means that the round increased by two.

3.8.5. Sessions and Order of Treatment

The finite treatment was conducted first. A session of three supergames were played, with a break between the first and second supergames. Each supergame consisted of five horizons. Each horizon had several repetitions (rounds) of the PD: with the first horizon containing 2 rounds; the second containing 4 rounds; the third containing 6 rounds; the fourth containing 8 rounds; and the fifth round containing 10 rounds.

The experiment was conducted in a university computer lab. Subjects sat spaced apart from each other behind computer terminals for the duration of the experiment. A short lunch break was provided between the first and second supergames in both treatments. The experiment and data collection was done through the computer terminal by the Z-tree program.

The instructions were verbally read to the subjects and each was given paper copies for reference during the game. After a two round practice period was played, the experiment session began. At the end of the session each players' points were totaled into a final score. This point total was then used to proportionally increase the probability each player had of doubling their earnings. The earnings were then distributed discreetly to the subjects.

The infinite treatment was conducted at a later date in a separate session with different subjects than those in the finite treatment. The infinite session consisted of three supergames. Each supergame had five periods of 10 rounds each. Each period represented a different horizon. Instructions were read and paper copies given to subjects. The players were also verbally told that the horizon increased for each subsequent period. They were explicitly made to understand that with each subsequent period of a supergame the probability that determines the number of rounds that will count for point allocation would increase and were reminded of this in the beginning of every period.

Each infinite period was matched with the appropriate horizon in the finite treatment and only the corresponding rounds were taken and counted for analysis.

3.8.6. Points and Pay

The points for each stage game were as follows:

- The payoff for mutual cooperation (CC) was 3 (R);
- the payoff for mutual defection (DD) was 1 (P);
- the payoff for a cooperating player that meets a defecting player was 0 (S);
- and a defecting player meeting a cooperating got 4 (T).

The experiment's payoff matrix is the same as that shown in Figure 3.2

At the end of the each period the players were shown the total points they attained over the rounds. The sum total of the accrued points for each supergame was shown to the player at the end of the supergame. At the end of the session the total points accrued over the 3 supergames were used to proportionally increase the likelihood of the player to increase their earnings by another 10GHC (Ghana Cedi).

The winnings were given out as follows: a random number generator was used to randomly select a number between 0 and 100. If the randomly generated number was below the total points the player accrued throughout the entire session, then they would gain the extra earnings. However if the number generated was above the player's total accrued points, then the player would not get the extra earnings. This methodology was clearly explained to the players before the game began so they understood that the outcomes of the games would directly affect their chances of gaining a higher pay-out thus creating significant motivation for subjects to maximize their point totals.

The pay-out was done at the end of the session, each player was given their total earnings (for participation in the study and/or the winnings) discreetly an invigilator.

3.9. Analysis of Results

To answer whether horizon is significant the evolution of cooperation is mapped by comparing the average cooperation as horizon and supergame increase. In order to capture any nuance lost by aggregating at the session level cooperation is also mapped at the individual supergame level. The differences between average cooperation are tested for statistical significance using the pairwise comparison of means[19]. Random Effects Probit regressions are also run to show what effect increasing the horizon[20], experience has on cooperation. The first round and last rounds are paid special attention to. In the former there is a "shadow of future", where subjects have more repetitions to play, while in the latter there is no future play and so is effectively a one shot game. We use the last partner's choice on the likelihood of a player cooperating as an additional variable for analysis and in addition to experience, we use it to test whether the Block method of inducing the infinite horizon was successful.

3.9.1. Regression and Heterogeneity Discussion

A random effects Probit regression is used to see the effect horizon has on the likelihood of a player cooperating. Unlike the ordinary Probit regression analysis of clustered data, the random effects regression does not assume that each observation is independent. There is the assumption that data within clusters are to some degree dependent. This type of regression takes advantage of the panel nature of the data by estimating and controlling for the degree of dependency resulting from the clustering of data, thus allowing the testing of relationships at the individual level. It provides the ability to estimate the effects and interactions of independent variables at either the

[19] We use the Dunnett Method which compares all means to that of the first instance
[20] A unit increase in horizon in this study means increasing the game length by two rounds of play

cluster or individual level (Hedeker, et al., 1994). This regression also controls for subject level random effects.

As the experiment uses the same subjects we expect there to be heterogeneity which needs to be addressed, a discussion regarding this type of heterogeneity can be found in the next section. The analysis uses cluster robust standard errors to address the heterogeneity in the data. The experiment conducted was under strict conditions to control for potential sources of heterogeneity, for example, the same invigilators, computer lab and even the same time of day were used to ensure consistency. The experiment also employs a large group of subjects[21] which provide many unique data points.

3.9.2. Discussion on Session Effects

As the data we rely on for the analysis stems from an economic experiment we expect there to be the session-effects problem. When subjects participate in different sessions in economic experiments, the observations across the subjects of a given session may display more correlation than those of subjects in different sessions, thus the problem is defined as a within session correlation in the variable of interest (or the residuals) once the appropriate factors are controlled for[22]. This issue results when experimenters do not identify or observable otherwise controllable salient factors. An example is in experiments conducted by multiple researchers. Session effects would occur when there are experimenter effects which they are unaware off, therefore do not control for.

[21] The experiment has 34 participants for each treatment

[22] This obviously excludes the source of the session effects, which are relevant however unknown or unobserved to the researcher

Two main problems arise in data analysis when session effects are ignored. The standard estimators will not be consistent nor unbiased if the observables and the session effects are correlated. Conversely the appropriate estimators may be consistent and unbiased, however the calculated variance of the estimate will be incorrect if session effects are ignored even if they are uncorrelated to the variables of interest. Therefore hypothesis testing is affected and complications arise in determining whether there is a treatment effect or not.

Assuming most session effects are likely to lead to within session positive correlation, it will be the case that the computed variance will be lower than the true variance of the estimator, causing the null hypothesis to be rejected too frequently (Fréchette, 2012). Despite concluding that it is difficult to find abundant situations where dynamic and static effects are large, it is nevertheless important to account for.

Fréchette (2012) makes a distinction between two types of session effects: static and dynamic session effects. While both types show a within session correlation to the variable of interest, the static session effect is created by something that is constant throughout the session and is not a function of subject behaviour during the session. While dynamic session effect is created due to factors that change throughout the session and the end result is a function of what happened during the session.

Common solutions to the session effects problem often come in one of two ways. The first is to use session averages for the variable of interest. The intuition is when session averages are treated as the unit of observation then the correlation due to the session effects will no longer present. Another posited solution to the session effects problem is not to replicate the task of interest in a session, i.e. have the subjects only play once. These solutions however do not always solve the session effects problem and come at a significant cost. Both solutions severely reduce the number of observations available for a given number of sessions and also reduce the power of the statistical tests to be performed. For example consider the case of a two treatment economic experiment with 4 sessions per treatment and 10 subjects in each session, playing 10 rounds of the stage game apiece (thus generating 800 observations but only 8 sessions). The smallest treatment effect that would lead to the null hypothesis being rejected needs to be significantly larger if the data is averaged by session than if each observation is used separately. In addition, if the behavior of interest only occurs after subjects have learned the game or practiced, then the second solution is not viable.

An easily implemented solution to various situations regarding the session effects problem is to cluster at the session level when computing the variance-covariance matrix. Although inefficient, this method provides a robust approach to hypothesis testing.

This study employs several techniques to address the session effects problem. We separate sessions into smaller subgroups which alleviates the problem by increasing the number of independent observations. We use random re-matching procedure which lessens the dynamic session effects. Clustering at the session level is employed when computing the variance-covariance matrix and we use random effects Probit regression to account for subject level heterogeneity.

3.10. Other Variables of Interest

Other variables such as experience and last player choice were also added in order to see how robust the horizon's effect is and how it moves with other factors. This analysis is conducted for both the finite and infinite treatments and the results were compared for any similarities and differences. In the next chapter the results of the experiment are presented and the findings are discussed.

4. CHAPTER FOUR:

RESULTS AND ANALYSIS

In this chapter we present the experiment results and highlight findings of interest. We begin with a general overview of how cooperation evolves in the finite and infinite treatments, analyse any differences and emphasising key findings.

4.1. Significance of Horizon and Evolution of Average Cooperation

In order to answer the question of whether the horizon has a significant effect on player cooperation we look to the differences between average cooperation for differing horizons. If game length is important we expect the average cooperation for each horizon to differ and for there to be a clear evolution of average cooperation as the horizon changes. The average cooperation aggregated by horizon is shown in Table 4.1 below:

Table 4.1 Average Cooperation Aggregated by Horizon

		Supergame			
	Treatment	1	2	3	Session
Finite	H = 2	0.078125	0.102941	0.117647	0.099571
	H = 4	0.09375	0.102941	0.110294	0.102328
	H = 6	0.09375	0.137255	0.137255	0.122753
	H = 8	0.097656	0.113971	0.139706	0.117111
	H = 10	0.090625	0.15	0.120588	0.120404
Infinite	H = 2	0.397059	0.117647***	0.117647***	0.210784
	H = 4	0.352941	0.058824***	0.117647***	0.127451
	H = 6	0.264706**	0.112745***	0.088235***	0.155229
	H = 8	0.158088***	0.110294***	0.055147***	0.107843***
	H = 10	0.129412***	0.102941***	0.055882***	0.096078***

Table 4.1 Mean Cooperation for Horizons by Supergame and for Session
Differences compared to the first instance of respective treatment (Supergame 1, H=2) for the supergame level and (H = 2) for the entire session. Significance levels ***1% **5% *10%

When aggregated by session the experiment results for the finite treatment show that the average cooperation generally tends to increase as horizon increases. Although this result mirrors the results from recent empirical literature the effect is not monotonic nor is it significant- the difference between the first horizon average cooperation of 0.10 and the final horizon average cooperation rate of 0.12 shown in Table 4.1 is not statistically significant from zero. The evidence here shows that horizon does not affect player cooperation in FRPD games. This result mirrors the findings of past game theoretic experiments that deemed the relationship between repetition and cooperation as inconclusive.

Our results show that players learn to cooperate more with experience which is unlike other recent PD experiments such as Embery, et al. (2016) that observe player cooperation in the early rounds and defection during the later rounds, with experienced players learning to defect sooner, ultimately causing average cooperation to fall with experience. For an general overview of how player cooperation evolved by round please refer to Table 4 found in Appendix I which shows the average cooperation by round for both finite and infinite sessions.

Figure 4.1 Average Cooperation per Round for the Finite Treatment

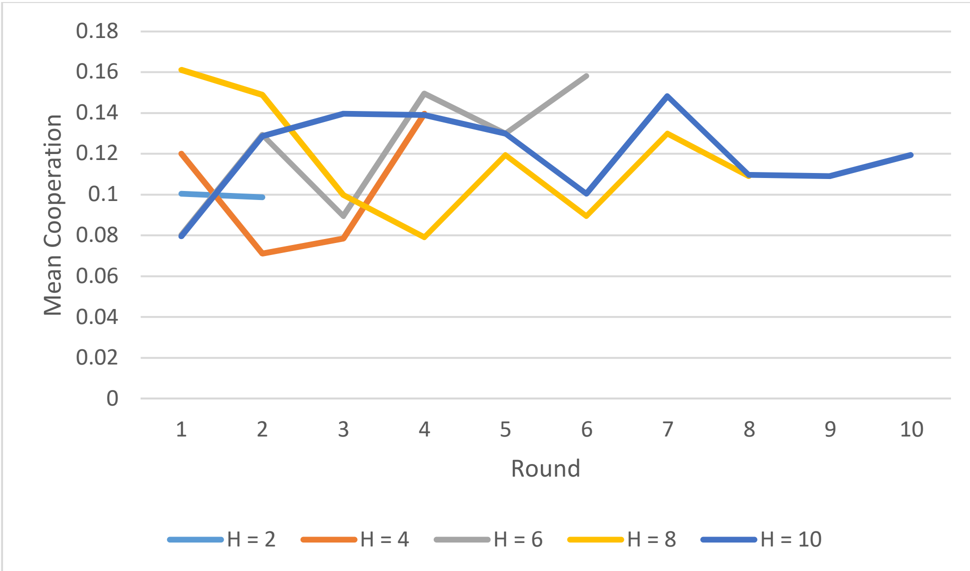

Figure 4.1 shows Average Cooperation aggregated for all supergames by round

As evidenced in other finite horizon studies we expect there to be cooperation in the early rounds then a "cut off" point were players defect near the end. As shown in Figure 4.1 above cooperation in the finite case does not follow the established evolution. In the finite treatment we find that for several last rounds of play the average cooperation actually increases from the previous rounds[23].

Contrary to recent infinite horizon studies such as Dal Bó (2005) where there is a positive relationship between horizon and cooperation, we find that under certain parameters this relationship is negative. The data in the infinite treatment shown in Table 4.1 indicate that average cooperation generally decreases when the horizon increases. Comparing the evolution of average cooperation aggregated through the entire session we find for each horizon there is a clear, general downward trend, with the last two horizons statistically different to the first. Looking further into the data we find this relationship remains even at the individual supergame level. For example in

[23] Although again this is not found to be statistically significant

Table 4.1 in the first supergame the first horizon average cooperation rate of 0.40 falls to 0.13 by the final horizon. This same trend follows in the subsequent two supergames. The decrease in average cooperation is not monotonic for the first two horizons, however becomes monotonic for all others. The differences are statistically significant for all horizons save in the first supergame where the horizon was 4 rounds.

Table 4.2 Average Cooperation by Supergame

Treatment	Supergame		
	1	2	3
Finite	0.0927083	0.1284314**	0.127451**
Infinite	0.2117647	0.1019608***	0.0745098***

Table 4.2 Differences compared to the first Supergame 1 of respective treatment
Significance levels ***1% **5% *10%

We see that with experience players cooperate more in FRPD games. Table 4.2 shows average cooperation aggregated by supergame, we see that for the finite treatment average cooperation rates increase from 0.09 in the first supergame to 0.13 in the last. The differences in average cooperation aggregated by supergame are compared are significant at the 5% level. Though this relationship is not strictly monotonic when aggregated at the supergame level when we drill down into the data in Table 4.1 we see that this effect is mostly monotonic as the average cooperation rates for each horizon increases from the first supergame to the last, save for the final horizon where average cooperation starts at 0.09 and increases to 0.15 in the second supergame, then decreases to 0.12 by the final supergame. This shows that with experience players in the finite treatment tend to cooperate more, however this effect is independent of game length.

In the infinite treatment the results in Table 4.2 show that with experience players choose to cooperate less. The average cooperation decreases monotonically across supergames, from 0.21 in the first supergame to 0.74 in the final supergame. These differences are significant to the 1% level. Not only do the average cooperation aggregated at the supergame level show that with experience players learn to cooperate less but when we drill down and look at the average cooperation for the individual horizons, we see the same result. For example, average cooperation in the first horizon of the first supergame decreases from an initial 0.397 in the first supergame to 0.118 by the final supergame. This difference constitutes a decrease in cooperation of almost 70.3% and is significant at the 1% level. Similarly, all other horizons of the infinite treatment have decreases in average cooperation rates across supergames.

Figure 4.2 Average Cooperation per Round for the Infinite Treatment

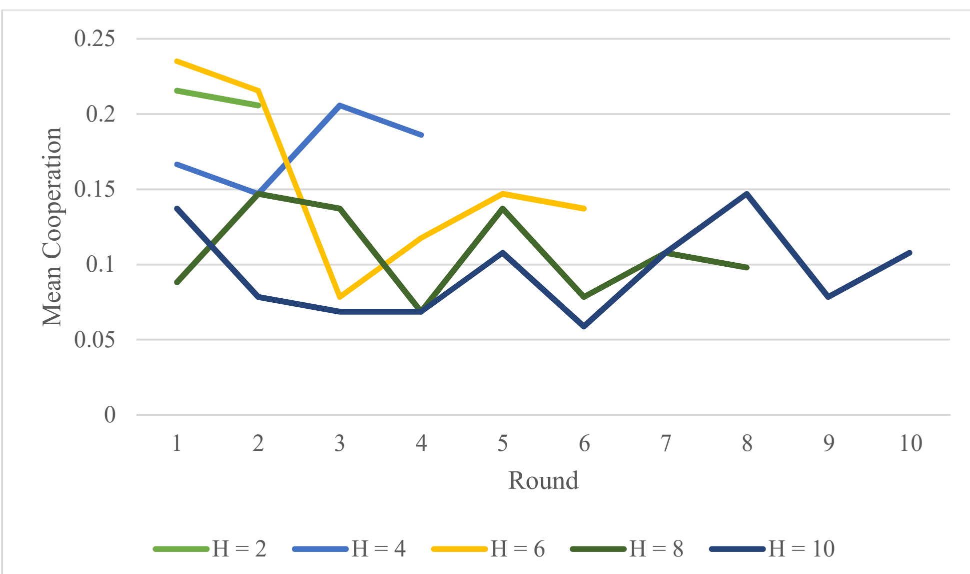

Figure 4.2 Average Cooperation per round aggregated by horizon and supergame

Opposite to what is evidence in other recent studies, Figure 4.2 shows that in the infinite treatment there are higher cooperation in shorter horizons and a general fall in cooperation as rounds increase.

When the finite and infinite treatments are compared we find that, similar to previous studies, in the finite case cooperation rates increase as the finite horizon increases, however the difference between cooperation rates was not significant.

In the infinite treatment we find the opposite result to the finite case: as the infinite horizon increases cooperation rates actually decrease. This result is divergent to the findings of other contemporary PD studies. Despite being contrary to more recent PD studies, this effect is statistically significant and echo the results found in initial PD experiments.

We also find that experience tends to mirror and perhaps magnify the relationship between average cooperation and horizon. In the finite case, we find that cooperation rates generally increase across supergames while in the infinite case the opposite is true; that cooperation rates decrease across supergames with these differences being statistically significant at the 5% and 1% level respectively.

Another interesting result is that except for the last two horizons (with 8 and 10 rounds respectively) the infinite case has higher average cooperation per horizon than in the finite case. For example when looking at the average cooperation aggregated by session for each horizon in Table 4.1 the data shows that for the first horizon the infinite case is 0.21 about twice as large as the finite case of 0.10. Similarly the infinite treatment has higher average cooperation than that of the finite treatment for the horizons with 2, 4 and 6 rounds.

Recent studies comparing the finite and infinite horizons have found similar results: that average cooperation tends to be higher in the infinite case than in the finite case, ceteris paribus. This result also mirrors theoretical predictions as the infinite horizon can support more cooperative equilibria. When comparing the evolution of average cooperation per horizon in each supergame, we find that the number of instances where the infinite treatment has higher average cooperation than the finite treatment are concentrated in the first supergame. However due to the decreases in average cooperation in the infinite treatment by the final supergame the finite treatment has a greater number of instances where average cooperation is higher than in the infinite treatment. Nevertheless throughout the entire session the frequency of instances where one treatment has higher cooperation by horizon than the other is fairly even: with both infinite and finite treatments having seven cases were the average cooperation is greater than the other.

This finding shows that the decrease in average cooperation as the horizon increases in the infinite treatment is quite substantial, so much so that when aggregated by supergame the average cooperation in the infinite treatment are only higher than in the finite case for the first supergame. As shown on Table 4.2 in the subsequent two supergames the aggregated average cooperation in the infinite treatment falls below that of the finite treatment.

These results highlight the importance of the experiment parameters when conducting game theoretic empirical studies.

4.2. First Round Analysis

We look at the first round cooperation rates of the finite case to see if game length, has any effect on initial cooperation for that horizon.

Conventional knowledge lends the expectation that longer horizons have higher cooperation rates in the first round of play. Players choose to cooperate rather than defect in the early rounds in order to foster more cooperative play in subsequent rounds. We anticipate that the evolution of cooperative play will follow this trend and we will see cooperation in early rounds and high levels of cooperation in the first round of each period. Longer horizons creates an incentive to cooperate early and we expect as the horizon increases the first round cooperation rates to increase as well. A significant caveat is that because this study employs a matching procedure that randomizes a player's partner after each stage game, the incentives to cooperate in this case are not caused by the significant threat of a player's partner choosing to defect for subsequent rounds of the period in response to the player defection, but by the general threat that the collective pool of players will choose to defect if they meet early defection from others. In effect it is dependent on subjects realizing that early defections can cause the pool of subjects to defect early and balancing that threat with the belief that they can be paired with cooperating and/or forgiving players in subsequent rounds.

Table 4.3 First Round Cooperation

		Supergame			
	Horizon	1	2	3	Session
Finite	2	0.125	0.117647	0.058824	0.10049
	4	0.125	0.058824	0.176471	0.120098
	6	0.09375	0.117647	0.029412	0.08027
	8	0.21875	0.088235	0.176471	0.161152
	10	0.0625	0.088235	0.088235	0.079657
Infinite	2	0.441177	0.117647***	0.088235***	0.215686
	4	0.352941	0.058824***	0.088235***	0.166667
	6	0.294118	0.205882*	0.147059***	0.235294
	8	0.117647***	0.058824***	0.088235***	0.088235*
	10	0.235294	0.117647***	0.058824***	0.137255

Table 4.3 Round 1 Average Cooperation Aggregated by Horizon for each Supergame and Session
Differences compared to the first instance of respective treatment (Supergame 1, H=2) for the supergame level and (H = 2) for the entire session
Significance levels ***1% **5% *10%

The results show that in the finite case the traditional evolution of high early cooperation followed by unravelling in later rounds is largely non-existent. This finding is interesting as it illustrates that under these parameters when the number of future rounds increases (thus increasing the gains from cooperating early) cooperation still remains sparse, regardless of experience, which gives players the opportunity to learn to cooperate early. Contrary to the expectation of seeing the highest levels of cooperation in the first round of the longest horizon than those of shorter horizons, we find no statistically significant difference. Thus showing no difference between average cooperation with differing horizons.

The observations shown in Table 4.3 where average cooperation is aggregated by horizon, the instances of one treatment having higher cooperation rates than the other are fairly even between the two cases: with the infinite treatment having 8 instances of higher average cooperation than the finite treatment, the finite having 6 instances of higher cooperation than the infinite treatment and 1 case where they have equal average cooperation. Again it is seen that in the infinite case the majority of these occurrences are concentrated in the first supergame, while the in the finite case they are spread throughout the last two supergames. Nevertheless, when aggregated for the entire session we see that the infinite treatment has higher average cooperation than in the finite treatment for all horizons save for the 8 round horizon.

4.3. Marginal Effects Analysis

We use a random effects Probit regression to answer the question regarding how horizon affects cooperation. This relationship is also viewed under the lens of experience. Rather than reporting the random effects Probit coefficients or we show the marginal effect. Several studies have shown that reporting the change in Pr(Y=1) or the marginal effect is the superior means of reporting the relationship between the dependent and the independent variable. When interpreting the results one must bear in mind that a unit change of horizon in this study refers to a two round increase of the game length while a unit increase in the supergame is the addition of another repetition of all the periods that constitutes a supergame.

Table 4.4 Marginal Effects of Horizon and Experience on Likelihood of Cooperating in the First Round

	Marginal Effect	P-Value	Cluster Robust SE	
Finite Horizon	-0.0001888	0.975	0.0540668	Session
Finite Supergame	-0.006944	0.719	0.1758698	Session
Infinite Horizon	-0.0225237	0.021	0.0433196	Session
Infinite Supergame	-0.100156	0.001	0.127999	Session

Table 4.4 Random effects Probit regression of Increasing Horizon and Experience on Likelihood of Cooperation in the First Round

The evidence showing there is no discernable difference in cooperation caused by changes in the horizon continues to mount in FRPD games. The marginal effects from the random effects Probit analysis in Table 4.4 shows the likelihood of a player choosing to cooperate in the first round for the finite treatment given an increase in horizon is not statistically significant, interestingly experience has no effect either. Collectively the data suggests that even allowing for experience players will disregard horizon as a variable when choosing to whether to cooperate in the first round. The experiment results so far heavily suggests that under these parameters, horizon has no effect on cooperation in finitely repeated games.

The infinite treatment first round average cooperation generally decreases when the horizon increases. In the first supergame the first horizon average cooperation falls from 0.44 to 0.11 by the fourth horizon (H = 8). This difference in average cooperation is significant to the 1% level.

Similarly the random effects Probit analysis of a player's likelihood to cooperate given a two round increase in the horizon shows players are less likely to choose to cooperate by about 2.3%. This inverse relationship is consistent with the decreasing evolution of cooperation rates in the infinite treatment and is significant at the 5% level.

The two treatments differ significantly as the first round results show that in the finite case horizon does not have any significant effect on average cooperation rates while in the infinite treatment the opposite is true. These results are consistent with the analysis of the results found in Table 4.1 and Table 4.2, which both draw the same conclusion.

These results again diverge from other recent studies as they find when horizon increases cooperation rates to increase as well, Dal Bó (2005) for example finds that first round cooperation rates increase with higher probability of continuation in the infinite case. When the two treatments are compared to each other[24], we find that the first round cooperation levels of the finite and infinite treatments are statistically different from each other at the session level as well as in the first, third and fifth horizon.

4.4. Last Round Analysis

We study the last stage game played for each horizon, as there is no future play after this round making them one shot PD games. In this way there is no incentive to cooperate as there are no payoff benefits to cooperating nor is there any fear of future punishment. Theoretically players

[24] We us a Two Sample Unpaired Proportions Test to compare the two treatments. Results show that the two treatments are different at the 5% level of significance with the infinite treatment being 0.061 larger than the average cooperation in the finite treatment.

should defect as it becomes the dominant strategy in this round; choosing to cooperate becomes the dominated strategy.

Table 4.5 Average Cooperation of Last Round Cooperation

	Horizon	Supergame			Session
		1	2	3	
Finite	H = 2	0.03125	0.088235	0.176471	0.098651961
	H = 4	0.125	0.205882	0.088235	0.139705882
	H = 6	0.0625	0.176471	0.235294	0.158088235
	H = 8	0.0625	0.058824	0.205882	0.109068627
	H = 10	0.09375	0.117647	0.147059	0.119485294
Infinite	H = 2	0.352941	0.117647*	0.147059	0.205882353
	H = 4	0.323529	0.117647*	0.117647	0.18627451
	H = 6	0.264706	0.088235**	0.058824***	0.137254902
	H = 8	0.117647*	0.117647*	0.058824***	0.098039216*
	H = 10	0.058824***	0.176471	0.088235**	0.107843137
True Infinite	H = 2	0.411765	0.088235***	0***	0.166667
	H = 4	0.264706	0.088235***	0.029412***	0.127451
	H = 6	0.205882**	0.029412***	0.058824***	0.098039
	H = 8	0.088235***	0***	0.058824***	0.04902*
	H = 10	0.058824***	0.176471**	0.088235***	0.107843

Table 4.5 Mean Cooperation for Last Round of Each Treatment by Supergame and Session
Aggregated for all supergames, differences of each horizon compared to the respective H = 2
Paired Means Comparison done for Infinite cases- Comparison and True
Significance levels ***1% **5% *10%

In the finite treatment we see that the last round differences in average cooperation for each horizon are not statistically different from zero. This result is consistent for the average cooperation aggregated from entire session and when we drill down into the individual supergames (shown in Table 4.5).

Despite the relatively high variance between the cooperation rates, for example in the first supergame the second horizon average cooperation of 0.13 is over four times larger than that of the first horizon average cooperation which it was compared, as there is no statistically significant difference in average cooperation between the last rounds the data suggests that the last rounds for this treatment are indeed treated as one shot games[25].

Table 4.6 Marginal effects of Horizon and Experience Random Effects Probit Regression showing Likelihood of cooperating in the Last Round

	Marginal Effects	P-Value	Cluster Robust SE	
Finite Horizon	0.0000521	0.995	0.055041	Session
Finite Supergame	0.0407148	0.108	0.1719043	Session
Comparison Infinite Horizon	-0.0266189	0.003	0.049076	Session
Comparison Infinite Supergame	-0.0649626	0.013	0.1450352	Session
True Infinite Horizon	-0.0161008	0.088	0.0611248	Session
True Infinite Supergame	-0.0781546	0.006	0.1847321	Session

Table 4.6 Marginal Effects of random effects Probit regression of horizon and supergame on likelihood of Cooperation in the last round

Indeed the Table 4.6 the random effects Probit result for the effect of the horizon on the likelihood of a player choosing to cooperate in the finite treatment reaffirms that there is no statistically significant effect on the last round cooperation rate caused by changing the horizon. This finding reiterates that despite differing game length, the last round of each horizon is treated as a one shot PD game.

[25] Again the two sample unpaired test of proportions is employed. Both the infinite cases are not statistically different from the last rounds of the finite treatment according to the test of proportions.

In the infinite case the true last round of the game is different from the last round that was counted for point allocation and analysis for each horizon. The cooperation rates labelled "Comparison" are for the last rounds of the infinite treatment chosen for analysis, which correspond to the last round of each horizon (2nd, 4th, 6th, 8th and 10th respective rounds) of the finite treatment. The cooperation rates labelled "True" are for the last round (10th) of each horizon of the infinite treatment the subjects played. The cooperation rates labelled "True" are the last rounds where subjects clearly understood that there is no longer any future play beyond that round for that particular horizon.

In both infinite cases; both the Comparison and "True", there is a mostly clear decline in average cooperation as horizon increases (see Figure 4.3). The Comparison case however tends to have higher average cooperation than that of the "True" case.

Figure 4.3 Last Round Cooperation for Each Treatment

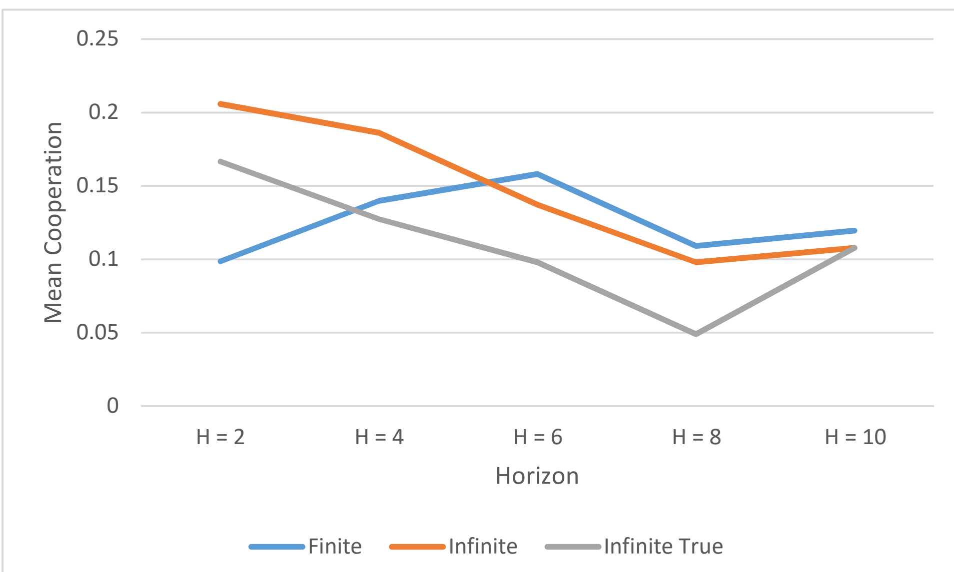

Figure 4.3 The average cooperation aggregated by horizon and supergame for the last round of each case

Except for the last round of the final horizon, where the Comparison and True infinite cases have the same last round, Table 4.5 and Figure 4.3 show that in all other rounds there is higher average cooperation in the Comparison last round case than in the True last round case.

The results also show a general decrease from the first to the fourth horizon for both infinite cases. However the differences between the average cooperation in the last game for both cases are significant only between the first horizon and the fourth horizon, the average cooperation for the last game of the other horizons do not differ from the first in a statistically significantly manner. This result shows that in the event subjects know the exact round the game will end, the round's outcome will be affected- as this last round will have lower cooperation levels.

As shown in Table 4.5, the differences between the average cooperation in the infinite last rounds for individual supergames are statistically significant, with near all differences of the True case being significant at the 1% level.

Interestingly the evolution of average cooperation is fairly sporadic, for example in the first supergame of the True case, as horizon increases there is a monotonic fall in average cooperation rates, while in the last supergame the opposite is observed and there is a general increase in average cooperation. In the Comparison case the differences in average cooperation between horizons are statistically insignificant for the first supergame (when each horizon's average cooperation rate is compared to that of the first horizon in the first supergame), however by the last supergame the differences in average cooperation become statistically significant and form a clear downward trend.

For the infinite Comparison case the random effects Probit results in Table 4 6 show that there is a statistically significant relationship between horizon and with players choosing to cooperate. It shows that as horizon increases, the likelihood of a player choosing to cooperate falls by 2.7%.

This negative relationship is mirrored by the statistically significant relationship between experience and player cooperation. In effect the inverse relationship of horizon and player cooperation is magnified with experience, showing that with longer games and successive supergames players learn to defect in the last game and the likelihood of players choosing to cooperate falls by 6.5%.

In the True case the experience also has a significant negative effect on the probability of subjects cooperating, however changes in horizon do not have a statistically significant effect on the choice to cooperate.

Directly comparing the finite and infinite treatments at the session level we largely find no significant differences in average cooperation in the last round for both treatments. This similarity becomes more nuanced when we look further into the data found in Table 4.5 and find significant differences in average cooperation in the infinite treatment (especially that of the True case) but not the finite treatment. In the infinite case we find that increases in horizon negatively affect subjects' likelihood to cooperate, however in the finite and True case this relationship is not significant.

Only in infinite treatments did experience factored into the subjects' decision and affected the likelihood of the player choosing to cooperate in the last round.

In the finite case there is no statistically significant relationship between horizon and the likelihood of players cooperating while in both infinite treatments there is a negative statistically significant relationship. In the True infinite case the experience has a larger effect than that of the Comparison infinite case. Theoretically horizon should not factor into the decision making process when it comes to the last round of play. Theory predicts players will choose to defect in the last round as there is no incentive to further cooperate. Our results show that in the cases where players are aware they are in the final round of the game or period (in the finite treatment or the True infinite case) the horizon does not factor into their choice to cooperate.

Figure 4.4 showing the First round versus the Last round Cooperation for the Infinite Treatment

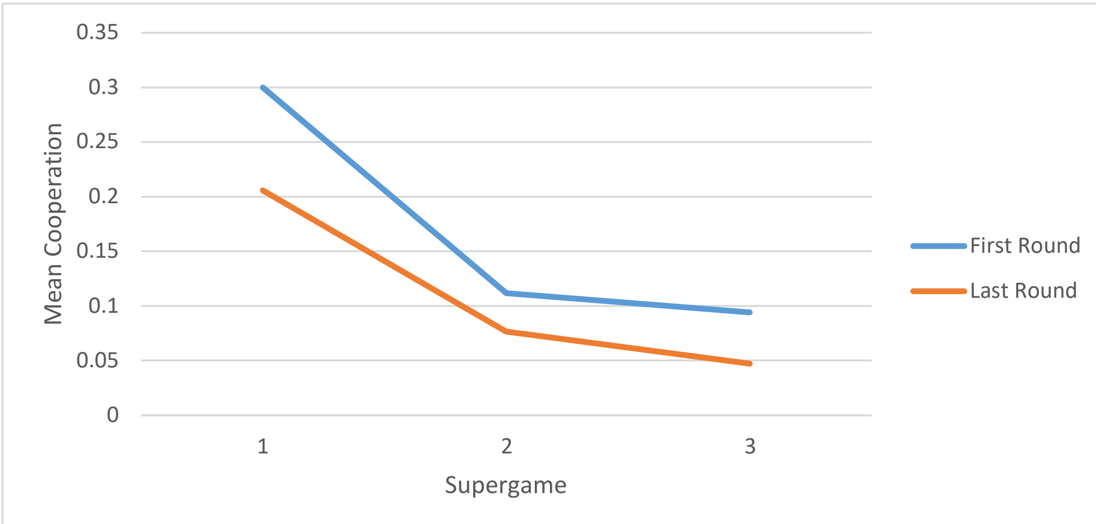

Figure 4.4 shows the Average Cooperation for the First and the Last Rounds Aggregated by Supergame for the Infinite treatment

We compare the first round average cooperation to the last round cooperation rates to see if there are any significant differences between the two. As there is no future in the last round of every period it plays as a one shot game. We expect first round cooperation levels to be higher than those found in the last round, if there are no significant differences between the cooperation levels in the two rounds then it suggests that players treat the first round as a one shot game. This case promotes the idea that under certain circumstances players will disregard game length and will treat varying horizons as a series of one shot games. In Figure 4.4 above we see that the first round of the infinite treatment is higher than that of the true last round of the infinite case, showing that for the infinite treatment earlier rounds were not seen as one shot PD games.

Figure 4.5 First Round versus Last Round Cooperation

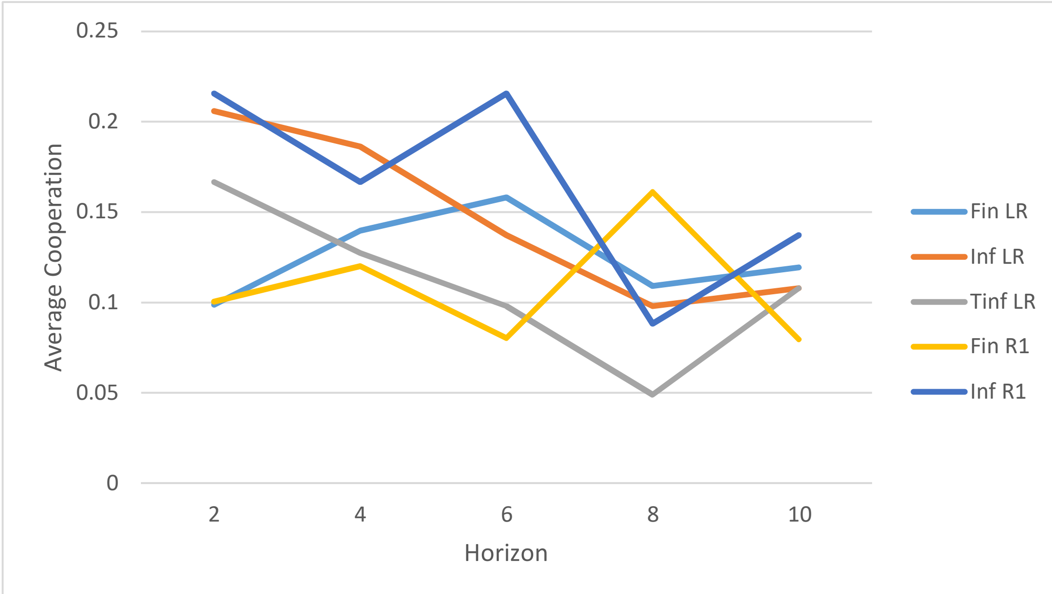

Figure 4.5 First Round Cooperation versus Last Round Cooperation for All Cases
Fin LR – Finite Last Round; Inf LR – Infinite Comparison Last Round; Tinf LR – True Infinite Last Round; Fin R1 – Finite First Round; Inf R1 – Infinite First Round

The first round and last round average cooperation rates are tested for statistically significant differences using the two sample test of proportions. We find there are significant differences in the first, third and fifth horizon cooperation levels of the first round, however for the last round cooperation levels we find no statistically significant difference between the two treatments.

4.5. Marginal Effects of Key Variables on Likelihood of Cooperating

We look deeper into the marginal effects of horizon, experience and last player choice on influencing player cooperation.

Table 4.7 shows the multivariate random effects Probit results for the effect game length, experience and the last partner's choice on the likelihood of players choosing to cooperate. For the finite treatment none of these variables affects the likelihood of a player deciding to cooperate while in the infinite case all three variables are significant.

Table 4.7 Marginal effects of Horizon, Experience and Last Partner's Choice on likelihood of Cooperating

Variable	Marginal Effect	P-Value	Cluster Robust S.E.	
Finite Horizon	0.0027116	0.347	0.0345242	Session
Finite Supergame	0.0124564	0.416	0.1293066	Session
Finite Last Partner Choice	0.0191776	0.143	0.1103929	Session
Infinite Horizon	-0.0243431	0.000	0.0263784	Session
Infinite Supergame	-0.0630449	0.002	0.1195382	Session
Infinite Last Partner Choice	0.0647094	0.000	0.0831857	Session

Table 4.7 shows the marginal effects of the variables in the multivariate random effects Probit on the likelihood of choosing to cooperate

We analyse the results for each individual variable more detail in the section below. We further use the random effects Probit regression to estimate the relationship between game length and cooperation for both treatments. Table 4.8 shows the results of this analysis.

Table 4.8 Marginal Effect of Increasing Horizon by Two Rounds on likelihood of Cooperating

Variable	Marginal Effects	P-Value	Cluster Robust S.E.	
Finite Horizon	0.0027116	0.316	0.0345242	Session
Infinite Horizon	-0.0243431	0.000	0.0263784	Session
Finite Horizon	0.0001991	0.879	0.0552621	Supergame 1
Finite Horizon	0.0060892	0.186	0.086685	Supergame 2
Finite Horizon	0.0006756	0.857	0.0609885	Supergame 3
Infinite Horizon	-0.06077	0.000	0.0554114	Supergame 1
Infinite Horizon	0.0027183	0.604	0.0526094	Supergame 2
Infinite Horizon	-0.0090124	0.000	0.0534387	Supergame 3

Table 4.8 Random Effects Probit Regression Results showing the marginal effect of increasing the treatment horizon by two rounds on the likelihood of a player to cooperate

In the finite case we see that the as game length increases the likelihood of subjects cooperating does not change in any statistically significant manner. This result remains true when aggregated by the entire finite session and for each individual supergame, showing that experience does not change this result: changes in horizon do not affect the player choosing to cooperate.

The random effects Probit regression of the infinite treatment analysis shows that changes in game length does in fact alter the subjects' choice in a statistically significant manner. The results show a negative relationship between increases in horizon and the likelihood of choosing to cooperate. This finding is unexpected as the literature establishes that the relationship should be positive due to the "shadow of future". We see that both when aggregated by session and by supergame this relationship is statistically significant at the 1% level. At the session level we see that increasing the length of a game by two rounds will decrease the likelihood of a player choosing to cooperate by 2.43%.

We see that experience plays an important variable when it comes to a player's decision making in the infinite treatment, an increase in supergame lowers the likelihood of a player cooperating by 6.3%. The random effects Probit analysis in Table 4.9 shows the effect increasing the supergame will have on subjects' likelihood to cooperate.

Table 4.9 Marginal Effect of Experience on the Likelihood of Cooperating

Variable	Marginal Effect	P-Value	Cluster Robust S.E.	
Finite Supergame	0.0128269	0.406	0.1290562	Session
Infinite Supergame	-0.0630449	0.002	0.1195382	Session

Table 4.9 the random effects Probit regression of increasing the supergame on the likelihood of cooperating

The finite case shows the effect of increasing the supergame is statistically insignificant. This means that subjects do not become more or less likely to cooperate through repeated play. This result infers that there are no significant learning effects through gaining of experience under these parameters or gives credence to past studies that find learning effects severely hindered in Stranger settings.

Experience in the infinite treatment is statistically significant however the coefficient is negative, showing that increasing the supergame decreases the likelihood of a player choosing to cooperate. This is an important finding as it shows that with experience the horizon's negative effect on the likelihood of cooperating is magnified. Previous studies have shown that experience indeed magnifies the horizon's effect on cooperation, however unlike these previous studies that find this relationship to be positive, our results show the relationship of horizon and cooperation to be negative.

The two treatments are different in that experience is statistically significant only in the infinite case where experience has a negative relationship with player cooperation. Similarly to the previous result the game length does not affect a player's likelihood of cooperating in the finite treatment while in the infinite treatment increases in horizon diminishes the likelihood of subjects cooperating, thus infinite treatment finds that a player's likelihood of cooperating will fall if the number of rounds in the game increases under the Stranger setting.

Surprisingly the results show the last partner choice has a significant effect in the infinite treatment. Although the threat of retaliation and defection in future rounds would help keep subjects cooperating subjects understood that after each round of play players are re-matched, making the last player's choice have little bearing on the current round. This suggests that in the infinite case subjects "pay it forward" and are open to creating a cooperative rapport in the pool of players, which could induce higher cooperation levels for the game. Given the decreasing nature of cooperation in the infinite treatment of this study this result is uplifting as it indicates that with enough cooperative subjects being paired there is a chance that cooperation proliferates. This idea of reciprocity explains the low and sharply decreasing nature of cooperation. Cooperating players that meet early defections would likely defect in the next rounds as punishment. Thus negative reciprocity and contagious punishment strategy will have a negative effect on cooperation.

Table 4.10 of the Marginal Effect of the Last Opponent's Choice on the Likelihood of Cooperating

Variable	Marginal Effects	P-Value	Cluster Robust S.E.	
Finite Last Partner Choice	0.0191776	0.143	0.1103929	Session
Infinite Last Partner Choice	0.0647094	0.000	0.0831857	Session
Finite Last Partner Choice	0.0161195	0.202	0.2125302	Supergame 1
Finite Last Partner Choice	0.0195452	0.254	0.1970166	Supergame 2
Finite Last Partner Choice	-0.0012154	0.941	0.1996063	Supergame 3
Infinite Last Partner Choice	0.0784107	0.001	0.101966	Supergame 1
Infinite Last Partner Choice	-0.0063117	0.630	0.1415763	Supergame 2
Infinite Last Partner Choice	0.0385445	0.000	0.2319763	Supergame 3

Table 4.10 showing marginal effects from random effects Probit analysis of the last partner's choice on the likelihood of cooperating

The last partner's choice does not have any statistical significance in the likelihood of a player choosing to cooperate in the current round in the finite treatment. This means that the outcome from the last game does not significantly affect the player's decision in their current round of play. This result is shown both throughout the finite session and at the supergame level, with the last player's choice only factoring into the decision in the first supergame. This is important as it suggests that players are set in their strategies or choose not to let past outcomes interfere with the present round.

In the infinite treatment the outcome of the last game does in fact have a significant effect on the likelihood of the player choosing to cooperate in the present round. The marginal effect is positive showing that if the last player the subject was paired with cooperated then the subject is 6.47% more likely to cooperate in their present round of play. Table 4.10 shows that the effect is significant at the session level and the individual supergame level, with the last partner's choice not effecting the likelihood of subjects cooperating only the second supergame. The significance of this variable under these parameters show that it is robust to changes in parameters and independent to different ways of creating the infinite horizon.

We find that the finite and infinite treatments differ when it comes to how the outcome of the previous round affected the likelihood of cooperation. While in the finite treatment the last partner's choice did not affect the likelihood of cooperating in the current round of play, the infinite treatment showed that the past cooperative outcomes spill over into the present round of play. These results are further contextualized and discussed in the next section of the chapter.

4.6. FURTHER DISCUSSION OF RESULTS AND ANALYSIS

We further continue the discussion of the results from Chapter Four. Here we focus more explicitly on the research questions raised in Chapter One.

4.6.1. Does Game Length Matter?

The results show there is no statistically significant difference between cooperation rates of differing horizon in finitely repeated PD games, neither does increasing the horizon by two rounds have any statistically significant effect on the likelihood on subjects choosing to cooperate. This result is robust even when subjects are allowed to gain experience and learn. The evidence shows that despite playing games of differing lengths subjects effectively treated the various finite horizons as a series of one shot games. This idea is further given credence when we study cooperation rates in the first round of each horizon. It is expected that players will be most likely to cooperate at this stage of the game, however as horizon increases the likelihood of subjects cooperating in the first round does not increase. This effect is independent of experience, suggesting that subjects do not learn to cooperate in the first round. Finally, because the last round of the finite horizon is treated as a one shot PD game as there is no future play, when first round cooperation is compared to that of the last round of play for each horizon we see no statistical difference between them. These consistent results shows that under the experiment's parameters, horizon has no effect on cooperation in finitely repeated PD games. Rather than accounting for the varying horizons, the subjects treat the game as a series of one shot PD games. This result may initially seem contrary to decades of established research that has found high levels of cooperation in the early rounds of finitely repeated games followed by unravelling near the end of the game however the result completely rational. Most studies in the literature pair subjects for the entire supergame, while this study paired subjects only for a round. This means that the continuing strategic relationship (discussed in Chapter Three) that enables reputation building and punishment is a necessary condition for the horizon to have an effect on cooperation in the finitely repeated games. These results suggest that it is not game length in and of itself that fosters increased cooperation in finite horizons: it is an enabling part of the toolkit subjects can use to cooperate the other part being repeated interaction with the same partner. Horizon essentially relies on the creation of reputations and threat of punishment to affect cooperation, without these two factors horizon's effect on cooperation is negligible. Further study is needed to test the robustness of this result and see under what conditions, if any, horizon can affect cooperation in a Stranger setting.

The results show that despite the existence of a shadow of future in the infinite treatment, as horizons increase cooperation falls. Although this effect is not necessarily monotonic it is statistically significant and substantial. The data suggests that horizon has a stronger effect on earlier supergames than the later ones. We also find that increases in horizon cause the likelihood of subjects cooperating in the first round to fall. This result is consistent with cooperation rates falling as horizon increases. This outcome is the opposite of what is found in recent literature. Much like what was observed in the finite treatment, this result illustrates how dramatic the changes in the evolution of cooperation can be in experiments with differing parameters.

Due to the method of inducing the infinite time horizon the infinite treatment had two last rounds for each horizon: the last round that counted towards the analysis for each horizon (Comparison) and the "true" final round played for the set of games from which each horizon was taken (True). In the Comparison case horizon matters when subjects are choosing whether to cooperate or not, while in the True case it does not. This result makes sense as in the Comparison case, the rounds correspond to earlier rounds in the period (where there is future play) while in the true case there are no rounds of play afterwards. This presents a key limitation of using the Block method to induce the infinite horizon, that the last round of play for the period is treated as a one shot game. This finding suggests that when using the Block method experimenters must be wary of when they set the last round and use caution when using the method in general for analysis. Nevertheless results show that there is indeed a shadow of future as there is a significant difference between cooperation in earlier rounds of play and the last round of play. In both cases there was a general trend for cooperation rates to decrease with horizon with the True case having lower cooperation than the Comparison.

4.6.2. Does the Block Method Induce an Infinite Horizon?

The block method seems to have been successful in inducing the infinite horizon in a laboratory setting. All three variables, horizon, experience and last partner choice were found to be significant. The horizon and experience having the same direction is also consistent with the literature on the subject that suggests in infinitely repeated games experience magnifies the horizons effect on cooperation.

The last partner choice being significant and positively correlated with cooperation is also consistent with game theoretic literature. We also find cooperation in the infinite treatment to be higher than in the finite treatment and when the cooperation in the final horizons of the finite and infinite treatments are compared, there are statistically significant differences. These outcomes suggest that the block method does in fact induce an infinite horizon in the laboratory setting, nevertheless questions remain regarding the robustness of the results gained. A study testing comparing the results gained from the random termination method and the block method could shed more light into the robustness and limitations of this novel method of inducing an infinite horizon in the lab.

4.6.3. Is the Effect Constant Across Supergames as Subjects Gain Experience?

Experience was found to have no statistically significant effect on cooperation in the finite treatment, showing that through repeated play subjects did not learn to cooperate any more or less. Game theoretic literature shows high levels of cooperation in FRPD games followed by a collapse of cooperation in the last rounds, with experience this collapse occurs earlier. In our study this pattern of cooperation was absent, the evidence suggests that players treated the game as a series of one shot games rather than finitely repeated games of differing horizons. Further research is needed to understand how experience and horizon can affect cooperation under differing parameters.

The negative relationship between experience and average cooperation mirrors the negative relationship between horizon and average cooperation in the infinitely repeated treatment. Consistent with the literature, experience is shown to magnify the horizon's effect, however in this case it lowers cooperation. The study does not capture whether the relationship between experience and cooperation as a general rule mirrors that of the horizon, however as repetition underlies both these variables we surmise that they affect cooperation in the same direction.

Experience in both last round infinite cases experience did play a significant role in a player's decision to cooperate. In the True case experience has a much stronger effect than that in the Comparison case.

4.7. Notable Findings and Other Variables

Average cooperation in the finite and infinite cases are statistically different from each other, with the former initially having lower cooperation. This result is consistent with recent literature such as Dal Bó (2005). However the data shows that this result is not absolute, but more nuanced. With repeated play the infinite treatment cooperation rates decrease so much so that the average cooperation in the second and final supergames are lower than those of the finite treatment.

Literature on infinitely repeated games have shown that players matched with a cooperating partner in the previous round are more likely to cooperate in their current round. In the infinite treatment subjects were more likely to cooperate if their previous partner chose to cooperate in the preceding round, despite their current partner being different. This result is robust as it holds even under this experiment's parameters, showing it is independent of reputation building and directed punishments. A likely reason for this effect is positive reciprocity, where cooperation in one the previous round is rewarded with cooperation in the next as well as the contagious punishment strategy.

The experiment parameters employed in this study mean subjects can only rely on the promise of "general" future play (based on their beliefs concerning the remaining rounds to be played), the outcomes of their previous rounds and their beliefs concerning the possibility of partner altruism to form their strategy. Subjects can be compelled to cooperate as a whole in order to avoid the case where defection proliferates throughout the entire subject pool, although the relatively low cooperation rates in data suggest subjects did not realize this while playing the game.

5. CHAPTER FIVE:

CONCLUSION

This study served to test and explore the relationship between horizon and cooperation in a novel way. An experiment of unique parameters was conducted in order to explore how cooperation evolved with horizon. The study found that average cooperation does not always behave as the literature suggests. In finitely repeated games the pattern of early cooperation followed by a collapse of cooperation rates near the end of the game does not always hold. The data finds that under a Stranger setting subjects will in fact disregard the horizon of finitely repeated PD games and treat the game as a series of one shot PD games. The study found that experience is not a significant factor when it comes to affecting the evolution of cooperation rates and allowing for it did not change horizon's insignificant effect.

Consistent with recent literature on infinitely repeated games horizon and experience do have a significant effect on cooperation. Unlike the wave of contemporary studies brought about by Dal Bó (2005) that find this relationship to be a positive one, we find that longer horizons have a negative impact on cooperation. This effect is magnified through repeated play.

Despite the fact that subjects changed partner after every round the previous partner's choice did affect the likelihood of a player cooperating in their current round of play in the infinite treatment. We attribute this effect to reciprocity as players will "pay it forward" even though the partner benefiting is different. The low level and decreasing nature of cooperation also suggests the opposite is true, as players who suffered from early defections will defect against others in subsequent rounds, thus spreading the punishment throughout the pool of players.

Parameters of the game are very important in determining player behaviour and how cooperation rates evolve. The experiment highlights that an integral parametric factor for the horizon to have an effect on player cooperation is a matching procedure that is able to take advantage of repetition, specifically pairing players in a manner that creates a continued strategic relationship between them, thus enabling reputation building and direct punishment of defection.

This study shows that in infinitely repeated games, parameters can significantly alter the relationships between factors, as evidenced by the negative relationship between horizon and player cooperation. More research is needed to discover how key variables and relationships are affected under differing conditions and parameters. The interesting question regarding this area of study is determining what mechanisms and rationales are behind why these relationships change and either positively or negatively affect cooperation.

This study highlights the need for more empirical examination regarding cooperation under differing conditions. Further investigation is needed on how these factors such as horizon, experience and cooperation interact under differing parameters.

REFERENCES

Andreoni, J., 1988. Why free ride? Strategies and learning in public goods experiments. *Journal of Public Economics,* pp. 291-304.

Andreoni, J. & Miller, J., 1993. Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence. *The Economic Journal,* p. 570–585.

Aumann, R. & Shapley, L., 1976. Long Term Competition: A Game Theoretic Analysis. *Hebrew University.*

Benoit, J.-P. & Krishna, V., 1985. Finitely Repeated Games. *Econometrica,* 53(4).

Bereby-Meyer, Y. & Roth, A., 2006. The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation. *American Economic Review,* p. 1029–1042.

Bernard, J. & Yuksel, S., 2014. The benefits of being able to fire (and replace) people: An experimental study. *Working Paper.*

Boone, C., Brabander, B. D. & van Witteloostuijn, A., 1996. The impact of personality on behavior in five prisoner's dilemma games. *Journal of Economic Psychology,* p. 343–377.

Camera, G. & Casari, M., 2009. Cooperation among strangers under the shadow of the future. *American Economic Review,* p. 979–1005.

Cooper, R., Dejong, D. V., Forsythe, R. & Ross, T. W., 1996. Cooperation Without Reputation: Experimental Evidence From Prisoner's Dilemma Games. *Games and Economic Behavior,* p. 187–218.

Dal Bó, P., 2005. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *American Economics Review,* p. 1591–1604.

Dal Bó, P. & Fréchette, G., 2011. The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review,* pp. 411-429.

Dal Bó, P. & Fréchette, G., 2016. On the Determinants of Cooperation in Infinitely Repeated Games: A Survey. *Journal of Economic Literature,* pp. 60-114.

Dreber, A., Fudenberg, D. & Rand, D. G., 2014. Who cooperates in repeated games: The role of altruism, inequality aversion, and demographics. *Journal of Economic Behavior & Organization,* pp. 41-55.

Ellison, G., 1994. Cooperation in the Prisoner's Dilemma with Anonymous Random Matching. *The Review of Economic Studies,* pp. 567-588.

Embrey, M., Fréchette, G. R. & Stacchetti, E., 2013. An experimental study of imperfect public monitoring: Renegotiation proofness vs efficiency. *Working Paper.*

Embrey, M., Fréchette, G. R. & Yuksel, S., 2016. Cooperation in the finitely repeated prisoner's dilemma. *Working Paper.*

Engle-Warnick, J. & Slonim, R. L., 2004. The Evolution of Strategies in a Repeated Trust Game. *Journal of Economic Behavior & Organization,* pp. 553-573.

Engle-Warnick, J. & Slonim, R. L., 2006a. Learning to trust in indefinitely repeated games. *Games and Economic Behavior,* pp. 95-114.

Feinberg, R. & Husted, T., 1993. An experimental test of discount rate effects on collusive behaviour in duopoly markets. *Journal of Industrial Economics,* pp. 153-160.

Fischbacher, U., 2007. z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics,* pp. 171-178.

Flood, M. M., 1952. *Some experimental games,* Santa Monica: Research Memorandum RM-789.

Fréchette, G. R., 2012. Session-effects in the laboratory. *Experimental Econommics,* p. 485–498.

Fréchette, G. R. & Yuksel, S., 2014. Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Working Paper.*

Friedman, D. & Oprea, R., 2012. A Continuous Dilemma. *American Economic Review,* p. 337–363.

Friedman, D. & Sinervo, B., 2016. Evolutionary Games in Natural, Social, and Virtual Worlds. *Oxford University Press.*

Friedman, J., 1971. A non-cooperative equilibrium for supergames. *Review of Economic Studies,* 38(1), pp. 1-12.

Gilboa, I. & Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics,* pp. 141-153.

Hedeker, D., Gibbons, R. D. & Flay, B. R., 1994. Random-Effects Regression Models for Clustered Data With an Example. *Journal of Consulting and Clinical Psychology,* pp. 757-765.

Kandori, M., 1992. Social Norms and Community Enforcement. *Review of Economic Studies,* pp. 63-80.

Kreps, D. M., P.Milgrom, Roberts, J. & Wilson, R., 1982. Rational Cooperation in the Finitely Repeated Prisoners' Dilemma. *Journal of Economic Theory,* pp. 245-252.

Lugovskyy V., P. D. S. A. W. J. W. A., 2017. An Experimental Study of Finitely and Infinitely Repeated Linear Public Goods Games. *Games and Economic Behavior.*

Mengel, F. & Peeters, R., 2011. Strategic behavior in repeated voluntary. *Journal of Public Economics,* pp. 143-148.

Murnighan, J. K. & Roth, A. E., 1983. Expecting Continued Play in Prisoner's Dilemma Games: A Test of Several Models. *Journal of Conflict Resolution,* pp. 279-300.

Normann, H. T. & Wallace, B., 2012. The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment. *International Journal of Game Theory,* pp. 707-718.

Palfrey, T. R. & Rosenthal, H., 1994. Repeated Play, Cooperation and Coordination: An Experimental Study. *Review of Economic Studies,* pp. 545-565.

Proto, E., Rustichini, A. & Sofianos, A., 2014. Higher intelligence groups have higher cooperation rates in the repeated prisoner's dilemma. *Working Paper.*

Rapoport, A. & Chammah, A. M., 1965. Prisoner's Dilemma: A Study in Conflict and Cooperation. *MI: University of Michigan Press.*

Roth, A., 1988. Laboratory experimentation in economics: A methodological overview. *Economic Journal,* 93(393), p. 974 –1031.

Roth, A., 1995. Handbook of Experimental Economics. In: *Handbook of Experimental Economics.* Princeton: Princeton University Press, p. 3–109.

Roth, A. E. & Murnighan, J. K., 1978. Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma. *Journal of Mathematical Psychology,* pp. 189-198.

Rubinstein, A., 1980. Strong Perfect Equilibrium in Supergames. *International Journal of Game Theory,* 9(1), pp. 1-12.

Schneider, F. & Weber, R. A., 2013. Long-Term Commitment and Cooperation. *University of Zurich Department of Economics Working Paper No130.*

Selten, R. & Stoecker, R., 1986. End Behavior in Sequences of Finite Prisoner's Dilemma Supergames: A Learning Theory Approach. *Journal of Economic Behavior & Organization,* pp. 589-591.

Sherstyuk, K., Tarui, N., Ravago, M.-L. V. & Saijo, T., 2013. Payment schemes in infinite-horizon experimental games. *Experimental Economics,* pp. 125-153.

StataCorp, 2013. College Station, TX, Patent No. Stata Statistical Software: Release 13.

Von Neumann, J. & Morgenstern, O., 1944. *The Theory of Games and Economic Behaviour.* New Jersey: Princeton University Press.

APPENDICES

APPENDIX I

Table 4 Average Cooperation by Round

Treatment	Horizon	Round									
		1	2	3	4	5	6	7	8	9	10
Finite	H = 2	0.10049	0.098652								
	H = 4	0.120098	0.071078	0.078431	0.139706						
	H = 6	0.08027	0.129289	0.089461	0.14951	0.129902	0.158088				
	H = 8	0.161152	0.148897	0.099877	0.079044	0.119485	0.089461	0.129902	0.109069		
	H = 10	0.079657	0.128676	0.139706	0.139093	0.129902	0.10049	0.148284	0.109681	0.109068627	0.119485
Infinite	H = 2	0.215686	0.205882								
	H = 4	0.166667	0.147059	0.205882	0.186275						
	H = 6	0.235294	0.215686	0.078431	0.117647	0.147059	0.137255				
	H = 8	0.088235	0.147059	0.137255	0.068627	0.137255	0.078431	0.107843	0.098039		
	H = 10	0.137255	0.078431	0.068627	0.068627	0.107843	0.058824	0.107843	0.147059	0.078431373	0.107843

Table 4 Average Cooperation per round aggregated for all supergames