CHARLES UNIVERSITY
FACULTY OF SOCIAL SCIENCES
Institute of Economic Studies

Learning in public goods game: alternative model performance

Bachelor's thesis

Author: Klára Grohmannová
Study program: Economics and Finance
Supervisor: RNDr. Michal Červinka, Ph.D.
Year of defense: 2024

Declaration of Authorship

The author hereby declares that he or she compiled this thesis independently, using only the listed resources and literature, and the thesis has not been used to obtain any other academic title.

The author grants to Charles University permission to reproduce and to distribute copies of this thesis in whole or in part and agrees with the thesis being used for study and scientific purposes.

During the preparation of this thesis, the author used ChatGPT 3.5 in order to debug and clarify the R scripts of model simulations. After using this tool/service, the author reviewed and edited the content as necessary and takes full responsibility for the content of the publication.

Prague, April 29, 2024

Klara Grohmannova

Abstract

Over the last fifty years the research field of the public goods game experiment became increasingly more convoluted. Especially in the area of learning models, which attempted to explain the experimental data, many models with different mechanisms were proposed and studied. However, when it comes to assessing the performance of the models in relation to each other, few conclusions were reached. This thesis attempts to further the comparative analysis conducted in Cotla (2015). Using R software simulations, three directional learning models are assessed on their ability to accurately predict experimental data. Among the considered models, K-strong equilibrium model is found to be the best predicting model for all of public goods game experiments considered in this thesis.

JEL Classification	C71, C73, C87, C92, D64, D90
Keywords	public goods game, behavioral game theory, learning
Title	Learning in public goods game: alternative model performance

Abstrakt

V průběhu posledních padesáti let se výzkum hry veřejných statků stal čím dál tím více komplexním. Zejména v oblasti modelů učení se, které se pokoušejí vysvětlit experimentální data, bylo navrženo a studováno velké množství modelů. Nícméně, co se týče porovnání výkonnosti těchto modelů vůči jiným, nebylo dosaženo mnoho závěrů. Tato práce se pokouší rozšířit komparativní analýzu provedenou v Cotla (2015). Za pomocí softwaru R, jsou tři modely učení se posouzeny na základě jejich schopnosti přesně předpovídat experimentální výsledky. Mezi těmito modely je K-strong equilibrium model tím nejúspěšnějším v předvídání všech výsledků experimentu hry veřejných statků, které jsou v této práci brány v potaz.

Klasifikace JEL	C71, C73, C87, C92, D64, D90
Klíčová slova	hra veřejných statků, behavioralní teorie her, učení se
Název práce	Veřejné statky: učení se a přesnost alternativních modelu

Acknowledgments

I would like to express my gratitude to my friends, who have kept me company throughout the process of writing this thesis. Namely my thanks goes to Andrew McCreary, Tereza Valešová, Tereza Paďourová, Lukáš Vernar and my colleagues at Advanto. I would also like to thank Ing. Lenka Dobešová for her valuable feedback. Last, but not least, my gratitude goes to my supervisor, RNDr. Michal Červinka, Ph.D., for the crucial feedback offered on the structure and wording of the thesis.

Typeset in FSV LaTeXtemplate with great thanks to prof. Zuzana Havrankova and prof. Tomas Havranek of Institute of Economic Studies, Faculty of Social Sciences, Charles University.

Bibliographic Record

Grohmannová, Klára: *Learning in public goods game: alternative model performance.* Bachelor's thesis. Charles University, Faculty of Social Sciences, Institute of Economic Studies, Prague. 2024, pages 62. Advisor: RNDr. Michal Červinka, Ph.D.

Contents

List of Tables **vii**

List of Figures **viii**

1 Introduction **1**

2 Public goods game **3**

2.1 Brief history of game theory 3

2.2 Game theory of the simple PGG 4

2.3 Setting description and variations 5

3 Research overview **10**

3.1 Basic experiment analysis 10

3.2 Social incentives and diversity 15

3.3 Introducing learning models 18

3.4 Real life significance and applications 23

4 Selected models **25**

4.1 K-strong equilibrium model 25

4.2 Noisy directional learning model 26

4.3 Reinforcement average model with a loss-aversion strategy . . . 27

5 Analysis **30**

5.1 Method 30

5.2 Data 31

5.3 Results 31

6 Conclusion **46**

Bibliography **54**

List of Tables

5.1 The overview of the data collected for the purpose of this analysis each listed with the set of parameters of the experiment in the original study. 32
5.2 The resulting mean quadratic distances from experimental data for each of the models, as well as the two bechnmarks RAND and EMP. The scores are written for both the whole dataset in 5.2 and subsets with certain parameter properties. 44

List of Figures

5.1 Simulated average contribution levels of the K-s model with different δ parameters and their comparison to empirical data (EMP) for the game parameters $N = 4, R = 0.4, T = 10$. 34

5.2 Simulated average contribution levels for the K-s model, in % of endowment, for different numbers of players for 10 round games. 35

5.3 Simulated average contribution levels for the K-s model, in % of endowment, for different values of MPCR for 10 round games. . 36

5.4 Comparison of the NDL model average contribution levels with its original ($\sigma^2 = t$) and adjusted ($\sigma^2 = \sigma_E^2$) variance against empirical data (EMP) for the game parameters $N = 4, R = 0.4, T = 10$. 38

5.5 Simulated average contribution levels for the NDL model, stated in % of endowment, for different number of players for 10 round games. 38

5.6 Simulated average contribution levels for the NDL model, stated in % of endowment, for different values of MPCR for 10 round games. 39

5.7 Simulated average contribution levels of the REL model, stated in % of endowment, for different number of players for 10 round games. 40

5.8 Simulated average contribution levels for the REL model, stated in % of endowment, for different values of MPCR for 10 round games. 41

5.9 Simulated average contribution level generated from uniform distribution on interval (0,1) for each agent in a game of 4 agents and 10 rounds. 42

5.10 Visualisation of the empirical average contribution levels constructed from the whole dataset in Section 5.2 43

Chapter 1

Introduction

Today the public goods game (PGG) setting is one of the classic laboratory paradigms of experimental economics. Although as an experiment only introduced in the 1970s (Bohm (1972)), the concepts that lead to its origin date as far back as philosophy itself. In the beginning of his second book of *Politics* Aristotle notes: “That which is common to the greatest number gets the least amount of care,” (Aristotle (1905/ca. 350 B.C.E.)). Thus, describing a phenomenon that is almost 2000 years later named by Garrett Hardin, an American ecologist, as the “Tragedy of Commons” in an article of the same name in relation to publicly accessible resources in Hardin (1968). Hardin’s essay inspires other scholars to approach this area in their research. A 2009 Nobel prize winner, economist Elinor Ostrom, showed in the 1970s, that although uncontrolled exploitation is a rational choice, varying degrees of cooperation appear in the real world to avoid the tragedy of commons scenario (Ostrom (1990)). Later, a complementary concept emerged, where instead of exploiting an existing resource, people voluntarily contribute to create a new one. Experimentally, that is what is most commonly referred to as the public goods game.

Although the game is studied extensively there is still no consensus as to what motivates people to contribute money towards the public good against the game theoretic Nash equilibrium strategy of free riding. Initially it was thought that it must be the good nature of people - their sense for the good of the whole society and joy from being selfless - that motivates their actions. Subsequent research however sewed seeds of doubt into this hypothesis. At the same time many other motives, frequently concentrated around learning, have been suggested and investigated, but none have managed to rise from the field of research to the sphere of common knowledge of behavioural economists and their students. As the public goods game framework is so simple and popular the scope of research conducted with its help or directly focused on it is immense and attempts at clarification of the findings and systematic overview are rare.

This thesis is focused on the evolution of average contributions in repeated PGGs and specifically how learning contributes to the decay in this contribution level across consecutive rounds of the game. The public goods game and its most popular variations are explored, and the mathematical description of the game is established. On these foundations a part of the PGG research is overviewed to help the reader see how learning became a prominent concept in PGG data explanation. Later a thorough description of three selected models is presented and these models are then analysed both to explore their characteristics as well as compare them amongst themselves in their ability to predict experimental data. The aim is to see whether there is a model that is significantly better at predicting the contribution levels in PGGs with different parameters. The inner workings of such model might also help in explaining the mechanisms through which individuals learn in economic experiments. The possibility of dropping some models in favour of better performing ones may also help clarify the research field in the long run.

In chapter 2 can be found the game theory necessary to understand the public goods game and some of its variations. Chapter 3 overviews the existing literature divided into four sections: the first one talking about experiments with descriptive goals, the second one being focused on experiments tying the explanation of the contribution levels in the public goods game experiments to prosociality and personal characteristics of the agents, and the third section presenting research focused on learning as an explanation of the experimental results. The last section offers a brief insight into how PGG frameworks are used in research, which extends to actual real life scenarios and can clearly highlight the significance of understanding this framework thoroughly.

Chapter 2

Public goods game

This chapter includes a brief summary of the history of game theory, mathematical description of the public goods game (PGG) and an overview of the most studied variations of the PGG experiment.

2.1 Brief history of game theory

When it comes to a rigorous approach to experiments like the public goods game, the underlying mathematical theory is known as game theory. Its mathematical foundations are attributed to John von Neumann, who cooperated with Oskar Morgenstern, to broaden the work of Emile Borel as well as his own article from 1928 (Borel (1921), Neumann (1928)). Morgenstern's and von Neumann's groundbreaking work on game theory was published in their book *Theory of Games and Economic behavior.* (von Neumann, Morgenstern 1944). A game in the sense of game theory is any interaction between at least two players, where their optimizing strategies are intertwined, i.e. the best choice for one of the players depends on the action taken by the other player(s) and vice versa. Hence the three major components of any game are the set of players (for uses in the field of economics generally referred to as agents), set of strategies (also referred to as actions), that they can choose, and a set of payoffs.

The solutions to games have been named equilibria, as a direct analogy of physical systems, where the system is said to be in equilibrium if its state remains constant unless acted upon by an external force. Initial solution approaches were limited and involved for example the simple elimination of all strategies which would lead to lesser payoffs for the agent regardless of the other agent(s)' choices. In Nash (1950) and Nash (1951) mathematician John Forbes Nash jr. proved the existence of a solution, possibly probabilistic, to any finite game, and described the solution procedure. His solution concept is known as *Nash equilibrium* and it is the most known and used equilibrium concept for judging the quality, stability and rationality of different strategies chosen by players in these games. It is a concept so well known that its definition is included also in the Oxford English Dictionary - "*noun* an array of strategies, one for each player, in which no player has an incentive (in terms of gaining an improved advantage) to change his or her strategy as long as none of the other players change theirs." In games where the *Nash equilibrium* takes on the form of agents choosing between two or more strategies with fixed probabilities, as opposed to choosing the same strategy every time, the equilibrium is said to be an equilibrium in mixed strategies.

Over the course of less than a century various concepts of equilibria have been described and studied. These are usually concepts made to specifically fit games with particular properties and real world applications, however other general solution concepts are researched as well. In particular the *Lindhal equilibrium* presents a direct analogy for the standard stability point on the competitive market - a state where the demand and supply of a good are in balance and thus the marginal cost equals the marginal revenue for every consumer. Isaac et al. (1985)

One more key concept for this thesis, which even predates the mathematical Game theory, and is well used in economics, is *Pareto efficiency*. It allows for evaluating the outcomes of a game in terms of their total social benefit. An outcome is said to be *Pareto efficient* if neither agent can achieve a better outcome for themselves without making the other agent(s) worse off. The incorporation of *Pareto efficiency* to the analysis of game outcomes is attributed to Schelling (1960).

2.2 Game theory of the simple PGG

To analyse the public goods game from the game theory perspective, let's first look at the mathematical description of the situation. The game is characterized by a set of N agents, payoff function $\pi(c), \pi : \mathbb{R}^N \to \mathbb{R}^N$, marginal rate of return R, initial endowment E_i of agent $i, (i = 1, ..., N)$ and an action set consisting of contribution c_i, such that $0 \leq c_i \leq E_i$, for each agent. Whether c_i is a real value or an integer depends on the setting of a particular experiment.

Some experiments may only have action sets consisting of two choices - that is - [contribute the full amount, not contribute]. The payoff function is defined as

$$\pi_i(c) = E_i - c_i + R \cdot \sum_{j=1}^{N} c_j, i = 1, ..., N \tag{2.1}$$

where A is the initial endowment each agent receives at the beginning of the experiment, c_i is the contribution of agent i, r is the marginal rate of return and c_j is the contribution of agent j. Most PGGs have utility equal to the payoff function and the agent's total utility is summed over rounds in the case of repeated PGG, but there are experiments which would let agents convert payoff tokens to cash or extra class credit - this would lead to a differentiation of the payoff and utility functions.

Trivially, the social optimum, i.e., Pareto efficient (optimal) point, here is when each agent contributes fully. However, as conjectured by Samuelson (1954) the private and public incentives in the environment with public goods are in conflict and the provision of public good will be suboptimal in the sense of social (Pareto) optimum. Brubaker (1975), Buchana (1968) and Olson (1974) go further and show that in ideal conditions of the simple PGG setup, rational players will not contribute at all - hence the Nash equilibrium of the game is to free-ride (i.e. contribute zero).

2.3 Setting description and variations

John O. Ledyard (1995) in the *Handbook of experimental economics* notes, that the scope of variations of a public goods game, as well as other social dilemma problems, is far wider than would seem at first sight. If we focus on the experiment as an environment in which two goods are produced - a public and a private one - it can easily be seen that any experiment that enables individuals to produce the public good from the private one via some function g and gain utility via some function u can be viewed as part of the public goods game experiment class. Among others, the Cournot duopoly model is also mentioned as a public good environment. The research covered in this thesis will focus mainly on the environment of *linear symmetric variable contribution* one (i.e. where g and u are linear, the numbering of the agents doesn't influence the outcome of the experiment, and the contributions by agents are chosen from a restricted set of natural numbers (Ledyard (1995)).

The simplest of settings of the PGG experiment was mathematically described earlier. A trivial modification of that setting is to have the experiment repeated for a finite number of rounds. From there arise endless options of modifications to the game. (Ledyard (1995)) Seven popular variations are described in this section to acquaint the reader with the complexity of mechanisms and variance of testable hypotheses:

Step-level public goods game

One of the easiest modifications to the game is the step-level modification presented by Kragt et al. (1983). In this modification, the public pool is only redistributed back to players with interest rate if a certain threshold of donations to the pool is reached. Should the threshold not be met, contributions are given back to the agents, who submitted them, and all agents keep their initial endowment. This enhances the uncertainty factor of the standard public goods game and can affect agents' behaviour differently with varying thresholds. Understanding these changes can be helpful in formation and setting of policies to achieve needed levels of contribution.

Volunteer's dilemma

This version of the game is essentially a simplification of the public goods game utility system and it is set up as such: players $i = 1, ..., N$ have zero utility to begin with and decide whether to cooperate (do a certain task such as taking out the trash in their house) or defect (let someone else do the job). If at least one player decides to cooperate and do the task at cost c, all players will be rewarded. If all players defect, they will gain zero reward. Volunteer's dilemma can be seen as a special case of the step-level public goods game. Such situations are very common in the real world and thus the experiment offers the possibility of testing different incentives, or perhaps personal characteristics of those most likely to volunteer, etc.

Tragedy of commons

As previously mentioned, tragedy of commons is a term coined even before public goods game was designed as an experiment. As a modification of the experiment itself, it is a term for a "reversal" of the game's rules. Initially, the endowment is concentrated in a public pool of money, from which each of the agents (individually, or as a group) can decide to withdraw a certain amount. This amount must be less or equal to the initial endowment.

When this game is played repeatedly for several rounds, the initial endowment recovers to the original "full" size at the rate of $r > 1$. In exact terms, the endowment in period $t + 1$

$$E_{t+1} = E_t - \sum_{j=1}^{N} c_j^t \cdot R, \tag{2.2}$$

where E_t is the endowment in period t, c_j^t is the depletion of the endowment by agent j in period t and r is the recovery rate. If the common resource is fully depleted ($E_{t+1} = 0$), the game ends. This introduces a temporal conflict to the game, where a high enough consumption of a resource in the current period leads to small or zero consumption in following periods. This setting is widely used in ecology, where depletion and protection of natural resources is experimentally investigated.

Asymmetrical endowments and/or payoffs

The next type of setting used within this experimental framework is the usage of different variants of asymmetrical endowments. Instead of every agent having the same endowment and choosing whether to put it towards the public fund, or choosing the amount to donate, there are different initial endowments, i.e., $\exists i, j \in (1, ..., N) : E_i \neq E_j$. These can take different forms from having two endowment groups, to every player having a different initial amount. These settings are used to explore how the endowment level of an agent influences the amount that they put towards the pool.

Another variation that can be added into this setting is also having asymmetrical payoffs - for example wealthier people benefiting less (having lower return rate) from the common pool and poorer people benefiting more. There it can be studied whether people will have prosocial tendencies and help raise the poor people's endowment by donating at a relatively lower cost.

Groups

Another set of variations of the experiment revolves around the number of agents participating and the information they have about other agents. Initial public goods game experiments had agents play in small groups of 4 to 10 agents (Andreoni (1988), 1995aa, 1995bb, Isaac & Walker (1988), 1994, Fehr & Gachter (2000)). Larger number of agents in a single game was then introduced to examine a possible difference between small- and large-scope experimental outcomes.

Other experiment involved studying the difference between contributions by individuals and contributions by groups instead of single agents. Groups allow for other game mechanisms such as those used by biologists to study the emergence of cooperation in populations, where networks are graphed between individual agents, and they can then only interact in potential PGGs happening in their proximity (i.e. with the agents they are connected to via the network). Groups can also be compared when they exhibit distinctive characteristics. A popular experimental treatment is to observe differences between "Partners" (PGGs where agents are assigned to the same PGG for the whole duration of the experiment and are aware of this) and "Strangers" (PGGs where each round agents are reassigned to different PGG and therefore cannot build reputation and coordinate actions with their co-players). The usage of this mechanism can be found for example in Andreoni (1988) and Fehr & Gachter (2000).

Two stage public goods game

Multiple experiments have been conducted where a particular variation of the simple PGG was played but constituted only a part of the experiment which before or after the PGG included also another phase (subgame). For example, in Kosfeld et al. (2009), the first stage had agents contribute towards an institution, that would then ensure compulsory full contributions of all agents in the PGG with existing institutions in the second stage. If the institution was rejected, agents were free to contribute any amount in the second stage PGG. Similar setups can serve to study the effects underlying institutions have on the ability to achieve socially optimal outcomes. The two-stage environment ,in general, offers many options as to the design of an experiment. There is a variety of hypotheses that can be tested using an appropriately designed two stage public goods game. Often, these will be concentrated around the possibility of designing a pre-stage to the PGG which will incentives the agents to contribute fully and achieve the social optimum.

Imperfect information

Lastly, the simple experiment in section 2.2 revealed all the relevant information regarding the mechanism of the game beforehand. However, in real life, perfect information is often unrealistic. Therefore, experiments have been conducted where the players of the game did not possess full information on the mechanism. These ranged from not knowing the number of rounds that the game is going to be repeated for to not knowing how probable it is that they will not be able to patrticipate in all the rounds or simply not knowing the return rate. Variations have been played where the players received no information, just an endowment and the option to contribute towards a public good. These experiments are used to bring the laboratory PGG setting closer to real life scenarios. In the first ever experiment resembling the PGG setting, agents also did not receive full information about the mechanism of the game played. Bohm (1972)

The reader will meet these variations of the PGG experiment and see how they contributed towards the knowledge we have about PGG today, in the following chapter.

Chapter 3

Research overview

As was previously explained, the research field of the public goods game (PGG) experiments is highly complex. Therefore, the literature review will be divided into four sections. Firstly, the reader will get to know the early history of the PGG experiments and other research aiming for descriptive analysis of the game. In a later section, experiments incorporating social preferences and other agent-specific characteristics into the explanation of the observed contributions. The third section is focused on proposed learning models explaining the deviation from the Nash equilibrium observed in experiments. In the last section a brief insight into the real life significance of PGG is looked at with the help of studies, which use the PGG framework, from fields such as ecology and medicine. Reader acquainted with these research areas will hopefully be able to appreciate the importance of the PGG research.

3.1 Basic experiment analysis

In this section initial PGG experiments and those that aggregate and analyse the PGG results, without regard to individual motivations or attempts to model and predict the experiment, are overviewed.

One of the first ever experiments involving a setup similar to the PGG was the one analysed by Bohm (1972). It aimed to estimate the demand of individuals for a certain public good via five different approaches. This experiment was conducted with the help of a TV station in Stockholm that recruited about 600 individuals to come for a paid experimental session to review a TV programme that was not shown before. Later, the participants who arrived for the session were told that the costs of running the experiment exceeded predictions and that they will only be shown the programme if in their group the total amount they are willing to pay for it exceeds the cost.

This set up directly translates to the step-level PGG setup if we see the programme as the public good, the maximum willingness to pay for it as the contribution to the public good, the costs of viewing the programme as the threshold and the benefit of seeing it as the reward. Bohm's experiment was directed towards comparing different treatments of what the individuals are actually expected to pay once the threshold is reached, but concluded that there were no statistically significant differences between the different groups. In other words, the valuation of the public good for an individual was not biased by the setup in which the individual was asked to state their valuation.

In Nunn & Watkins (1978), the free rider concept is closely investigated and special conditions reveal a possible non-cooperative contribution equilibrium. Firstly, the authors point out that free riding can be of different substance and formulate two distinct motivations of a free rider. One is that the individual's action is insignificant compared to the whole population and, therefore, the individual does not contribute as he believes it will not change anything. In the other free riding case, individuals believe that the costs of the public good will be shared among other individuals anyway and so it is not rational for him to contribute. The first case can only occur in PGGs with large enough player groups, but the latter is present everywhere.

Secondly the aim of the authors here is to prove the existence of an equilibrium different to freeriding. For that purpose, they introduce the case of a divisible and an indivisible public good. For the former, no positive contribution (non-cooperative) equilibrium is found, but for the latter (essentially a step-level PGG), such mixed strategy equilibrium can be found at probability of contribution π^* which equals the perceived probability of the $(k-1)$th player contributing over the perceived probability of other $(n-k)$ players contributing, in a game where the threshold is reached if k players contribute. Such equilibrium leads to positive contributions, although no explicit predictions are made. Essentially, this is the first attempt at finding a game-theoretical explanation for the positive contribution level observed in experiments.

A year later, in Marwell & Ames (1979) deepen the insight into PGG contribution levels by investigating different parameters of the game and their effect on the public good provision. Primarily, their source is The Logic of Collective Action (Olson (1974)), but they also conduct a wider research overview mostly on the Volunteer's dilemma to arrive at the following notions mentioned in the literature, that are valuable to state here:

1. Group size effects
    (a) PGGs played in small number of players should have higher PG provision as the investor "makes up" for less free riders in case he is the only one to contribute.
2. Interest in the good
    (a) The more the interest in the PG good in question varies, the more likely it is that someone's interest will be high enough for them to contribute.
    (b) This effect will have relatively larger significance in small groups of players.
3. Distribution of resources
    (a) If resources vary so that some individuals have enough resources left after they have contributed to towards the public pool of money, while others will be left with very little when contributing, risk avoidance becomes a relevant factor in the contribution levels.

These findings lead Marwell and Ames to pose hypotheses that they later test experimentally. The results show there is no evidence that individuals would contribute more in smaller groups than in larger ones (disproving not only the group size hypothesis but also the first free-riding "motivation" stated by Nunn & Watkins (1978), and that groups where interest in the good is unequal in deed lead to a more likely provision of the public good, while the distribution of resources is not relevant. As a side note, the article states that participants in the experiment were asked on their beliefs regarding fairness in the PGG setup and most responded that they believe in the fairness of contribution, however, there was no consensus on what constitutes a fair contribution.

Around that point in time *Lindhal equilibrium* was often mentioned with connection to the PGG experiment. Isaac et al. (1985) explore the validity of this equilibrium from the game theoretic and empirical point of view as their main objective is an attempt to design an experiment where free riding will become a reasonably valid experimental prediction. They note 3 main issues with previous experiments that might have skewed the results in favour of an alternative equilibrium. The first is lack of repetition as most experiments conducted up until that time were one shot PGG experiments. The second is the fact that a lot of these experiments also involved “all or nothing” mentality when the public good demand was discrete. Lastly, that previous experiments did not rely solely on the mechanism of voluntary contribution, but for example went through an iterative process of announcing and evaluating a proposal on the donation towards a public good. Isaac et al. (1985) conduct 9 experiments with varying returns to contributions. They make the following observations:

1. For period 5 and subsequent periods, any considered alternative equilibrium (notably *Lindhal equilibrium*) can be rejected in favour of Nash equilibrium.

2. Contribution levels are near zero but remain positive.

3. Throughout the rounds of the experiment, the provision level falls.

4. The contributions made by individuals in round 1 are uncorrelated with their decisions in later rounds.

5. Individuals contribute more under the condition of higher return on contribution than their peers with lower return. The authors also attempt to manipulate the experiment towards the Lindhal equilibrium, by announcing the prices and quantities, which indeed marks an increase of the provision of public good, but this effect is only temporary.

Andreoni (1988) explores the second observation made by Isaac et al. (1985) and aims to test if a PGG experiment of enough rounds can cause the so-called decay of contributions to zero. The article also suggests a game strategic motivation for positive contribution in the early rounds of a PGG by which players can attempt to conceal their rationality and the fact that they will eventually free ride so as not to teach other players to play the Nash equilibrium too. Therefore, Andreoni designs a two-group experiment where one group can engage in strategic behaviour as the players in one PGG remain the same and known throughout the experiment (Partners treatment) and the other one where this is not possible as each round the players in one PGG change (Strangers treatments). A restart of the experiment - after finishing the initial 10 rounds subjects are then asked to play another 10 rounds but end up being interrupted after three rounds - is used to further isolate learning and strategic effect. Contradicting the strategic play expectations, contributions by Strangers are higher throughout the whole experiment than by Partners who choose more to free ride and their total contribution decays more rapidly than that of Strangers although it does not reach a zero level even in round 10. While Strangers are only temporarily affected by the restart, Partners' decisions in round 11 largely mirror those made in round 1 and the effect is lasting. As noted, the strategies hypothesis is disproved by less free riding in the Strangers group. Learning hypothesis, in the sense of considering early rounds contributions as solely error caused, is also disproved as given the opportunity, Partners repeat their actions in the restart.

Further research by Isaac et al. (1994) investigates the relationship between the size of the group in which the game is played and the provision levels. So far, the experiment was only carried out with small sized groups and Isaac et al. (1994) aims to see whether the previous experimental results hold in large groups. At that time, it was a very common belief that larger group would lead to decreased efficiency of provision. Through a computer-based experiment conducted on 1908 subjects in 87 separate datasets, they study the provision levels in groups of 100, 40, 10 and 4 with marginal per capita (MPCR) return either of 0.3 or 0.75. For the lower MPCR, the provision rates are higher in the larger ($n = 100, 40$) groups compared to the smaller ones ($n = 10, 4$). As the MPCR rises to 0.75, there is no significant difference between large and small groups but results from Isaac & Walker (1988) are confirmed as higher MPCR leads to higher provision levels.

Group effects are further explored in an article by Sollow & Kirkwood (2000). There, group identity and gender's influence on the provision level is studied. For that purpose, players were recruited and sorted into three groups - control group of strangers where the subjects are randomly selected from the experiment, "questionnaire" group of strangers that are asked to cooperate on a task and therefore have a chance to establish interpersonal connections, and lastly the community group where players are recruited from a marching band with long tradition and prestige. Purely female groups are also compared to purely male groups. The results show that men contribute significantly more than women and community members contribute more than both the strangers and questionnaire groups, while it cannot be rejected that strangers and questionnaire groups provide the same level of the PG at 10% confidence. However, in the article, a PGG experiment is then also conducted among university sorority and fraternity members and these groups achieve significantly lower levels of provision than the strangers group. Solow and Kirkwood conclude that although group identity has an effect on the provision, determining the size and direction of the effect is not straightforward and is possibly subject to self-selection bias of particular groups.

Baptista et al. (2013) explore new technology options for conducting game-theoretical experiments by introducing a virtual space in which serious consequences can be simulated. Their INVITE (social Identity and partNership in VIrTual Environments) framework is essentially a survivor video game and was found to present a viable option for the study of PGG experiments.

Feige & Ehrhart (2015) engage in exploration of voting and transfer payments impact on the provision of a public good. They do so by designing a step-level PGG experiment where the treatment group has the option to transfer payments among players and the control group does not, while all players have different marginal contribution costs. In addition, groups where the game is played repeatedly are compared to one-shot groups. It is observed that the option to transfer payments to players with lower marginal costs leads to the social welfare maximizing outcome more often. In the control groups, this outcome is never achieved throughout the experiment.

Lastly, in this group of experiments, the author would like to mention Bailey et al. (2022). There, the PGG experiment is revisited to see whether the house-money effect is present in this environment. This effect refers to the possible differences in experiments when subjects are given a starting capital as opposed to when the money is considered by the subjects as their own. To test the hypothesis of the contribution level differing among these two groups of subjects, the authors give money to a part of their subjects well ahead of time, so that in the game they face a risk of losing money that they treat as their own. Bailey et al. (2022) however draw the conclusion that there is no evidence for a difference in the treatment and control group provision levels.

3.2 Social incentives and diversity

Differentiating slightly from the original line of questioning, Andreoni (1995a) pays attention to the possible motives behind contributions in PGG games. Three possibilities for the driving factors behind agents' decisions are proposed:

1. The warm-glow effect

   This effect refers to the good feeling people have when they do something they consider to be morally/socially right and that in consequence can steer individual decisions towards a more selfless behaviour regardless of whether someone else actually benefits from that decision.

2. Altruism

   In the article, altruism refers to the joy individuals feel when someone else, who usually is in a worse situation than them, gets to enjoy some benefit they would normally not be able to access, i.e. the joy of seeing someone else benefit. This will also motivate people to act more selflessly, but only if they know the consequences of their actions for the other people and possibly witness them.

3. Errors in decisions

   The third effect group is the chaos that comes with the human nature - all the mistakes and errors people make. This effect, if found to be predominant, would support the opinion that the decay in contribution levels in PGG might be due to learning.

Through altering the PGG mechanism, the article compares two groups of players - in one group the players can cooperate out of both kindness and error, while in the other group, the incentives to kindness (both warm-glow and altruism motivation) are removed and so presumably the only cooperative motive remaining is error in decisions. By comparison, Andreoni (1995a) notes that when the incentives for kindness are removed, there is a higher percentage of Nash equilibrium players (full free riders), and the contribution levels are about $\frac{1}{3}$ of what is usually observed in PGG experiments. The conclusion drawn from these observations is that about fifty percent of the contributions normally observed can be attributed to kindness motives. The article also mentions that the decay in contributions may be partially attributed to frustration from unreciprocated attempts at kindness.

Anderson et al. (1998) presents a theoretical continuation to Andreoni (1995b) as they attempt to create a model that captures both kindness and error motives and is able to accurately determine the contribution distribution when compared to experimental data. Their model is in the form of a probability distribution for the contribution level adopted by an individual based on particular initial parameters of the game such as susceptibility to altruism and error, initial endowment, marginal rate of return and the number of players in the game. They form three predictions of the model:

1. Contribution increases with the MPCR of the public good.

2. Total contribution increases with the number of players.
3. Mean contributions lie between the Nash equilibrium and half the initial endowment.

All of the above predictions are found to be consistent with the data considered in the article. Furthermore, maximum likelihood estimates of altruism and error parameters are significant and plausible.

A different approach to studying the factors behind the level of individual contributions is taken by Burlando & Guala (2005) where they experimentally approach categorizing people as free riders, cooperators, and reciprocators. Later, PGGs are played among homogenous agents. In their findings they observe that the average contribution level is higher when the game is played in the homogenous group rather than in the heterogenous one. While cooperators and reciprocators are seen to contribute highly and consistently through the repeating rounds and their contributions only decay in the last few rounds, free riders have lower contribution levels to begin with and there is a quick decay to the Nash equilibrium strategy.

Salahshour (2021) asks the question whether some people are consistent in cooperation/defection no matter the circumstances. And if they are, what role do they play in the resulting cooperation rates in the experimental environments. This is done through executing a PGG experiment with 2 stages. In the first stage, the PGG games are played to select the people who play in a cooperative manner. In the second round, the cooperators play the game separately from the defectors, who have “their own” PGG. This underlying mechanism is known to the players. Salahshour (2021) shows that when cooperation is rewarded with more social interaction, it helps to solve the social dilemma and consistent cooperation emerges. Even if defection is rewarded too, by advancing to the second stage, but the interaction is assorted, cooperation still evolves and sometimes even to a better level. I.e. the reward for defection may foster cooperation.

To demonstrate the complexity and diversity of opinions in the PGG field - let us look at two articles that can be considered directly clashing in opinion - Burton-Chellew & West (2012) and Neugebauer et al. (2007). The former uses the same game mechanism on three different groups of subjects and only changes what information is available to them. Namely whether they

1. have not played the game before and do not receive any information as to how the behaviour of agents influences the payoffs of other agents,

2. have not played the game before and are informed of the influence each agent has on the other agents,

3. have first received information according to 2. and are now repeating the experiment (unsure of whether it is the same game or not) but without any information being directly (re)specified again.

Burton-Chellew & West (2012) conclude that when people are informed of their influence on others (and hence there is room for prosociality and altruism in the game) they actually contribute lower amounts. The authors interpret the results as a direct contradiction to the prosociality hypothesis as the driver of the non-Nash-equilibrium strategies played in the PGG research. On the other hand, Neugebauer et al. (2007) ask participants of their experiment to first guess the total amount of contribution made and only then make their contribution. They observe that contributions steadily follow the individuals guesses and on the bases of that they argue that there exists no error (random) based motive driving the contribution levels in the PGG environment.

3.3 Introducing learning models

The presence of errors naturally implies the possibility to learn. Hence, Nax et al. (2016) explore how learning takes place in repeated games with no information other than the resulting payoff of the actions that the player took. They form and test 5 different hypotheses as to the phenomena that can be seen with agents' learning mechanism in repeated games.

1. Assymetric inertia: Agents are less likely to change their decisions drastically when payoff is steady or increasing.

2. Asymmetric volatility:The variance of decision adjustment is higher when payoff declines than when it is steady or increasing.

3. Asymmetric breadth: The deviation from a previously chosen strategy is higher when payoff declines than when it is steady or increasing.

4. Reversion: When payoff decreases, it becomes more likely that the agent will revert back to their pre-decrease choices.

5. Directional bias: If the strategies can be ordered linearly, the direction in which agents adjust their decisions tends to be reinforced throughout the repetitions.

In the experiment, subjects play two consecutive repeated PGGs, first without any information on the game but the received payoff for their actions, then with the standard PGG experiment information available to them. MPCR is also manipulated in different groups. Nax et al. (2016) conclude that all 5 phenomena are present with high statistical significance. Even if the subjects later get to know the rules and gain experience, these phenomena still continue to be present. This provides credibility to some learning models like the near-far search, “win-stay-lose-shift”, and others.

In the realm of learning models, a major overview was included in Cotla (2015), a case study, which this thesis aims to further. Eight prominent learning models are considered and compared on their predictive accuracy within Cotla (2015). All of the models are described here, to capture the mechanisms that proved less capable of predicting PGG data, alas only the important features of each model are considered in detail:

1. Reinforcement learning model

   This model is considered in Erev et al. (2008). It is based on an earlier model proposed in Erev & Roth (1998). Agents are modelled to have attractions towards certain strategies, which they then choose on the basis of a stochastic choice rule. The probability P_j^i of agent $i = 1, ..., N$ choosing a strategy $j = 0, ..., A$ is

   $$P_j^i(t) = \frac{e^{\lambda A_j^i(t)}}{\sum_{l=1}^{m_i} e^{\lambda A_l^i(t)}}, \tag{3.1}$$

   where $A_j^i(t)$ is the attractivity of strategy j to agent i in round $t = 1, ..., T$, and $\lambda > 0$ is the attraction sensitivity parameter. For high λ, strategies that are more attractive to the agent are chosen with a higher probability, while for $\lambda \to 0$, the stochastic choice rule becomes random choice.

2. Normalized reinforcement learning model

   An alternation of the previous model where the attraction sensitivity parameter becomes adaptive and decreases as the agent observes increasing variety in payoffs. Essentially, as the observed payoff variability increases, the sensitivity attraction parameter leads the agent to a random choice. Erev & Roth (1995)

3. Reinforcement average model with a loss-aversion strategy (REL)

This model is yet again an adapted version of the Erev & Roth (1998) model. Proposed by Erev et al. (1999) to address some of the problems present in the former studies, payoff variability is added. As this model is the most accurate in predicting experimental results according to the article and is considered in the practical part of this thesis, it is later described more comprehensively in section 4.1.

4. Stochastic fictious play

   The article chose to include a specification of this model from Erev et al. (2008), however, Stochastic fictious play is a learning model that has also been used in other studies such as Fudenberg & Levine (1998), Cheung & Friedman (1997) and Cooper et al. (1997). It is a model that has a belief learning nature. Rather than agents learning to make decisions solely on observations, they also employ some assumptions about the behaviour of others.

   In the first round all strategies are chosen with equal probabilities, however, the attraction of a strategy A_j^i is then updated by the attraction weight parameter times the payoff u_i agent $i = 1, ..., N$, gets from strategy j, given others have chosen strategies according to what they have chosen in round $t-1$. The probability function is then given as

$$P_j^i(t) = \frac{e^{\lambda A_j^i(t)}}{\sum_{l=1}^{m_i} e^{\lambda A_l^i(t)}}, t = 1, ..., T. \tag{3.2}$$

5. Normalized fictious play

   Described in Ert & Erev (2007), normalized fictious play is analogous to the stochastic fictious play model, however, the attraction sensitivity parameter is now transformed into a scaled version - $\lambda / PV_i(t)$ - where $PV_i(t)$ is the observed variability in payoff by agent i. Thus, PV changes every round in such a way, that the agent will trust the payoff from strategy j which is the most recent one and update their observed payoff variability by the difference of the highest overall attained payoff and payoff from the last period.

6. Experienced Weighted Attraction learning (EWA)

   Described by Camerer & Ho (1999), this model offers a more general take on a learning model as it presents a combination of reinforcement and belief learning. EWA is a three-component model comprised of Initial attraction setting, Attraction updating a Stochastic choice rule. The initial attractions are set to have zero attractions leading to equal probability of being chosen. Throughout the game, they are then updated as follows

$$A_i^j(t) = \frac{\phi N(t-1)A_i^j(t-1) + [\delta + (1-\delta)\mathbb{I}(s_i^j, s_i(t-1))]u_i(s_i^j, s_{-i}(t-1))}{N(t)}, \tag{3.3}$$

where ϕ discounts previous attractions, $\delta \in [0,1]$ allows the model to value virtual possible payoffs and actual payoffs differently, the payoff obtained by agent i when choosing strategy j while others have chosen strategy profile s_{-i} in the previous round is $u_i(s_i^j, s_{-i}(t-1))$, $\mathbb{I}(s_i^j, s_i(t_1))$ is an indicator function which returns 1 if s_i^j is equal to the strategy agent i chose in the last round. Finally, attractions are weighted by $N(t)$ function which is defined as

$$N(t) = 1 + N(t-1)\phi(1-\kappa), \text{ with } N(0) = \frac{\eta}{1-\phi(1-\kappa)}, \tag{3.4}$$

where κ determines whether the attraction updating will happen either in cumulative or in averaging fashion and η determines the strength of the initial experience weight. The stochastic choice rule is the same as in the case of the reinforcement learning model and uses the final parameter λ, which represents the sensitivity to attractions. Through different parameter settings, it can take on special forms such as the cumulative reinforcement model (Erev & Roth (1995)) for $\delta = 0, \phi = 1, \kappa = 1$; average reinforcement model (Erev & Roth (1998), for $\delta = 0, \phi = 1, \kappa = 0$); cournot learning model ($\delta = 1, \phi = 0, \kappa = 1$) and weighted fictious play ($\delta = 1, \phi = 1, \kappa = 0$).

7. Self-tuning experience weighted attraction learning (STEWA)

   The EWA model updated by Ho et al. (2007), STEWA fixes the κ parameter of EWA to 0, and the N(0) parameter to 1. This ensures the attraction updating to only consider averages. Furthermore, ϕ and δ parameters are transformed to time (round) dependant variables. As EWA, due to its high number of parameters, lead to substantial overfitting issues, STEWA was formed to avoid this issue. Ho et al. (2007) shows the superiority of STEWA to EWA in multiple games, however PGG was not included, although using STEWA for PGG does not have technical limitations. The parameters are functionalized.

$$\phi(t) = 1 - \frac{1}{2}S_i(t) \tag{3.5}$$

where $S_i(t)$ is a function that records the quadratic distance between a cumulative history vector and an immediate history vector, and therefore mimics the parameters function of discounting previous experiences.

$$\delta_i^q(t) = \begin{cases} 1, & \text{when } u_i(s_i^q, s_{-i}(t-1)) \geq u_i(s_i(t-1), s_{-i}(t-1) \\ 0, & \text{otherwise,} \end{cases} \tag{3.6}$$

so, the parameter directs the attention to strategies which would have generated a greater outcome.

8. Impulse matching learning

   The impulse matching learning model introduced by Chmura & Goerg (2012) is a regret-learning based model as in each round agents receive impulses guiding them away from strategies, which resulted in a lower payoff, and towards strategies, which resulted in higher payoffs. These impulses are translated into impulse sum vector $R_i(t)$, where $R_i^j(t)$ is the impulse sum of strategy j of agent i in round t. In this respect impulse sums can be viewed as an analogy of the attractions concept, which are a part of some of the previous models, with different attraction updating system. Impulse sums are then used to calculate the probability distributions of agents' strategy choices. Chmura & Goerg (2012) does this via simple normalization, while Cotla (2015) chose to use the logit function resulting in a stochastic choice rule analogous to the one in Reinforcement learning model and Stochastic fictious play model, introducing a free parameter λ representing the sensitivity to impulse sum.

Several descriptive measures derived from log-likelihood are considered to judge the models in Cotla (2015). Maximum likelihood estimates (MLEs) are used for the free parameters of the considered models. Random choice benchmark (RAND) and empirical frequencies benchmark (EMP) are also introduced. The former makes for a trivial benchmark to beat, while EMP is actually a hardly attainable benchmark, as it is constructed ex-post from the experimental observations. REL is found to be the best in terms of descriptive fit by a great margin over the second best performing model - the EWA model. In terms of predictive success, using log-likelihood and mean quadratic score (MQS) methods REL is again the best performing model, even beating EMP on the MQS method. It is worth noting that only reinforcement learning model, normalized reinforcement learning model, REL and EWA outperform the RAND benchmark for predictive success. Thus, we can comfortably judge the remaining models to be unsuccessful prediction models for PGG experiments. It is tested and concluded that both the components of REL are important for its comparative success and that REL is the most accurate predictor among these models.

3.4 Real life significance and applications

Not only does the PGG experiment serve to reveal the how (i.e. how public projects achieve necessary funding from people) and the why (i.e. why people would contribute more than theory would predict) of privately funded public goods, but its use is far wider.

In the field of ecology, the PGG experiment is often used in its "Tragedy of Commons" variation. Kraak (2010) uses the PGG setup and thought experiments based on previous PGG experimental data to suggest new ways of preventing overfishing (leading to the depletion of the fish stock) and promoting cooperative behaviour of fishermen. Hasson et al. (2010) present the dilemma of either investing into mitigation of climate change or into adapting to it as a PGG experiment with mitigation investment being a public good and adaption investment a private good. They conclude that more than the immediate danger from climate change, simulated as a stochastic cost to all agents mitigated by the contribution level in the public good, it is trust between people, that incentivises higher contribution levels towards the mitigation fund. Bosetti et al. (2017) explore formation of coalitions of individuals promoting clean energy and environmental preservation using the PGG set up and reach the conclusion that the presence of a mechanism, which allows for internalization of the benefits to the investing party, leads to a faster formation of a sizeable coalition. On the other hand, if the benefits are hard to internalize and lead to spillovers, it might incentivise formation of competing coalitions and lure more individuals into combating climate change.

Another field where the PGG setup has found its use is medical research. Archetti (2016) uses a non-linear public goods game to study the cooperation of cancer cells. As cancer cells produce (invest their energy into) growth factors which sustain tumour proliferation (public good beneficial to all cancer cells) the public goods game framework is a suitable one, although due to the heterogeneity of cancer cells complex network effects have to be accounted for. Archetti (2016) concludes that investigating this framework in this particular context can shed light on, for example, how resistance to therapies targeting cell growth factors can evolve and therefore help in the prevention of resistance development. Renton & Page (2021) use the same framework as Archetti (2016) to explain how greater global heterogeneity among cancer cells leads to cooperation more easily than in well-mixed cancer cell populations.

Of course, social sciences are the most common area of usage of PGG experiments. Politological research Gallier (2017) explores how the policy system, under which contribution rules punishable by a non-deterrent fine are formed, influences contribution levels. It is found that endogenously chosen rules, via the democratic system, lead to a higher contribution, not solely because of the democracy itself, but rather as a result of the information agents learn about each other in the democratic process. Hiraishi et al. (2015) use the public goods game to explore the heredity of prosociality and cooperation tendencies and Clist & Hill (2021) attempt to answer the question of whether the language you speak can affect your choices. They do this by having bilingual speakers of two rural African languages play the PGG, choosing the language of instructions randomly for each individual, and observing the differences between the two groups. The conclusion is that the language of instructions can significantly affect the social preference of the agent in the experiment. In Tomassini & Antonioni (2020) PGG frameworks are used to simulate emergence of cooperative social networks in society, contributing to explanation of history of early human civilizations development.

As the reader can see, the public goods game has found its way into many areas of research and consequently plays a large role in what is known about cooperation in its many forms. To find the models, that are most accurately able to predict PGG data is essential.

Chapter 4

Selected models

In this chapter three selected learning models for analysis will be introduced. Two of these models were introduced to the analysis as they were designed with PGG directly or at least potential games, which PGG belongs to, in mind, and they were not a part of the previous analysis conducted in Cotla (2015). The third model is then the best performing model from the 2015 article. The R code for simulations of the models is part of the appendices of this thesis.

4.1 K-strong equilibrium model

First out of the three selected models is from Nax & Perc (2015). The article defines a new equilibrium concept where a pure strategy s^* is a *k-strong equilibrium* of the public goods game if it holds that for all subsets of the population of players, the size of which is less than k, it leads the players from the subset to a greater payoff if they play the strategy s^* rather than any alternative strategy. (For formal definition, see Nax & Perc (2015). To study the predictions of the *k-strong equilibrium concepts* they build a learning model for the PGG experiment. That is why the model will be used under the name K-strong equilibrium model in this analysis. The technical description of the model follows.

The population of N agents plays the public goods game. They face a binary choice of contributing ($c_i = 1$) or defaulting ($c_i = 0$). Probability p_i represents the chance that agent i will choose to contribute. Their choices result in a private payoff $u_i = (1 - c_i) + (r/n) \sum_{j \in N} c_j$, where $r > 0$ is the fixed rate of return and r/n is the marginal per-capita rate of return which can be for simplicity denoted by R. In their simulations the game is repeated infinitely for time steps $t = 0, 1, 2....$ Nax & Perc (2015) initialise the game parameters at $t = 0$ as $p_i^t \in [\delta, 1 - \delta]$, subsequently probability of $c_i = 0$ as $(1 - p_i)$, where δ is the incremental probability by which an agent adjust their p_i in response to their observed private payoff. The author of this thesis however chooses to initialise the model at $p_i = 0.5$ for all agents. This better reflects the empirical experience from previous experiments and enables direct comparison with other models. Furthermore, restricting p_i to the $[\delta, 1 - \delta]$ interval would also corrupt the possibility to compare this model to other models. Thus p_i in the model as used by this thesis is restricted only to the usual interval $[0, 1]$. Throughout the game, p_i of agent i is then adjusted according to three rules:

Upward directional bias: If $u_i(c_i^t) > u_i(c_i^{t-1})$ and $c_i^t > c_i^{t-1}$, or if $u_i(c_i^t) < u_i(c_i^{t-1})$ and $c_i^t < c_i^{t-1}$, then $p_i^{t+1} = p_i^t + \delta$ (if $p_i^{t+1} > 1$, then $p_i^{t+1} = 1$). I.e. if the agent observes an increase in their payoff for a switch from defaulting to contributing or a decrease in payoff for a switch from contributing to defaulting, their probability of contribution will increase.

Neutral directional bias: If $u_i(c_i^t) = u_i(c_i^{t-1})$ and/or $c_i^t = c_i^{t-1}$, then p_i^{t+1} is adjusted to p_i^t, $p_i^t + \delta$ or $p_i^t - \delta$ with equal probability. Should p_i^{t+1} be < 0 or > 1 it is adjusted to $p_i^{t+1} = 0$ or $p_i^{t+1} = 1$ respectively. I.e. if the agent observes the same payoff and/or does not switch strategies their probability is equally likely to increase, decrease and stay the same.

Downward directional bias: If $u_i(c_i^t) > u_i(c_i^{t-1})$ and $c_i^t < c_i^{t-1}$, or if $u_i(c_i^t) < u_i(c_i^{t-1})$ and $c_i^t > c_i^{t-1}$, then $p_i^{t+1} = p_i^t - \delta$ (if $p_i^{t+1} < 0$, then $p_i^{t+1} = 0$). I.e. if the agent observes an increase in payoff for a switch from contributing to defaulting, or if they observe a decrease in payoff for a switch from defaulting to contributing, their probability of contribution will decrease.

Furthermore, the adjustment algorithm is subject to random perturbation ϵ of degree 0.1. This means that with a 10% chance the adjustment will not proceed according to the directional bias rules stated above but will rather adjust to p_i^t, $p_i^t + \delta$ or $p_i^t - \delta$ with equal probability. Again, should p_i^{t+1} be < 0 or > 1 it is corrected to $p_i^{t+1} = 0$ or $p_i^{t+1} = 1$ respectively, to fulfil the probability restriction.

The δ parameter is a free parameter that has to be estimated.

4.2 Noisy directional learning model

Another model featured in this analysis was described by Anderson et al. (2004). As an extension of their work on the PGG model capturing susceptibility to error, mentioned previously in Section 3.2, Anderson et al. (2004) introduce a general dynamic model for potential games. The rationale behind this model is rather simple - agents examine the total contribution level and calculate what strategy would lead them to the highest payoff. That strategy becomes their desired strategy for the next round. However, these decisions are subject to error. Unlike in Anderson et al. (1998), the error is not given in estimated parameter but rather by a standardised white noise process. This model is not explicitly specified for the public goods game, but this experiment belongs to the potential game class and thus here it will be adjusted into forming a specific model for the PGG experiment. The adjustments have to consider the discrete nature of the PGG experiment and thus a discrete equation describing the action of agent i in round t is established:

$$c_i(t) = c_i^*[\max_{c_i}(\pi_i^e(c_i, c_{-i}(t-1)))] + \sigma_i \cdot w_i(t) \tag{4.1}$$

The function $c_i^*[max(\pi_i^e(c_i, c_{-i}(t-1)))]$ describes the strategy that yields the highest payoff, when other players accumulate the same contribution level that they did in the previous round, and $\sigma_i \cdot w_i(t)$ is the standard one dimensional white noise process weighted by its variance.

4.3 Reinforcement average model with a loss-aversion strategy

Lastly the best performing model from Cotla (2015) is also a part of this analysis. As previously noted in Section 3.3, the REL model is an adaptation of a previous model and it was proposed by Erev et al. (1999). Its technical description is comprised of several components - attractions, payoff variability, cumulative payoff average and stochastic choice rule.

Attractions: Each agent $i = 1, ..., N$ has certain attractions A_j^i towards each of the possible strategies $j = 1, ..., E$. These attractions are initialised as zero for all strategies for all agents. Throughout the game, they are updated based on the observed payoff $u_i(s_i(t), s_{-i}(t))$, accumulated payoff average $PA_i(t)$, observed payoff variability $PV_i(t)$ and the strength of initial attractions $N(1)$. The updating is ruled by the following formula. Strategy j of agent i

has attraction $A_i^j(t)$ in round t:

$$A_i^j(t) = \begin{cases} \frac{A_i^j(t-1)(C_i^j(t-1)+N(1))+u_i(s_i^j(t-1),s_{-i}^j(t-1))}{C_i^j(t-1)+N(1)+1}, & \text{if } s_i^j = s_i(t-1), \\ A_i^j(t-1), & \text{otherwise.} \end{cases} \tag{4.2}$$

Payoff variability: Another component that influences the probability of agent i choosing strategy j is the payoff variability that the agent observes. This value is initialised as the expected standard deviation of actual observed payoff from a random choice from the mean payoff obtained from a random choice. The mean payoff from a random choice is trivially about half the initial endowment and the expected standard deviation is a quarter of the initial endowment in absolute value. This initialisation setting makes the agent indifferent to the size of the endowment as the payoff variability is observed in relative and not absolute measures in the first round. In the following rounds, the indifference to payoff magnitude is ensured by the accumulated payoff average component $PA_i(t)$ of the updating mechanism. After each new payoff observation (marked by round t), agent i adjusts their payoff variability.

$$PV_i(t) = \frac{PV_i(t-1)(t-1+mN(1)) + |u_i(s_i(t-1), s_{-i}(t-1)) - PA_i(t-1)|}{t+mN(1)}, \tag{4.3}$$

where m is the number of strategies available to the agent and $N(1)$ is the parameter determining the strength of initial strategies attraction.

Cumulative payoff average: The payoff average component $PA_i(t)$ is initialised as 0 and its accumulation throughout the rounds of the game is governed by the equation:

$$PA_i(t) = \frac{PA_i(t-1)(t-1+mN(1)) + u_i(s_i(t-1), s_{-i}(t-1))}{t+mN(1)}, \tag{4.4}$$

where again t signifies the round, m the number of strategies available and $N(1)$ the strength of attraction towards initial strategies.

Stochastic choice rule: Finally, all the above-mentioned components combine in the form of a stochastic choice rule to determine the probabilities with which agents will play their strategies. Probability that agent i chooses strategy j in round t is driven by the following formula

$$P_j^i(t) = \frac{e^{\frac{\lambda}{PV_i(t)} A_i^j(t)}}{\sum_{l=1}^{m} e^{\frac{\lambda}{PV_i(t)} A_i^l(t)}}, \tag{4.5}$$

where λ is the sensitivity to attractions parameter and m the number of available strategies.

This model has two free parameters λ and $N(1)$ that need to be estimated.

Now that the theoretical descriptions of selected models were established, and the reader can see the different mechanisms in play, results of simulations and comparison with experimental data will be discussed in the following chapter.

Chapter 5

Analysis

The objective of the analysis is to further the previous comparisons done by Cotla (2015) by comparing the Reinforcement average learning with loss-aversion model (REL)(Section 4.3) with the previously comprehensively described K-strong equilibrium model (K-s)(Section 4.1) and the noisy directional learning model (NDL)(Section 4.2), as these were not a part of the 2015 article. In addition to simple comparison of the performance of each model, deeper exploration of the characteristics and trends in results of each model is also conducted.

We believe that due to the high interest in PGG research this field has become increasingly convoluted and establishing viable routes for further research is essential. Therefore, finding the model that's closest to reality, even if that would still the most accurate model from Cotla (2015), can be considered an important result, which can progress this field towards a clearer overview and consensus on the topic.

5.1 Method

Firstly, all selected models are rewritten into scripts in the R statistical software, that allow for simulating the results of potential PGG games with varying parameters such as the number of rounds, number of players, marginal per capita return and endowment. Then the simulations are run with parameters equivalent to those of experiments, which are a part of the used dataset, and the resulting average contribution levels throughout the rounds of the games of the simulations are compared to the actual experimental data (see Section 5.2). The models are then given scores based on the distance of the predicted simulated results from the actual experimental results. This allows for comparison of the models in terms of how closely they can predict the outcomes of experiments. This can shed light on the possible mechanisms that affect players in their decisions, but more importantly, may be useful in selecting what model(s) are worth exploring more deeply, so they can potentially be used in real life scenarios where it is useful to know the contribution level achieved before the fund raising phase commences.

5.2 Data

The dataset gathered for the purpose of this analysis is largely based on the 18 datasets used in Cotla (2015). Only data from Botelho et al. (2009), which were included in Cotla (2015), were not readily available and therefore are not taken into consideration here. On the other hand, extra two datasets from Bailey et al. (2022) are a part of this analysis. The dataset can be seen in Table 5.1

All datasets include information on the average contribution level for each round of the experiment. Only experiments with the same basic repeated PGG experiments where included. Some of these datasets come from articles that aim to study some other mechanisms such as punishments for uncooperative behaviour, however, from there, only the control groups data were considered for this analysis. Regarding the data from (Bailey et al. (2022)), as the treatment was concluded to have no effect on the contribution levels, also the treatment groups' data is considered. These datasets may also be subject to other treatments (extra incentives to (not) cooperate, different nationalities of participants, etc.) however these treatments do not change the mechanisms of the game and, if needed, their effects are taken into account by using the random effects method of estimation of free parameters.

5.3 Results

For each unique set of parameters that enter the simulation, ten simulated round series of contribution levels were averaged to get a mean series of contribution level over rounds. The number of simulated series of average contributions was selected to be ten as it presented a reasonably low number for the computational needs while also presenting stable results. When one set

Table 5.1: The overview of the data collected for the purpose of this analysis each listed with the set of parameters of the experiment in the original study.

Dataset	Article	Treatment	Agents	Rounds	Endowment	MPCR
1, 2, 3	Andreoni (1995a)	Strangers	5	10	1	0.5
4, 7	Andreoni (1988)	Strangers	5	10	50	0.5
5, 6	Andreoni (1988)	Partners	5	10	50	0.5
8, 9, 10, 11	Andreoni (1995b)	Strangers	5	10	60	0.5
12	Isaac et al. (1988)	Partners	4	10	62	0.3
13	Isaac et al. (1988)	Partners	4	10	25	0.75
14	Isaac et al. (1988)	Partners	10	10	25	0.3
15	Isaac et al. (1988)	Partners	10	10	10	0.75
16	Isaac et al. (1994)	Partners	4	10	1	0.3
17	Isaac et al. (1994)	Partners	10	10	1	0.3
18	Isaac et al. (1994)	Partners	40	10	1	0.3
19	Isaac et al. (1994)	Partners	100	10	1	0.3
20	Isaac et al. (1994)	Partners	4	10	1	0.75
21	Isaac et al. (1994)	Partners	10	10	1	0.75
22	Isaac et al. (1994)	Partners	40	10	1	0.75
23	Isaac et al. (1994)	Partners	100	10	1	0.75
24, 25	Fehr et al. (2000)	Strangers	4	10	20	0.4
26, 27	Fehr et al. (2000)	Partners	4	10	20	0.4
28, 29, 30, 31	Sefton et al. (2007)	Partners	4	10	1	0.5
32, 33	Kosfeld et al. (2009)	Partners	4	20	20	0.4
34, 35	Kosfeld et al. (2009)	Partners	4	20	20	0.65
36, 37	Bailey et al. (2022)	Partners	3	20	50	0.5
38 - 53	Gachter et al. (2010)	Partners	4	10	20	0.4
54	Keser et al. (2000)	Strangers	4	25	10	0.5
55	Keser et al. (2000)	Partners	4	25	10	0.5

of parameters was simulated more times, the mean contribution levels did not proceed to change significantly anymore. Hence ten simulation repetitions were decided to be used for the mean of every simulated model given certain parameters. Firstly, the models are investigated on their own and later compared with real data and amongst themselves.

K-strong equilibrium model

The model presented by Nax & Perc (2015) contains a free parameter δ that represents the step in probability by which an agent adjusts their probability of contribution. At first, the strategy to estimating the correct δ value was to use random effects estimation. This would be advantageous as the random effect model would integrate any oddities of the specific experiments, such as additional incentives. Merrett (2012) also speaks in favour of the RE method of estimation when looking at PGG data. Since only data on the average contributions and not individual contribution levels are a part of the dataset used here, there was the presumption, that the δ parameter could be estimated as the effect the previously observed change in average payoff has on the average probability of contributing to the public pool of funds. This estimation was done in an out-of-sample way where each experimental dataset for which δ was estimated was taken out of the dataset and δ was inferred only from the remaining observations. Using this approach would help avoiding overfitting issues. The estimated values for most datasets could be rounded to 0.003 and ranged from 0.001 to 0.005.

However, when the simulation is run with the δ parameter of this magnitude (≤ 0.005), the adjustment in probability is unable to direct the agents away from the initial randomizing strategy (average contribution levels stay very close to the 50% of endowment level) in the 10 rounds that most experiments were conducted for. In the actual experimental data, a significant decay in average contributions can be observed already within the 10 round limit and thus δ parameters of magnitude 0.005 and less would not produce a simulation model which would be able to predict experimental outcomes in any accurate way and would not have been significantly different from the random choice benchmark (RAND). Hence it was decided to test another approach to choose a δ value that would ensure simulation results suitable for comparison. This was achieved by selecting the parameters most common in the dataset ($N = 4, R = 0.4, T = 10$) and averaging over the experimental results in experiments with these parameters. The resulting series of average contribution levels consitutes the EMP line seen in Figure 5.1). Although the simulations are initialised at the uniform distribution mean of 50% of endowment contribution for each agent, the EMP curve starts at a slightly higher level in the first round. This is believed to be due to the dataset not being large enough to get to the 50% level, which is assumed to be the true empirical mean of the first round. Concurrently, for the same parameters, average probability series were simulated (and again averaged over to achieve a more stable average probabilities) with δ parameters 0.05, 0.075, 0.1, 0.125 and 0.15. As can be seen in Figure 5.1, for the first two parameter values, the decay was too slow to follow the empirical probabilities and for 0.15 the decay became too quick. For values 0.1 and 0.125 the closeness to the empirical data (EMP) was about the same, but the 0.125 simulated averages reflected the decay in empirical data (EMP) a bit better in the last round. Thus, δ was chosen to be used as 0.125 in the subsequent analysis. The model is utilized here with the perturbation parameter $\epsilon = 0.1$ as is used in the

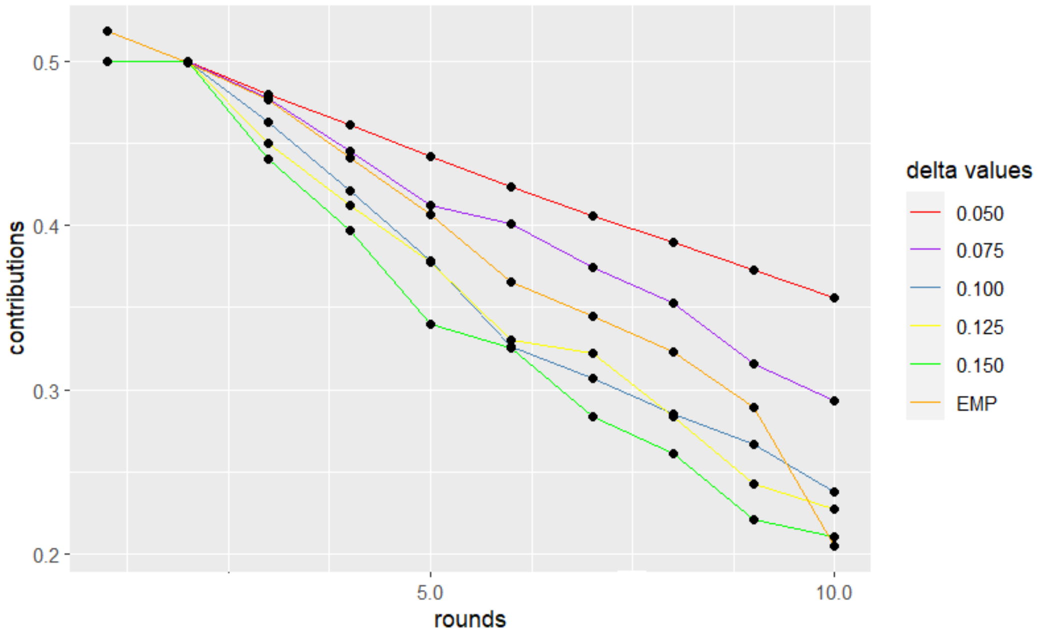

Figure 5.1: Simulated average contribution levels of the K-s model with different δ parameters and their comparison to empirical data (EMP) for the game parameters $N = 4, R = 0.4, T = 10$.

general simulation in Nax & Perc (2015).

A note has to be made on the fact that the model by Nax & Perc (2015) is designed to handle binary decision making by the agent. In order to use it to predict data in experiments, where the agents have more discrete choices, it is better to turn to the average underlying probability rather than the average contribution simulated. To rationalize using this method, the author of this thesis considered the probability of contributing fully from the model as the desire to contribute an amount less than 1 and more than 0. However, that is not possible and the agent mixes between strategies 0 and 1 with the probability to at least achieve her desired contribution hypothetically on average over multiple games. Therefore in games with more strategies, when average contribution levels are divided by the agents' endowment, they are a direct analogy for the simulated average probability series. This is the key characteristics enabling comparison of simulated data form this model with experimental data.

The adjustment algorithm of the model presents a built-in randomizing probability of contribution for the first two rounds as only after receiving two payoffs, the agents can infer what effect their actions might have on their payoffs. In the third round, the probability of contribution - which can also be interpreted as the contribution level in percentage of endowment - deviates from 0.5 to lower numbers. Over the next rounds a slow decay towards the Nash equilibrium contribution level of 0 can be observed on average, which is however only very rarely reached in single runs of the simulation.

To observe how the properties of the model change when parameters change, the simulation results were grouped by the parameter in question and averaged to then be plotted against other parameter values. Only parameter values that appear in the gathered experimental data were considered for these simulations. The change in simulated contribution levels when the number of players playing the PGG changes can be seen on Figure 5.2.

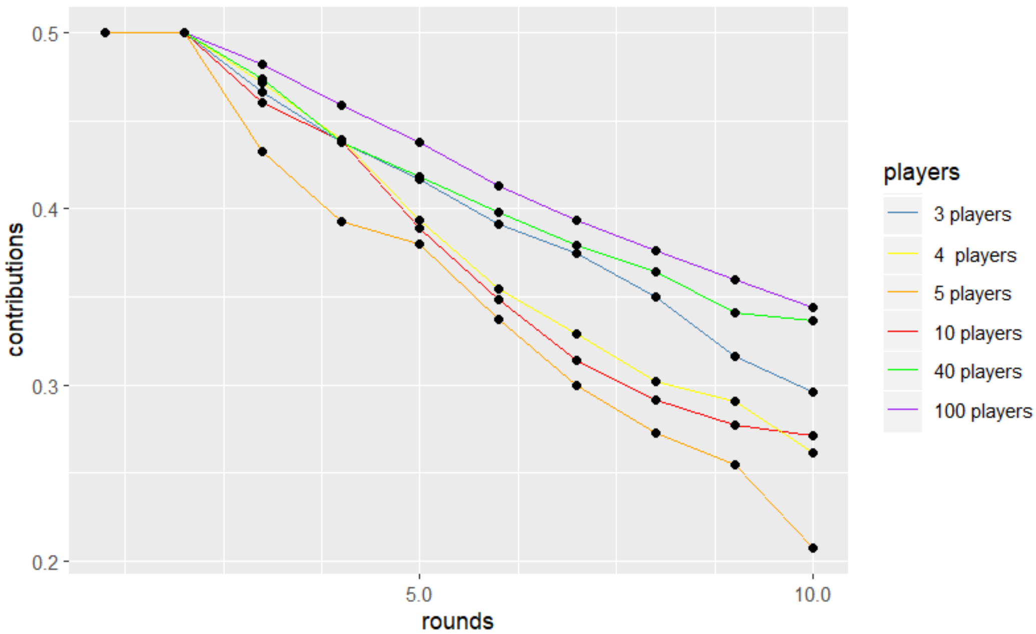

Figure 5.2: Simulated average contribution levels for the K-s model, in % of endowment, for different numbers of players for 10 round games.

All the mean contribution levels produced by the simulations in Figure 5.2 lie in the open interval $(0.2, 0.35)$. The simulated results for 40 and 100 players lead to higher contribution levels than results for 3, 4, 5 and 10 players. However, no direct relationship can be inferred as e.g. a 3-player game yields a higher end-round contribution level than both a 5-player and a 10-player game does.

The response of the mean simulated contribution levels for different values of the MPCR was investigated in an analogous manner. The results can be seen in Figure 5.3.

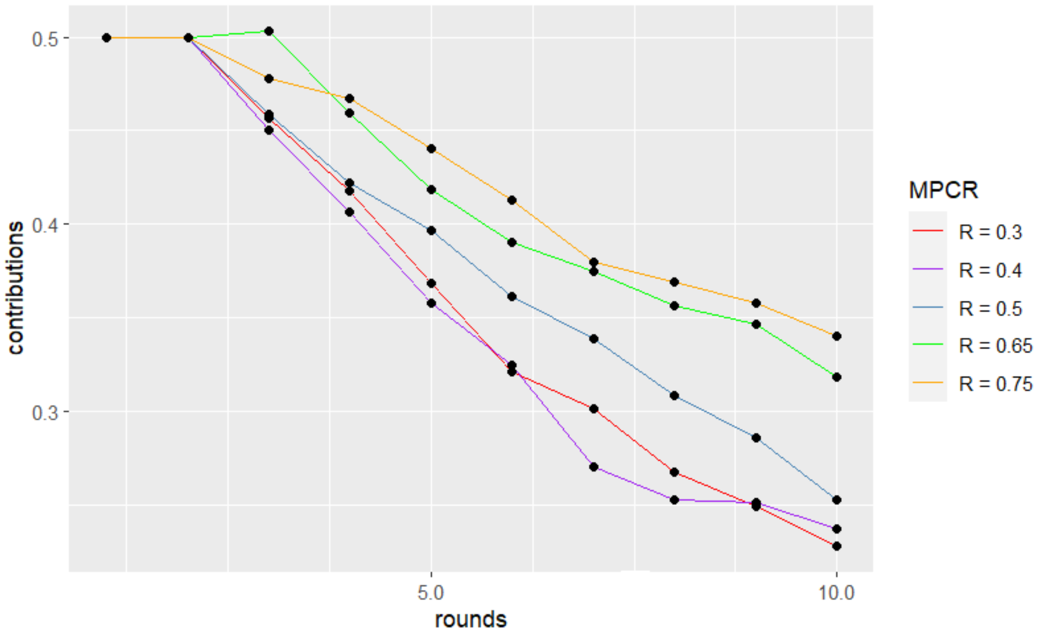

Figure 5.3: Simulated average contribution levels for the K-s model, in % of endowment, for different values of MPCR for 10 round games.

Here, a clear relationship between the end round contribution level and the marginal per capita rate of return can be seen as $R = 0.75$ yields the highest contribution level of 0.34% in round 10 and $R = 0.3$ the lowest level of 0.23% with the in-between values yielding lowering contribution levels as the MPCR drops.

As this model is only yielding contribution levels in terms of probability of contributing fully, the absolute volume of the endowment agents have at hand is not relevant at all for the simulated results.

In general, this model produces contribution level time series that are not at first glance inconsistent with experimental results as it shows a slow decay of contributions throughout the 10 rounds which is the most common experimental length of the repeated PGG.

Noisy directional learning model

For the noisy learning model, it was first attempted to estimate it according to the original equation from Anderson et al. (2004) (Equation 4.1). However as the noise is weighted by the variance of the white noise process and the process itself takes on the form of the normal distribution with the mean $\mu = 0$ and variance $\sigma^2 = t$ in one dimension, the noise turned out to be insufficient to direct the agents choices away from the Nash equilibrium fast enough and also showed a clearly increasing rather than decreasing series of average contributions. After observing the results for the first couple of sets of parameters, it became clear that this model does not perform nearly as well as the other models. Especially in cases where the endowment did not present a binary decision for the agent. This could be solved by either standardizing all experimental datasets for this binary decision making or adjusting the noise for the endowment. However, if the variance weighting remained, the noise would lead to an increase in contribution levels over rounds of the game which is inconsistent with experimental data, that shows a significant declining trend. Hence the author of this thesis decided to weight instead the noise process by the variance of the endowment, where it is viewed as a uniform probability distribution, which fixes both the adjustment of the noise to the size of the endowment and the increasing contribution level inconsistency.

For better illustration of the inability of the original model to compete in predictive ability with the other models and why the adjusted model is a better fit see Figure 5.4. The EMP data used here is the same as in Figure 5.1. While the model with original variance results in predictions that are increasing even beyond the 10th round, while the model with adjusted variance stabilises aproximately around the EMP contribution level in later rounds.

The model was investigated on the response of mean simulated contribution levels to changing parameter values. Again, the experimental dataset's set of parameters was grouped by a particular parameter value and averaged over it to gain one mean contribution level series for the 10 round PGG. For the purpose of comparison of the simulated results, all average contributions were transformed to average contributions in % of endowment.

Looking at Figure 5.5, immediately visible is the direct drop from the 0.5 contribution level at which the model is initialised to contribution levels around 0.2. This is an inherent characteristic of the model and would be true for any initialisation value. As all agents realise, that since in the first round all have contributed half of their endowment, they can be better off by not contributing at all, given that the assumption, of other agents keeping their contributions fixed, is true. The model does not allow for higher order thinking. Subsequently, through the noise component average contribution levels end up around the 20% level instead of 0 in round 2. Throughout the remaining rounds, the contribution level rises steadily to the 10th round level of about 30% of endowment. This happens irrespective of the number of players and all

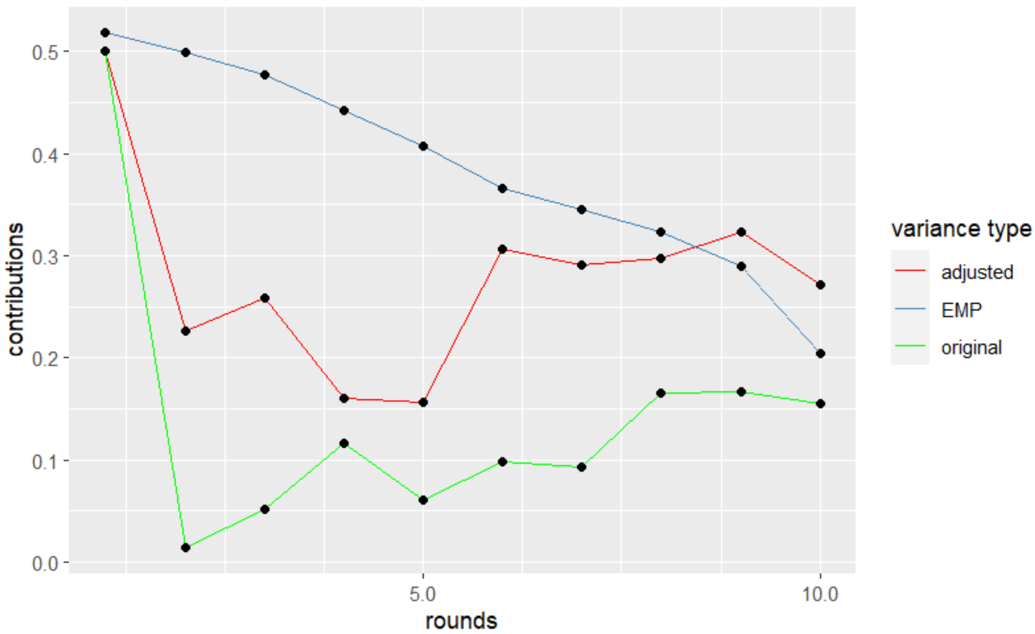

Figure 5.4: Comparison of the NDL model average contribution levels with its original ($\sigma^2 = t$) and adjusted ($\sigma^2 = \sigma_E^2$) variance against empirical data (EMP) for the game parameters $N = 4, R = 0.4, T = 10$.

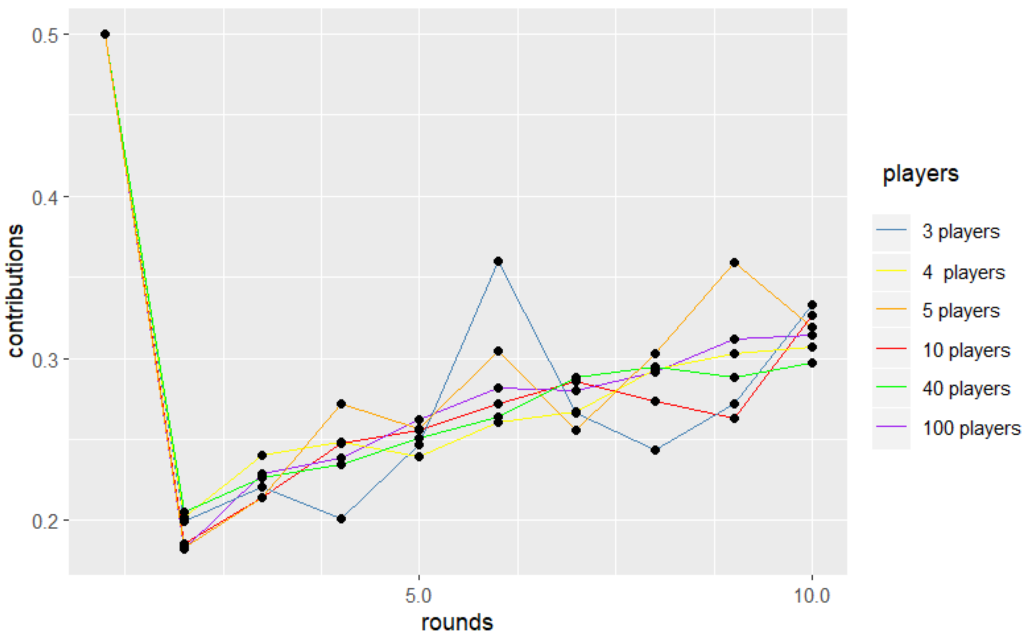

Figure 5.5: Simulated average contribution levels for the NDL model, stated in % of endowment, for different number of players for 10 round games.

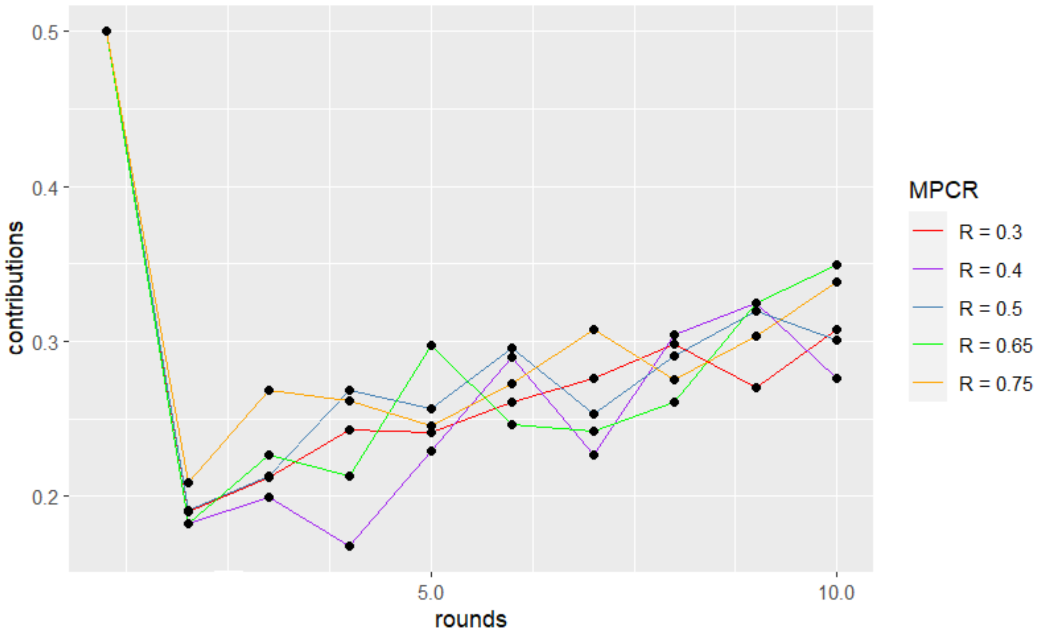

Figure 5.6: Simulated average contribution levels for the NDL model, stated in % of endowment, for different values of MPCR for 10 round games.

the mean simulated contribution levels follow approximately the same trend line. It can be noted also that the results for 100 and 40 players appear to be less volatile around the trend than results for lower number of players. This rising trend is in direct contrast with the experimentally observed decay in contribution levels and immediately suggests that this model will be a worse predictor of the experimental results than a model which will have a stable or decreasing trend.

Similarly to the previous figure, Figure 5.6 shows an increasing trend in all parameter values without any stable distinction between them. Although $R = 0.75$ and $R = 0.65$ results achieve the highest contribution levels in the 10th round, the course of the evolution of all of the simulated contribution level series suggests that this is not a given characteristics but rather a random event.

Reinforcement average model with a loss-aversion strategy

The REL model has two free parameters that need to be estimated in order to run the simulations. Since both λ and $N(1)$ parameters are a part of the individuals stochastic choice rule, it is impossible to infer them from the average contribution level dataset. Since the dataset used for this analysis is largely

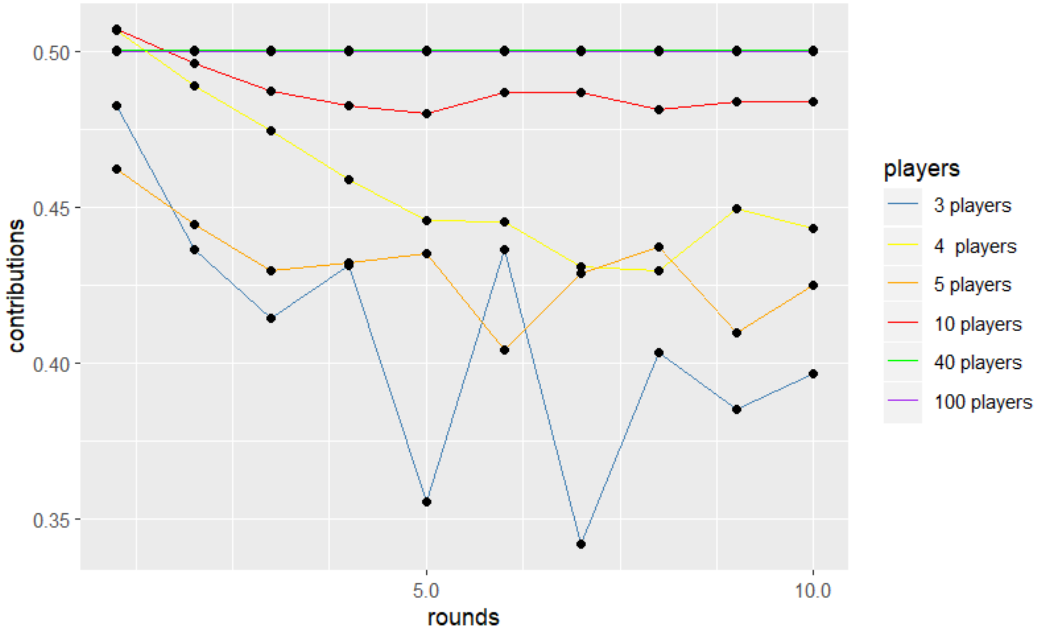

Figure 5.7: Simulated average contribution levels of the REL model, stated in % of endowment, for different number of players for 10 round games.

overlapping with the one used in Cotla (2015), the parameter values were taken over from their estimation and set to be $\lambda = 2.2983$ and $N(1) = 1.0839$. Other parameter values were shortly looked into as well, but none seemed to produce results noticeably better fitting to experimental data, than the results produced with the parameters from the 2015 article. Therefore, the analysis is continued only for this specific value combination of the free parameters.

Figure 5.7 shows how the model reacts to different number of players playing the game. Although one of the features of the REL model should be invariance to the size of the endowment and is therefore invariant also to the absolute size of the payoffs, the simulation results hint that this invariance has a limit. When there are enough players in the game, the accumulated public fund offers a return that is big enough to update the attraction of the initial strategy to a remarkably high number. This leads the initial agents' strategies (that were chosen randomly from a uniform distribution) to absorb all probability on themselves making them a sure event. As the number of players is high enough, the randomized first choice leads to a stable 50% of endowment average contribution level. Already at 10 players the simulation results to a less volatile process that decays under the 50% level only slightly. More decay can be observed for 3, 4 and 5 agents playing the game, however a 4-player game

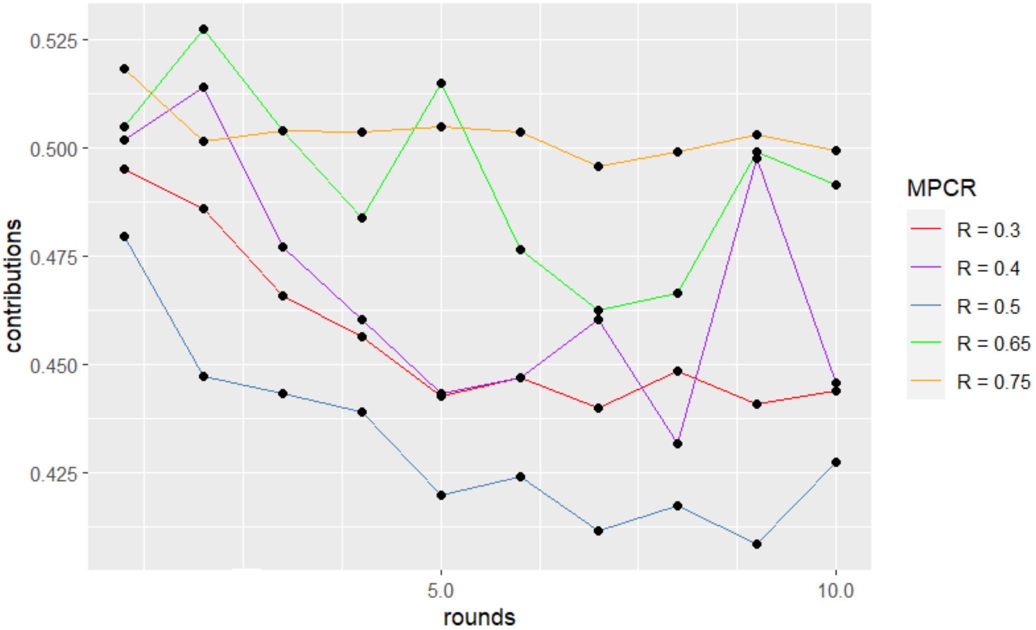

Figure 5.8: Simulated average contribution levels for the REL model, stated in % of endowment, for different values of MPCR for 10 round games.

yields a higher 10th round contribution level than the 5-player game. More importantly all REL simulations yield a relatively high average contribution levels in the 10th round starting from approximately 40% of endowment contribution level.

Figure 5.8 shows that also high value of MPCR has the same effect on the model as a high number of players has. Since the simulation was run for 4, 10, 40 and 100 players and then averaged, the player factor has negligible impact on the results for $R = 0.75$. Again, players' attractions are updated so graciously towards the initial strategy, that deviating from it becomes exceedingly rare, although compared to the 40 and 100 players simulations averaged over different return rates there are at least some change in agents' decisions over time. Here, the simulation yields even higher end round contribution levels not getting significantly under the 42.5% average contribution level. There is some more volatility in response to MPCR changes than was seen previously with changes invoked by changing number of players.

As in the experimental data, some decay in contribution levels can be seen even in games with high MPCR and/or high number of players it is expected that REL will perform worse in the prediction of average contribution levels in these games.

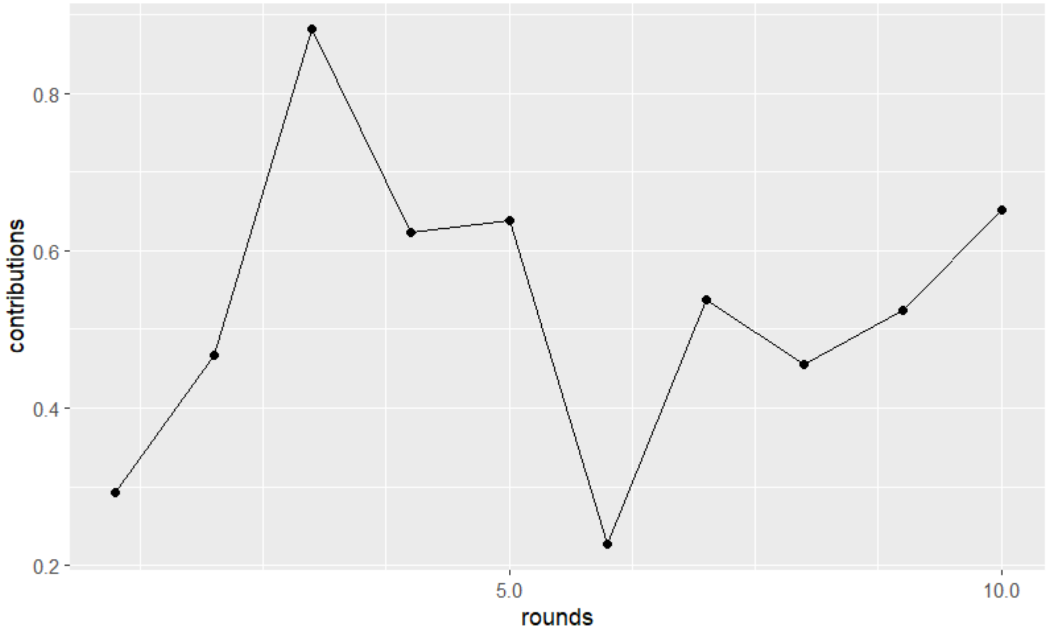

Figure 5.9: Simulated average contribution level generated from uniform distribution on interval (0,1) for each agent in a game of 4 agents and 10 rounds.

Comparison of models with experimental data

Before the models are compared against themselves in their ability to fit with the actual average contribution levels recorded in experiments, two benchmarks need to be introduced. First, it is the randomizing benchmark (RAND) that provides a basic check that the model is able to beat random choices in terms of distance of predictions from experimental data. Since many of the experiments were played among small groups, the virtual randomizing behaviour notably does not have to (and most likely will not) lead to a steady 0.5 average contribution level and can deviate from it quite significantly. To allow for a direct analogue of the selected models' analysis, randomizing simulations are run for each set of parameters and replicated respectively in the RAND simulated dataset. An example of a RAND benchmark for parameters $N = 4, T = 10$ and contribution levels in terms of percentages of endowment is shown in Figure 5.9.

For curiosity purposes an empirical benchmark (EMP) is also introduced, which is represented by averaging average contribution levels over all experimental datasets. This yields the following average contribution level series seen in Figure 5.10. Note that this EMP curve is different to the one used previously in Figures 5.1 and 5.4 as it is constructed out of the whole dataset presented

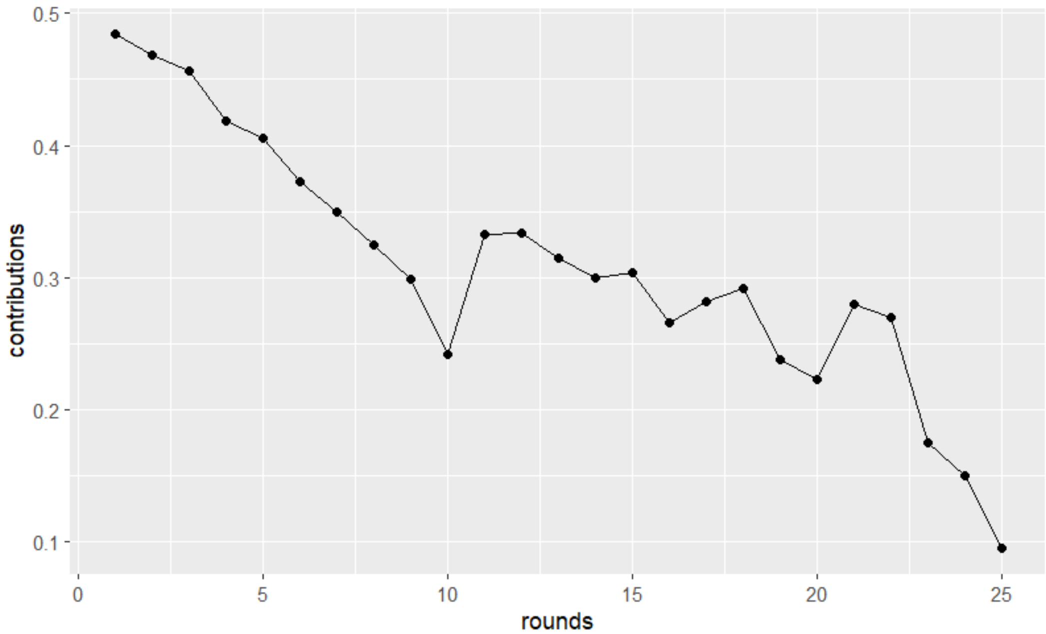

Figure 5.10: Visualisation of the empirical average contribution levels constructed from the whole dataset in Section 5.2

in Section 5.2. The first round EMP mean is however different from 0.5 for the same reasons as mentioned with the previous EMP curve.

Like previous cases, for averaging and comparison purposes all data have to be transformed into average contribution levels in terms of percentages of the agents' endowment. Smooth decay from about the 50% contribution level to around 25% in round 10 is followed by a jump back up to a 33% contribution level and then again a decay is observed to around 10% contribution level in round 25. The sudden jump is caused by the averaging process over all experimental data as most experiments end in round 10 and the last two rounds tend to accelerate the decay as was described by Andreoni (1988).

To compare the models with experimental data the author takes the simulated results for each set of parameters and count their distance from the average contribution levels observed in each experiment. The distance is measured in terms of difference of simulated and experimental average contribution level squared. Distances are then summed over all rounds of each experimental dataset resulting in a squared distance score of the model. These scores can then be used to compare the performances of different models. To get a general predictive ability score the resulting squared distance scores are averaged to get a single mean squared distance (MSD) score.

Since from previous analysis of the trends in simulations of each model it

Table 5.2: The resulting mean quadratic distances from experimental data for each of the models, as well as the two bechnmarks RAND and EMP. The scores are written for both the whole dataset in 5.2 and subsets with certain parameter properties.

	EMP	K-s	REL	NDL	RAND
all datasets	0.02001	0.18005	0.39079	0.48470	0.58088
high # of players	0.00837	0.06504	0.19548	0.24342	0.30593
low # of players	0.02199	0.19962	0.42404	0.52694	0.62767
high MPCR	0.01807	0.13638	0.24423	0.61689	0.20767
low MPCR	0.02026	0.18748	0.41574	0.46336	0.64439
high endowment	0.01936	0.14877	0.39420	0.33207	0.75538
low endowment	0.01865	0.19237	0.41250	0.57705	0.56535
Strangers	0.02510	0.19932	0.63655	0.37777	0.90873
Partners	0.01843	0.17408	0.31472	0.51911	0.47939

is known that some models (namely REL) are likely to be worse at predicting games with certain set of parameters over others also comparison of the models only on subsets of experiments with different properties is explored. In addition to the previously explored varying parameters, such as MPCR and number of players, subsets formed on endowment levels and Strangers/Partners treatment (see chapter 2) are formed and explored.

The table of results, Table 5.2, includes mean quadratic distance scores for each studied model as well as for the empirical average contribution level (EMP) and the simulated average contribution level from randomizing uniformly over strategies (RAND). Scores are listed not only for the overall analysis of the whole dataset, but also for the analysed subsets, main characteristics of which are noted on the left-hand side. Thresholds for high/low subset categories were chosen this way: high number of players ($N = 10, 40, 100$), low number of players ($N = 3, 4, 5$), high MPCR ($R = 0.65, 0.75$), low MPCR ($R = 0.3, 0.4, 0.5$), high endowment ($E = 50, 60, 62$), low endowment ($E = 10, 20, 25$). Data collected in form of percentages of endowment without reference to the absolute size of the endowment were disregarded for the purpose of analysis of low and high endowment subsets.

From the Table 5.2 it is clear that EMP as a benchmark is unchallenged. This presents minor inconsistency with Cotla (2015) as they find that on some occasions the REL mode is able to beat the EMP benchmark. In this analysis, however, it does not score well enough to even come close to the benchmark. The ex-post empirical average is however a very tough benchmark. In all datasets K-strong equilibrium model (K-s) achieves the lowest MQD score and thus beats all other models in terms of the accuracy of its predictions. This happens not only for the dataset as a whole, but also for every analysed subset of the dataset. On the general dataset level REL model comes in second and the noisy directional learning model (NDL) performs the worst although it still manages to beat the RAND benchmark. The fact, that REL beats the RAND benchmark for the whole dataset is consistent with the findings in Cotla (2015). In two special cases - when only data from games with high endowment or only games where the players are randomly re-matched into games every round (Strangers treatment) - REL scores worse than NDL does (NDL scores highlighted in yellow). Since it observed that NDL trends in the opposite direction to EMP, NDL beating REL in certain subsets means that REL is a bad predictor of games with these parameters. Although REL notably does not score as bad as would be expected in the high number of players and high MPCR subsets given the observation made in Section 5.2.3 (5.8,5.7), this is mainly caused by the fact that the average contribution levels observed in the experiments are very close to those arising from randomising behaviour (as seen from the low RAND score, highlighted in red) and REL tends to repetition of the initial rounds' choices for every agent under these parameters it yields similar average contribution levels as RAND. In the case of high MPCR subset both REL and NDL are even beaten by the RAND benchmark.

In the setting of this analysis the K-strong equilibrium model, inititialised at $p_i = 0.5$ for every agent and with learning parameter $\delta = 0.125$ and distrubance parameter $\epsilon = 0.1$, is the unquestioned winner against the other two models in terms of its ability to accurately predict average contribution levels of PGGs with different parameter sets.

Chapter 6

Conclusion

This thesis presented three prominent models of learning in the public goods game and compared them to actual experimental data. In particular the Reinforced average model with loss-aversion strategy, the Noisy directional learning model and the K-strong equilibrium model were thoroughly explained, studied for their characteristics and results, and judged on their ability to accurately (with minimum deviation distance) predict experimental data. This research found that the K-strong equilibrium model with 10% chance of perturbation occurrence and learning parameter $\delta = 0.125$ is the best performing model in predicting the chosen dataset among these models. In relation to Cotla (2015) the results regarding the REL model are consistent in the aspect of scoring better than the random choice simulations, but a minor inconsistency regarding the relation to the benchmark constructed from empirical data was also found.

In other sections of this thesis, it was explained how learning models became a supported mechanism behind the decay in contributions observed in experiments. The empirical dataset is in accordance with these findings. It also discussed the importance of understanding PGG results and the ability to predict them, as the PGG framework is a staple experiment in high variety of fields studying any form of cooperation. Assuming Cotla (2015) is correct in their analysis of the models' performance all explanations of the PGG data based on these models can be dropped in favour of not only the best performing model - REL - from the 2015 article, but also in favour of the even better performing K-strong equilibrium model.

It is the natural implication of this research, that more effort should be directed towards studying this model. Perhaps finding supporting data for the underlying decision-making mechanism of the individual agents used in this model, or finding variations of the model, which are able to predict PGG average contribution data with greater detail. Dropping alternative routes stemming from models that perform worse than the K-strong equilibrium model may result in some degree of clarification in the complex research field. The model itself may, with modifications, be helpful in predicting real life data of contributing towards some public good, but that is also a question yet to be answered by follow up research.

The results of this thesis are valuable to all researchers interested in the public goods game, especially when prediction of the resulting average contribution levels is vital. Hopefully, these results can also help those, whose work revolves around setting up policies for the collection of funds for privately funded public goods using frameworks similar to the linear public goods game.

References

Anderson, S., Goeree, J., & Holt, C. (1998). A theoretical analysis of altruism and decision error in public goods games. *Journal of Public Economics, Vol. 70, issue 2, p. 297-323*. doi: https://doi.org/10.1016/S0047-2727(98)00035-8

Anderson, S., Goeree, J., & Holt, C. (2004). Noisy directional learning and the logit equilibrium. *Scandinavian Journal of Economics*, *Vol. 106, issue 3, p. 581-602*. doi: https://doi.org/10.1111/j.0347-0520.2004.00378.x

Andreoni, J. (1988). Why free ride?: Strategies and learning in public goods experiments. *Journal of Public Economics*, *Vol. 37, issue 3, p. 291-304,*. doi: https://doi.org/10.1016/0047-2727(88)90043-6

Andreoni, J. (1995a). Cooperation in public-goods experiments: Kindness or confusion? *The American Economic Review, Vol.85, issue 4, p.891-904*. doi: http://www.jstor.org/stable/2118238

Andreoni, J. (1995b). Warm-glow versus cold-prickle: The effects of positive and negative framing on cooperation in experiments. *The Quarterly Journal of Economics, Vol. 110, issue 1, p. 1-21*. doi: https://doi.org/10.2307/2118508

Archetti, M. (2016). Cooperation among cancer cells as public goods games on voronoi networks. *Journal of Theoretical Biology, Vol. 396, p. 191-203*. doi: https://doi.org/10.1016/j.jtbi.2016.02.027

Aristotle. (1905/ca. 350 B.C.E.) *Aristotle's politics.* Oxford: Claredon Press, Translated by Jowett, B., original work published cca 350 B.C.E.

Bailey, N. T., Ramalingam, A., & Stoddard, B. V. (2022). Experimental (re)analysis of the house-money effect in a public goods game. *Journal of* *the Economic Science Association, Vol. 9, issue 1, p. 1-14*. doi: https://doi.org/10.1007/s40881-022-00122-2

Baptista, M., Damas, H., Dimas, J., Raimundo, G., & Prada, R. (2013). A serious game based on a public goods experiment. *Social Computing Conference, Proceedings - SocialCom/PASSAT/BigData/EconCom/BioMedCom 2013. p. 774-781*. doi: http://dx.doi.org/10.1109/SocialCom.2013.116

Bohm, P. (1972). Estimating demand for public goods: An experiment. *European Economic Review, Vol. 3, issue 2, p. 111-130*. doi: https://doi.org/10.1016/0014-2921(72)90001-3

Borel, E. (1921). La theorie du jeu et les equations integrales a noyau symetrique. *Comptes Rendus de l'Academie des Sciences, Vol. 173, p. 1304-1308*. doi: http://www.jstor.org/stable/1906946

Bosetti, V., Heugues, M., & Tavoni, A. (2017). Luring others into climate action: coalition formation games with threshold and spillover effects. *Oxford Economic Papers, Vol. 69, issue 2, p. 410-431*. doi: https://doi.org/10.1093/oep/gpx017

Botelho, A., Harisson, G., Costa Pinto, L., & Rutstrom, E. (2009). Testing static game theory with dynamic experiments: A case study of public goods. *Games and Economic Behavior, Vol. 67, issue 1, p. 253-265*. doi: https://doi.org/10.1016/j.geb.2008.11.013

Brubaker, E. (1975). Free ride, free revelation, or golden rule? *The Journal of Law and Economics, Vol. 18, issue 1*. doi: https://doi.org/10.1086/466808

Buchana, J. M. (1968). *The demand and supply of public goods.* Chicago: Rand, McNally & Company.

Burlando, R., & Guala, F. (2005). Heterogenous agents in public goods experiments. *Experimental Economics, Vol. 8, p. 35-34*. doi: https://doi.org/10.1007/s10683-005-0436-4

Burton-Chellew, M., & West, S. (2012). Pseudocompetition among groups increases human cooperation in a public-goods game. *Animal behaviour, Vol. 84, p. 947-952*. doi: https://doi.org/10.1016/j.anbehav.2012.07.019

Camerer, C., & Ho, T. H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica, Vol.67, issue 4, p. 827-874.* doi: https://doi.org/10.1111/1468-0262.00054

Cheung, Y., & Friedman, D. (1997). Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior, Vol. 19, issue 1, p. 46-76.* doi: https://doi.org/10.1006/game.1997.0544

Chmura, T., & Goerg, R., S.J. Selten. (2012). Learning in experimental 2x2 games. *Games and Economic Behaviour, Vol. 76, issue 1, p. 44-73.* doi: https://doi.org/10.1016/j.geb.2012.06.007

Clist, P., & Hill, J. (2021). Bilinguals in the lab: (why) does randomising language affect cooperation? *AEA RCT Registry.* doi: https://doi.org/10.1257/rct.4616-1.0

Cooper, D., Garvin, S., & Kagel, J. (1997). Adaptive learning vs. equilibrium refinements in an entry limit pricing game. *Economic Journal, Vol. 107, issue 442, p. 553-75.* doi: https://doi.org/10.1111/j.1468-0297.1997.tb00027.x

Cotla, C. R. (2015). Learning in repeated public goods games - a meta analysis. *SSRN.* doi: http://dx.doi.org/10.2139/ssrn.3241779

Erev, I., Bereby-Meyer, Y., & Roth, A. E. (1999). The effect of adding a constant to all payoffs: experimental investigation, and implications for reinforcement learning models. *Journal of Economic Behavior & Organization, Vol. 39, issue 1, p. 111-128.* doi: https://doi.org/10.1016/S0167-2681(99)00028-1

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making, Vol. 21, issue 5, p. 575-597.* doi: https://doi.org/10.1002/bdm.602

Erev, I., & Roth, A. E. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior, Vol. 8, issue 1, p. 164-212.* doi: https://doi.org/10.1016/S0899-8256(05)80020-X

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review, Vol. 88, issue 4, p. 848-881*. https://www.jstor.org/stable/117009

Ert, E., & Erev, I. (2007). Replicated alternatives and the role of confusion, chasing, and regret in decisions from experience. *Journal of Behavioral Decision Making, Vol. 20, issue 3, p. 305-322*. doi: https://doi.org/10.1002/bdm.556

Fehr, E., & Gachter, S. (2000). Cooperation and punishment in public goods experiments. *The American Economic Review, Vol. 90, No. 4, p. 980-994*. doi: https://www.jstor.org/stable/117319

Feige, C., & Ehrhart, K.-M. (2015). Voting and transfer payments in a threshold public goods game. *KT Working Paper Series in Economics, Vol. 73,*. doi: https://doi.org/10.5445/IR/1000050442%0A

Fudenberg, D., & Levine, D. (1998). Learning in games. *European Economic Review, Vol. 42, issue 3-5, p. 631-639*. doi: https://doi.org/10.1016/S0014-2921(98)00011-7

Gallier, C. (2017). Democracy and compliance in public goods games. *ZEW - Centre for European Economic Research Discussion Paper, Vol. 17, issue 38*. doi: http://dx.doi.org/10.2139/ssrn.3057268

Hardin, G. (1968). The tragedy of the commons. *Science*, *Vol. 162, issue 3859, p. 1243-1248*. doi: http://dx.doi.org/10.1126/science.162.3859.1243

Hasson, R., Lofgren, A., & Visser, M. (2010). Climate change in a public goods game: Investment decision in mitigation versus adaptation. *Ecological Economics, Vol. 70, issue 2, p. 331-338*. doi: https://doi.org/10.1016/j.ecolecon.2010.09.004

Hiraishi, K., Shikishima, C., Yamagata, S., & Ando, Y. (2015). Heritability of decisions and outcomes of public goods games. *Front. Psychol., Vol. 6-2015, sec. Decision Neuroscience*. doi: https://doi.org/10.3389/fpsyg.2015.00373

Ho, T. H., Camerer, C. F., & Chong, J. (2007). Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory, Vol. 133, issue 1, p. 177-198*. doi: http://dx.doi.org/10.1016/j.jet.2005.12.008

Isaac, R., McCue, K., & Plott, C. (1985). Public goods provision in an experimental environment. *Journal of Public Economics, Vol. 26, issue 1, p. 51-74*. doi: https://doi.org/10.1016/0047-2727(85)90038-6

Isaac, R., & Walker, J. (1988). Group size effects in public goods provision: The voluntary contribution mechanism. *Quarterly Journal of Economics, Vol. 103, p. 179-200*. doi: https://doi.org/10.2307/1882648

Isaac, R., Walker, J., & Williams, A. (1994). Group size and the voluntary provision of public goods: Experimental evidence utilizing large groups. *Journal of Public Economics, Vol. 54, issue 1, p. 1-36*. doi: https://doi.org/10.1016/0047-2727(94)90068-X

Kosfeld, M., Okada, A., & Riedl, A. (2009). Institution formation in public goods game. *American Economic Review, Vol. 99, issue 4, p. 1335-1355*. doi: https://doi.org/10.1257/aer.99.4.1335

Kraak, S. B. (2010). Exploring the public goods game model to overcome the tragedy of the commons in fisheries management. *Fish and Fisheries, Vol. 12, issue 1, p. 18-33*. doi: https://doi.org/10.1111/j.1467-2979.2010.00372.x

Kragt, V. d., Orbell, J., & Dawes, R. M. (1983). The minimal contributing set as a solution to public goods problems. *American Political Science Review, Vol. 77, issue 1, p. 112-122*. doi: http://dx.doi.org/10.2307/1956014

Ledyard, J. O. (1995). *Handbook of experimental economics: 2 public goods: A survey of experimental research*. Princeton University Press. doi: https://doi.org/10.2307/j.ctvzsmff5.6

Marwell, G., & Ames, R. E. (1979). Experiments on the provision of public goods. i. resources, interest, group size, and the free-rider problem. *American Journal of Sociology, Vol. 84, issue 6, p. 1335-1360*. doi: https://www.jstor.org/stable/2777895

Merrett, D. (2012). Estimation of public goods game data. *Working Papers, University of Sydney, School of Economics, Vol. 09-2012*. https://econ-wpseries.com/2012/201209.pdf

Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings Nat. Academy Sciences USA, Vol. 36, p. 48-49*. doi: https://doi.org/10.1073/pnas.36.1.48

Nash, J. F. (1951). Noncooperative games. *Annals Math., Vol. 54, p. 289-95*. doi: https://www.jstor.org/stable/1969529

Nax, H., Burton-Chellew, M., West, S., & Young, H. (2016). Learning in a black box. *University of Oxford, Department of Economics Discussion Paper Series, Vol. 653*. doi: https://doi.org/10.2139/ssrn.2664646.

Nax, H., & Perc, M. (2015). Directional learning and the provisioning of public goods. *Scientific Reports, Vol. 5, article 8010*. doi: https://ssrn.com/abstract=2555610

Neugebauer, T., Perote, J., Schmidt, U., & Loos, M. (2007). Selfish-biased conditional cooperation: On the decline of contributions in repeated public goods experiments. *Journal of Economic Psychology, Vol. 30, issue 1, p. 52-60*. doi: https://doi.org/10.1016/j.joep.2008.04.005

Neumann, J., von. (1928). Zur theories der gesellschaftsspiele. *Mathematische Annalen, Vol. 100, p. 295-320*.

Neumann, J., von, & Morgenstern, O. (1944). *Theory of games and economic behaviour.* Princeton University Press. doi: https://doi.org/10.1515/9781400829460

Nunn, G., & Watkins, T. (1978). Public goods games. *Southern Economic Journal, Vol. 45, issue 2, p. 598-606*. doi: https://doi.org/10.2307/1057688

Olson, M. (1974). *The logic of collective action.* Harvard university press.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action.* Cambridge University Press.

Renton, J., & Page, K. (2021). Cooperative success in epithelial public goods games. *Journal of Theoretical Biology, Vol. 528, article 110838*. doi: https://doi.org/10.1016/j.jtbi.2021.110838

Salahshour, M. (2021). Evolution of cooperation and consistent personalities in public goods games. *Sci Rep, Vol. 11, article 23708*. doi: https://doi.org/10.1038/s41598-021-03045-w

Samuelson, P. A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics, Vol. 36, issue 4, p. 387-389*. doi: https://doi.org/10.2307/1925895

Schelling, T. (1960). *The strategy of conflict*. Harvard University Press.

Sollow, J., & Kirkwood, N. (2000). Group identity and gender in public goods experiments. *Journal of Economic Behavior & Organization, Vol. 48, issue 4, p. 403-412*. doi: https://doi.org/10.1016/S0167-2681(01)00243-8

Tomassini, M., & Antonioni, A. (2020). Public goods games on coevolving social network models. *Front. Phys., Vol. 8, issue 58*. doi: https://doi.org/10.3389/fphy.2020.00058