SOCIAL PREFERENCES, LEARNING, AND THE DYNAMICS OF COOPERATION IN NETWORKED SOCIETIES: A DIALOGUE BETWEEN EXPERIMENTAL AND COMPUTATIONAL APPROACHES

by

Chenna Reddy Cotla
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
In Partial fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Computational Social Science

Committee:

_________________________ Dr. Robert Axtell, Dissertation Chair

_________________________ Dr. Ragan Petrie, Committee Member

_________________________ Dr. Daniel Houser, Committee Member

_________________________ Dr. Qing Tian, Committee Member

_________________________ Dr. Kevin Curtin, Acting Department Chair

_________________________ Dr. Donna M. Fox, Associate Dean,
Office of Student Affairs & Special Programs,
College of Science

_________________________ Dr. Peggy Agouris, Dean, College of Science

Date: _________________________ Summer 2016
George Mason University
Fairfax, VA

Social Preferences, Learning, and the Dynamics of Cooperation in Networked Societies: A Dialogue Between Experimental and Computational Approaches

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

Chenna Reddy Cotla
Master of Technology
Indian Institute of Information Technology and Management, 2008
Bachelor of Technology
Indian Institute of Information Technology and Management, 2006

Director: Dr. Robert Axtell, Professor
Department of Computational and Data Sciences

Summer 2016
George Mason University
Fairfax, VA

Dedication

I dedicate this dissertation to my mother Janakamma Cotla.

Acknowledgments

First, I would like to thank my mother, Janakamma, and my sister, Neeraja, for their love, patience, support, and encouragement. Second, I would like to thank my co-advisors Dr. Rob Axtell and Dr. Ragan Petrie for making this dissertation possible through their exceptional insight, patience, and advice. Rob introduced me to the research methodology of agent-based modeling and guided my research throughout my career as a PhD student. Ragan encouraged my interest in pursuing research at the intersection of agent-based modeling and experimental economics for my doctoral dissertation. She introduced me to the world of experimental economics and nurtured my interest in empirical research. I cannot thank her enough for the time she spent discussing and reviewing my research work, responding to my questions, and being my mentor at Interdisciplinary Center for Economic Science (ICES). Third, I thank my dissertation committee members Dr. Daniel Houser and Dr. Qing Tian for their professional support and the feedback on my dissertation research. Last but not the least I am greatly indebted to Dr. Marco Castillo for offering very helpful suggestions on my research, providing numerous learning opportunities, and mentoring me during the final stages of my PhD.

I would like to extend my gratitude to Krasnow Institute of Advanced Study for funding me during the first year, Center for Social Complexity (CSC) for funding me during the second and third years, and ICES for funding me during the final years of my PhD. I would like to thank Claudio Cioffi-Revilla, Bill Kennedy, and Tim Gulden at CSC for the chance to work with them and the mentoring I have received. Working on several research projects with Dr. Marco Castillo and Dr. Ragan Petrie at ICES was a tremendous learning experience for me and it directly influenced my dissertation research.

Finally, I would like thank my personal friends, many wonderful teachers I was fortunate to have during my time as a student, and my colleagues at ICES and Computational Social Science. Special thanks go to Vishnu Karanam, Aditya Trivedi, Karen Underwood, Omar Guerrero, Ates Hailegiorgis, Ernesto Carrella, and Ahrash Dianat.

Table of Contents

References . . . 205

List of Tables

List of Figures

Abstract

SOCIAL PREFERENCES, LEARNING, AND THE DYNAMICS OF COOPERATION IN NETWORKED SOCIETIES: A DIALOGUE BETWEEN EXPERIMENTAL AND COMPUTATIONAL APPROACHES

Chenna Reddy Cotla, PhD

George Mason University, 2016

Dissertation Director: Dr. Robert Axtell

In this dissertation, I empirically investigate cooperative behavior in networks using the framework of network public goods games. To do so, I use a dialogue between behavioral experiments and agent-based models. I design and conduct behavioral experiments to generate data to construct boundedly rational agents that behave like humans and reproduce stylized facts in public goods environments. The human-like agents are deployed in a small-scale agent-based model to make novel quantitative predictions that can be statistically tested using a new set of behavioral experiments. This ensures that the behavioral specification of agents carries predictive value so that quantitative predictions made using it can be reproduced with human subject experiments. The high fidelity agent-based model is then extended to study the dynamics cooperation in networked environments. The dissertation is organized into three chapters.

In the first chapter, using experimental data from a number of published studies on repeated public goods games, I examine the ex-post descriptive fit and ex-ante predictive accuracy of several learning models. I show that choices in repeated public goods games are best explained by an averaging reinforcement learning model. The second chapter builds on the experimental evidence that social preferences are necessary alongside with learning to explain contribution patterns in repeated public goods games. I design and conduct novel behavioral experiments to disentangle the roles of social preferences and learning in explaining repeated game choices. I find that choices in the repeated public goods games are best described by social preferences affecting the choice of first round contributions and then subsequent contributions based on payoff-based averaging reinforcement learning. I deploy the behavioral specification thus obtained in an agent-based model. Simulations using the agent-based model demonstrate a novel result that reducing the price of cooperation in the first round alone is sufficient to sustain significantly higher contributions over the later rounds of a repeated game. The quantitative predictions of the empirical agent-based model are successfully reproduced using follow-up behavioral experiments substantiating the predictive value of agents' behavioral specification.

In the third chapter, using the empirical agent-based model constructed from the experimental data, I show that network size, network density, degree heterogeneity, and the average path length of a network have no significant effect on the cooperation levels in public goods games. These results stand in contrast to the existing findings in the agent-based modeling literature and demonstrate that agent-based models based on empirical micro-foundations can lead to different conclusions than that of agent-based models based on micro-foundations extrapolated from other domains like that of biology. This dissertation illustrates that empirical understanding of determinants of behavior in social dilemmas is instrumental to identify mechanisms that can promote cooperation using agent-based models.

Chapter 1: Learning in Repeated Public Goods Games - A Meta Analysis

Abstract

I examine the generalizability of a broad range of prominent learning models in explaining contribution patterns in repeated linear public goods games. Experimental data from twelve previously published papers is considered in testing several learning models in terms of how accurately they describe individuals' round-by-round choices. The experimental data is split into 18 datasets. Each of these datasets is different from the remaining in at least one of the following aspects: the marginal per capita return, group size, matching protocol, number of rounds, and endowment that determines the number of stage game strategies. Both ex-post descriptive fit of learning models and their ex-ante predictive accuracy are examined. The following learning models are included in the study: reinforcement learning, normalized reinforcement learning, reinforcement average model with loss aversion strategy (REL), stochastic fictitious play, normalized stochastic fictitious play, experience weighted attraction learning (EWA), self-tuning EWA, and Impulse matching learning. REL out-performs all other learning models in both within dataset descriptive fit and out-of-sample dataset predictive accuracy. While all the learning models out-perform the random choice benchmark, only REL performs at least as well as the model that reflects dataset level overall empirical frequencies. The results suggest that learning in repeated linear public goods games is more in line with reinforcement learning than that of belief learning or regret-based learning. Finally, REL also outperforms individual evolutionary learning (IEL) in predicting the full distribution of contributions. Average reinforcement learning that is sensitive to the observed payoff variability and insensitive to the payoff magnitude underlie the success of REL in explaining contributions in repeated public goods games over a broad spectrum of game parameters.

1.1 Introduction

In this Chapter, I use data from controlled lab experiments and standard econometric methods to compare a number of competing behavioral specifications of agents that are specified as different learning models in the context of repeated public goods games. The best performing learning model identified here will be used for further analysis in conjunction with social preferences in Chapter 2. The analysis here and Chapter 2 will inform the behavioral specification of agents in an empirical agent-based model of public goods games in networked communities in later chapters. Thus, the investigations and results in this Chapter create a foundation over which the later chapters of the dissertation are developed.

A handful of earlier studies have used data from laboratory experiments on public goods games to test competing learning models (Cooper & Stockman, 2002; Janssen & Ahn, 2006; Arifovic & Ledyard, 2012; Wunder, Suri, & Watts, 2013). However, the approaches used in these studies have limitations in terms of model generalizability either due to the data used or the set of learning models considered or both. Cooper and Stockman (2002) use the cumulative reinforcement learning model of Roth and Erev (1995) to model learning behavior in step level public goods games. They have used the data from their own experiments to fit the model and derive conclusions. Previous research shows that cumulative reinforcement learning is outperformed by average reinforcement learning in predicting behavior (Erev & Roth, 1998; Ho, Camerer, & Chong, 2007). Arifovic and Ledyard (2012) and Janssen and Ahn (2006) consider data from multiple experiments conducted by Mark Isaac and his coauthors in studying learning in public goods games. However, similar to Cooper and Stockman (2002), they also considered only one learning algorithm in their studies. The former exclusively considers their own Individual Evolutionary Learning (IEL) model while the latter considers only Experience Weighted Attraction (EWA) learning model of Camerer and Ho (1999). Both of these papers, examine learning models in terms of their within sample descriptive fit.Thus, these analyses are subject to the potential problem of model overfitting. Wunder et al. (2013) consider a number of learning models, however, they take a purely machine learning approach by considering models that are difficult to interpret from the cognitive plausibility perspective. Furthermore, they consider data from the experiments conducted in Suri and Watts (2011) which involve public goods game with only one set of game parameters. While their stochastic discounted two-factor model does very well in both descriptive fit and predictive accuracy compared to a number of other learning models considered in their paper, it is unclear if the model will also be successful in explaining behavior in public goods games with different values for the game parameters like the endowment, marginal return, number of rounds, and matching. This is a valid concern since previous studies have supported different models of learning in different games and the relationship among distinct conclusions obtained in different studies is often not clear (Erev & Haruvy, 2013). Therefore, it remains an open question if any of the models that were considered in these earlier studies are general enough to explain behavior in public goods environments with wide ranging parameters.

The obstacles faced in understanding the learning mechanism individuals use in repeated public goods environment, or for that matter in any repeated strategic environment, using data from experiments can be listed as follows:

1. **Small datasets due to the small number of experimental subjects**: Small sized data sets due to the small number of experimental subjects can induce selection biases in econometric investigations. Often, datasets that are generated using controlled lab experiments contain individuals in sub hundreds. They may not have enough variation to compare competing learning models. Furthermore, an experimental dataset considered in isolation may consist of individuals drawn from a very narrow spectrum of the population. Thus, a model that is favored by a single data set may not generalize over.

2. **Small number of rounds of in repeated games:** Small number of rounds could potentially lead to misleading conclusions regarding which learning model performs best. Learning models generally require sufficiently large number of rounds before learning process unfolds. When there are data for only a couple of rounds learning

models' performance could depend strongly on initial conditions. Different learning models could be made favorites by cleverly choosing suitable initial conditions.

3. **The problem of model overfitting**: Salmon (2001) and Hopkins (2002) have shown that learning models can be very easily overfitted to the experimental data. The problem of overfitting is exacerbated when one is dealing with small data set sizes and a small set of learning models (Erev, Ert, & Roth, 2010).

4. **Consideration of a relatively small set of models**: Since every learning model is a simplified representation of an underlying cognitive or neural mechanism, it is fair to say each of them is misspecified. If one considers only a very small set of learning models, it is difficult to get a relative sense of what constitutes a good learning description in a given environment. Therefore, the choice of the set of learning models that one chooses to study can influence what conclusions one reaches. This point can also be demonstrated using the results in this Chapter. If one considered only belief learning models, one could potentially argue, using the results in this Chapter, that learning does not perform better than a random choice model and thus is not useful in predicting behavior in public goods environments. By considering a large set of learning models that span across different families of learning, one can reduce the potential model selection bias.

5. **Choice of model comparison metric:** The choice of a metric that one uses to compare the performance of different learning models can potentially influence the conclusion one reaches. Feltovich (2000) studies the performance of reinforcement and belief learning models in the context of multi stage asymmetric information games. When he compares models using maximum log-likelihood achieved by them, a belief learning model outperforms a set of other models that includes a reinforcement learning model. However, when he considers Mean Squared Deviation (MSD) as the comparison metric, the reinforcement learning model outperforms the belief learning model. In the work of Chmura, Goerg, and Selten (2012), the authors find that impulse matching

learning outperforms a number of other models in replicating aggregate frequencies of observed choices, but, self-tuning EWA performs the best in predicting individuals' round-by-round behavior in twelve 2×2 games. Thus, the choice of comparison metric can have significant influence on the conclusions one reaches.

In this chapter I address these obstacles in understanding learning in public goods games by:

1. Considering a large dataset involving 1201 individuals and 17250 decisions that is constructed using the experimental data from 12 different published studies. These studies are conducted with subjects from different nations, educational backgrounds, and are conducted by different researchers. This reduces the potential data selection bias in my investigation. As Erev and Roth (1998) states, there is danger that researchers treat the models they propose as their toothbrushes, by using one's model on one's own data. Experimenters may also unconsciously make some decisions in their experimental design that could favor some learning models over others. My study seeks to overcome this danger by considering data from multiple experiments.

2. Considering data on public goods games that involve at least 10 rounds of play and thus sufficiently allowing the learning process to unfold.

3. Comparing models using both descriptive fit and prediction accuracy. I divide the experimental data on 1201 individuals into 18 separate datasets each of which involves a public goods game with a unique set of game parameters. I compare how generalizable each learning model is in two ways. First, using the parameter estimates of each learning model estimated from the pooled data of 1201 individuals, I test how well the same parameters explain behavior across datasets with varying parameters of the game. Second, for each dataset I predict choices in it using each learning model with parameters estimated using the remaining 17 datasets. I compare models based on how well they do in the prediction task. In both of these comparisons testing the generalizability of models, I use only non-parametric exact tests.

4. Considering nine learning models that cover a broad spectrum of types of learning. The following learning models are considered to model individual level adaptive behavior: reinforcement learning (RL), normalized reinforcement learning (NRL), reinforcement average model with loss aversion strategy (REL), stochastic fictitious play (SFP), normalized stochastic fictitious play (NFP), experience weighted attraction learning (EWA), self-tuning EWA (STEWA), Impulse matching learning (IM), and Individual Evolutionary Learning (IEL). These models cover the ideas of reinforcement learning, belief learning, regret-based learning, hybrid learning that incorporate elements from both belief and reinforcement learning, and individual evolutionary learning. This choice of wide ranging models, helps to minimize methodological bias that can enter into the analysis due to the consideration of a very narrow set (often of size one or two) of learning models.

5. Using a wide ranging comparison metrics in analyzing the effectiveness of learning models in replicating the observed behavior in the data. I consider log-likelihood (LL), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and Mean Quadratic Score (MQS)[1] achieved by learning models in comparing their descriptive fit and prediction accuracy. I also compare the most successful model REL with the evolutionary learning model IEL using the distance between the observed and predicted distribution of choices.

The main finding in this study is that REL outperforms all other learning models in terms of both descriptive fit and prediction accuracy. It is the only model to outperform dataset level overall empirical frequencies in both descriptive fit and predictive accuracy when model performance is compared using MQS. In terms of LL and BIC, REL performs as well as the empirical frequencies in both comparisons. None of the remaining models outperform the empirical frequencies in both comparisons according to any of the metrics. EWA and NRL jointly take the second place. They are outperformed by the dataset level

[1]Note that MQS is equivalent to Mean Squared Deviation (MSD) by an affine transformation. MQS = 1 - MSD for one dimensional distributions.

empirical frequencies in terms of LL and BIC. But, they do as well as the dataset level empirical frequencies in terms of MQS. The ordering of learning models in both comparisons is identical: REL, NRL, EWA, RL, NSFP, SFP, IM, STEWA (the ordering is not strict in all cases. REL performs strictly better than all other models). Belief learning models and regret-learning models perform poorly in comparison to reinforcement models. In public goods environments, learning is more in line with reinforcement learning. Average reinforcement learning that is adaptive to observed payoff variability and insensitive to the payoff magnitude underlies the success of REL. REL also outperforms IEL in predicting full distribution of contributions across all data sets.

The chapter is organized as follows. In Section 1.2, I describe the linear public goods environment. Section 1.3 describes the data that is used in econometric investigations. Section 1.4 presents a brief description of learning models. Section 1.5 presents results. A brief discussion is presented in Section 1.6. Section 1.7 concludes the chapter by summing up the findings and discussing future directions.

1.2 Linear Public Goods Game

Before I proceed to the data and learning models that are considered in my analysis, I first briefly describe the linear public goods environment. A linear public good game consists of N_g individuals. Index these individuals as $i = 1, 2, .., N_g$. Each individual is endowed with an amount w_i and must decide how much to contribute to the public good from his endowment and how much he wants to keep to himself. In this paper, I study public goods games where each individual is endowed with the same amount $w_i = E, \forall i$. Say an individual's contribution is denoted as $c_i \in [0, E]$, then his payoff from his contribution decision is given as:

$$\pi^i = E - c_i + M \sum_{j=1}^{N_g} c_j$$

Where M is the marginal per capita return (MPCR) from the public good. As long as $\frac{1}{N_g} < M < 1$, it is individually rational to contribute nothing to the public good but the social optimum is achieved when everybody contributes their entire endowment. In this way, a linear public goods game epitomizes the tension between private and public interests that is the basis of any standard social dilemma.

In a repeated game, the stage described above is repeated over T number of rounds. In a finitely repeated game, the number of rounds, T, is determined beforehand. It is easy to see that zero contribution in each round is individually rational in a finitely repeated public goods game.

1.3 Data

A number of published studies involving repeated public goods game experiments with at least 10 rounds were considered for data collection. Since the main aim of this study is to evaluate different learning models in terms of their descriptive and predictive power in explaining choices over time, I considered only the baseline studies where no additional mechanism like punishment was in place. For the studies where the data was not publicly made available at the time of data collection, data was requested from the corresponding authors. Data was obtained for experiments conducted in the following twelve published studies: Andreoni (1988, 1995b, 1995a); Isaac and Walker (1988); Isaac, Walker, and Williams (1994); Fehr and Gächter (2000); Nikiforakis and Normann (2008); Gächter, Renner, and Sefton (2008); Sefton, Shupp, and Walker (2007); Kosfeld, Okada, and Riedl (2009); Botelho, Harrison, Pinto, and Rutström (2009); Keser and Van Winden (2000). While a few more studies had data on baseline repeated public goods games of length 10 or more rounds, the data was not readily available at the time of this study.

The experimental data from 12 published papers under consideration is reorganized into 18 datasets. Data from identical treatments across studies is combined and put into a single dataset to obtain 18 distinct datasets that are different from each other in terms of at least one parameter of the public goods game itself. The parameters of the game that are used to distinguish among datasets are: the number of rounds of the repeated game (T), group size (N_g), endowment (E), marginal per capita return (M), and matching protocol. The matching protocol can be Partners, Strangers, or Perfect Strangers. In Partners matching, group composition stays fixed over the rounds of a repeated game. In Strangers' matching, groups are reshuffled in each round using the subject pool a given experimental session. However, Strangers' matching does not guarantee that any two subjects will not be matched more than once over the rounds of a repeated game. Perfect Strangers' matching is identical to the Strangers' matching except that it ensures that no two subjects will be matched more than once over the rounds of a repeated game. Table 1.1 presents 18 datasets that are constructed from the data obtained from published studies. No two datasets contain the same individual. These datasets cover a broad range of game parameters. For example, E ranges from 6 to 60, M varies between 0.03 to 0.8, N_g varies between 3 to 40, T ranges from 10 to 50. All three matching protocols are observed across the datasets.

1.4 Learning Models

In this section, I briefly describe each learning model that is considered in this Chapter. In total, I consider eight learning models.[2] These eight models cover a broad spectrum of learning sophistication. Models that use reinforcement learning occupy the lower end of the learning sophistication spectrum since they characterize a very basic principle of reinforcing actions that led to higher payoffs in previous rounds. I consider three variations of reinforcement learning: Reinforcement Learning Model (RL), Normalized Reinforcement Learning Model (NRL), and Reinforcement Average Model with a Loss Aversion Strategy (REL). Belief learning models and regret-based learning models rank high on learning sophistication.

[2]in the Discussion section of this chapter, I also consider Individual Evolutionary Learning (IEL) which makes the count nine. I do not consider IEL in this Section and the Results section since maximum likelihood methods cannot be used in the estimation of IEL.

Table 1.1: Data

Dataset	Source	T	N_g	E	M	N	Matching
DS1	Andreoni 1988	10	5	50	0.5	40	Strangers
DS2	Andreoni 1988	10	5	50	0.5	30	Partners
DS3	Andreoni 1995a, 1995b	10	5	60	0.5	80	Strangers
DS4	Isaac & Walker 1988	10	10	10	0.75	30	Partners
DS5	Isaac & Walker 1988	10	10	25	0.3	30	Partners
DS6	Isaac & Walker 1994	10	40	50	0.3	120	Partners
DS7	Isaac & Walker 1994	10	40	50	0.03	40	Partners
DS8	Fehr & Gachter (2000), Nikiforakis (2008)	10	4	20	0.4	96	Partners
DS9	Fehr & Gachter (2000), Nikiforakis (2008)	10	4	20	0.4	72	Strangers
DS10	Gachter et al. (2008)	10	3	20	0.5	60	Partners
DS11	Gachter et al. (2008)	50	3	20	0.5	51	Partners
DS12	Sefton et al. (2007)	10	4	6	0.5	144	Partners
DS13	Kosfield et al. (2009)	20	4	20	0.4	40	Partners
DS14	Kosfield et al. (2009)	20	4	20	0.65	36	Partners
DS15	Botelho et al. (2009)	10	4	20	0.8	56	Strangers
DS16	Botelho et al. (2009)	10	4	20	0.8	116	Perfect Strangers
DS17	Keser & van Winden(2000)	25	4	10	0.5	40	Partners
DS18	Keser & van Winden(2000)	25	4	10	0.5	120	Strangers
$Ntotal$: No. of Total Subjects						1201	
No. of Total Decisions						17250	

Belief learning and regret-based learning models capture a notion of counterfactual thinking by considering foregone payoffs from unobserved strategies. Hybrid models like that of Experience Weighted Attraction (EWA) learning model of Camerer and Ho (1999) and Self-tuning Experience Weighted Attraction learning (STEWA) of Ho et al. (2007) combine the elements of both reinforcement and belief learning.

All of the learning models presented below have two elements in common. They assume that strategies have propensities and propensities are linked to the probabilities of choice using a choice rule. The learning algorithm of a given model dictates how the propensities of strategies are updated. Propensities of strategies are often referred to as attractions. Given this framework, each learning model can be succinctly described by defining the specification of initial attractions, attraction updating rule, and choice rule. All investigated learning models will start with a randomization of available strategies in the first round. This is done by setting initial attractions to zero. Alternatively, one could initialize attractions to be equal to the payoff from a random choice. However, in my estimations it did not make any difference. One could also estimate the initial attractions based on the first round data. This would improve the fit of the learning model since it provides a better description of the first round. For games with large strategy spaces like that of public goods games, estimating the first round attractions would mean estimating too many parameters. To maintain model parsimony, I did not pursue estimating the first round attractions. A brief description of each learning model follows in the remainder of this section.

I refer to strategy j available to individual i as s_i^j. There are m_i strategies available to individual i. The chosen strategy of i in round t is denoted as $s_i(t)$. $u_i(s_i^j, s_{-i}(t))$ is the payoff obtained by i in round t when it chooses s_i^j and others choose strategies given in $s_{-i}(t)$. The attraction of strategy j of agent i in round t is denoted as $A_i^j(t)$.

1.4.1 Reinforcement Learning Model (RL)

This is the variation of reinforcement learning model considered in Erev, Roth, Slonim, and Barron (2007). It differs from its earlier version proposed in Erev and Roth (1998) in the way attractions are related to choice probabilities. In the earlier version proposed in Erev and Roth (1998), attractions are translated to choice probabilities via simple normalization, whereas, in the version considered in Erev et al. (2007) choice probabilities are computed from attractions using a logit function as described below. The three components of this learning model are described below:

- **Initial Attractions:** In the first round, attractions of all strategies are zero. Therefore, $A_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen.

- **Attraction Updating:** Strategy j of agent i has attraction $A_i^j(t)$ in round t:

$$A_i^j(t) = [1 - w.\mathbb{I}(s_i^j, s_i(t-1))].A_i^j(t-1) + [w.\mathbb{I}(s_i^j, s_i(t-1))].u_i(s_i^j, s_{-i}(t-1))$$

  Where $\mathbb{I}(s_i^j, s_i(t-1))$ is an Indicator function and equals 1 when $s_i^j = s_i(t-1)$. Otherwise it is zero. $u_i(s_i^j, s_{-i}(t-1))$ is the payoff obtained when the agent chooses strategy s_i^j given other have chosen strategies given by $s_{i-1}(t-1)$ in round $(t-1)$.

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\lambda A_i^j(t)}}{\sum_{l=1}^{m_i} e^{\lambda A_l^j(t)}}$$

There are two free parameters to be estimated in the RL model: the attraction sensitivity parameter λ and the attraction weighting parameter w. For large λ, strategies with higher attractions are chosen with a higher probability where as when $\lambda \to 0$ choice rule is equivalent to a random choice.

1.4.2 Normalized Reinforcement Learning Model (NRL)

Normalized Reinforcement Learning (Erev et al., 2007) is similar to the RL model described earlier with one exception: attraction sensitivity is adaptive and decreases with increasing observed payoff variability. When observed payoff variability increases decision making moves closer to the random choice behavior.

The three components of this learning model are as below:

- **Initial Attractions:** In the first round, attractions of all strategies are zero. Therefore, $A_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen.
- **Attraction Updating:** Strategy j of agent i has attraction $A_i^j(t)$ in round t:

$$A_i^j(t) = [1 - w.\mathbb{I}(s_i^j, s_i(t-1))].A_i^j(t-1) + [w.\mathbb{I}(s_i^j, s_i(t-1))].u_i(s_i^j, s_{-i}(t))$$

  Where $\mathbb{I}(s_i^j, s_i(t-1))$ is an Indicator function and equals 1 when $s_i^j = s_i(t-1)$. Otherwise it is zero.

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\frac{\lambda}{PV_i(t)} A_i^j(t)}}{\sum_{l=1}^{m_i} e^{\frac{\lambda}{PV_i(t)} A_l^j(t)}}$$

  Where $PV_i(t)$ is the measure of observed payoff variability for agent i. $PV_i(0)$ is set equal to λ. In each round $PV_i(t)$ is updated as below:

$$PV_i(t) = (1-w).PV_i(t-1) + w.|max\{recent_i^1, ..., recent_i^m\} - u_i(s(t-1), s_{-i}(t-1))|$$

Where $recent_i^j$ is the most recent observed payoff from choosing the strategy s_i^j. Before the first observation of the payoff of strategy s_i^j, $recent_i^j = A(1)$.

There are two free parameters to be estimated in the RL model: the attraction sensitivity parameter λ and the attraction weight parameter w.

1.4.3 Reinforcement Average Model with a Loss Aversion Strategy (REL)

The REL model was proposed by Erev, Bereby-Meyer, and Roth (1999) to circumvent the problems faced by the reinforcement model of Erev and Roth (1998) in explaining behavior in games when a constant is added to all payoffs. It introduced two modifications to the original reinforcement model: sensitivity to payoff variability and insensitivity to payoff magnitude. The model is similar to the NRL model described above in the sense that attraction sensitivity is adaptive and decreases with increasing observed payoff variability. However, unlike the NRL model, it accounts for the insensitivity for payoff magnitude by initializing the first round payoff variability as the expected absolute difference between the observed payoff from random choice and the average payoff from random choice. The REL model, though originally was proposed to explain the choice behavior in the presence losses, was observed to provide a better account for the games considered in Erev and Roth (1998) that did not involve losses.

The three components of this learning model are as below:

- **Initial Attractions:** In the first round, attractions of all strategies are zero. Therefore, $A_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen.

- **Attraction Updating:** Strategy j of agent i has attraction $A_i^j(t)$ in round t:

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

where $C_i^j(t)$ is the number of times s_i^j has been chosen in the first t rounds and $N(1)$ is a free parameter that determines the strength of the initial attractions. A large $N(1)$ means that effect of actual payoffs in later rounds on attractions will be smaller. The attractions of unchosen strategies are not updated.

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\frac{\lambda}{PV_i(t)} A_i^j(t)}}{\sum_{l=1}^{m_i} e^{\frac{\lambda}{PV_i(t)} A_l^j(t)}}$$

where λ_i is a free parameter that determines the reinforcement sensitivity of the individual i. $PV_i(t)$ is the measure of payoff variability.

The payoff variability is updated according to:

$$PV_i(t) = \frac{[PV_i(t-1)(t-1+m_iN(1)) + |u_i(s_i(t-1), s_{-i}(t-1)) - PA_i(t-1)|]}{[t+m_iN(1)]}$$

where $PA_i(t)$ is the accumulated payoff average in round t and m is the number of strategies. $PV_i(1) > 0$ is initialized as the the expected absolute difference between the obtained payoff from a random choice and the average payoff from a random choice.

$PA_i(t)$ is calculated in a similar manner.

$$PA_i(t) = \frac{[PA_i(t-1)(t-1+m_iN_i(1)) + u_i(s_i(t-1), s_{-i}(t-1))]}{[t+m_iN_i(1)]}$$

I initialized $PA_i(1)$ as initial attraction of a any strategy which is 0. There are two free parameters to be estimated in the REL model: the attraction sensitivity parameter λ and the strength of initial attractions $N(1)$.

1.4.4 Stochastic Fictitious Play (SFP)

Stochastic Fictitious Play (SFP) (Fudenberg & Levine, 1998; Cheung & Friedman, 1997; Cooper, Garvin, & Kagel, 1997) is a prototypical characterization of belief learning. Here I consider the variation of the model described in Erev et al. (2007).

- **Initial Attractions:** In the first round, attractions of all strategies are zero. Therefore, $A_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen.

- **Attraction Updating:** Strategy j of agent i has attraction $A_i^j(t)$ in round t:

$$A_i^j(t) = [1 - w].A_i^j(t-1) + w.u_i(s_i^j, s_{-i}(t-1))$$

  $u_i(s_i^j, s_{-i}(t-1))$ is the payoff obtained when the agent chooses strategy s_i^j given other have chosen strategies given by $s_{-i}(t-1)$ in round $(t-1)$.

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\lambda A_i^j(t)}}{\sum_{l=1}^{m_i} e^{\lambda A_l^j(t)}}$$

There are two free parameters to be estimated in the SFP model: the attraction sensitivity parameter λ and the attraction weight parameter w.

1.4.5 Normalized Fictitious Play (NFP)

Normalized Fictitious Play (NFP) is similar to the SFP with the exception of the scaled attractions sensitivity. I follow its description in Ert and Erev (2007):

- **Initial Attractions:** In the first round, attractions of all strategies are zero. Therefore, $A_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen.

- **Attraction Updating:** Strategy j of agent i has attraction $A_i^j(t)$ in round t:

$$A_i^j(t) = [1 - w].A_i^j(t-1) + w.u_i(s_i^j, s_{-i}(t-1))$$

  $u_i(s_i^j, s_{-i}(t-1))$ is the payoff obtained when the agent chooses strategy s_i^j given other have chosen strategies given by $s_{-i}(t-1)$ in round $(t-1)$.

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\frac{\lambda}{PV_i(t)} A_i^j(t)}}{\sum_{l=1}^{m_i} e^{\frac{\lambda}{PV_i(t)} A_l^j(t)}}$$

  Where $PV_i(t)$ is the measure of observed payoff variability for agent i. $PV_i(0)$ is set equal to λ. In each round $PV_i(t)$ is updated as below:

$$PV_i(t) = (1-w).PV_i(t-1) + w.|max\{recent_i^1, ..., recent_i^m\} - u_i(s_i(t-1), s_{-i}(t-1))|$$

  Where $recent_i^j$ is the most recent observed payoff from choosing the strategy s_i^j. Before the first observation of the payoff of strategy s_i^j, $recent_i^j = A(1)$.

There are two free parameters to be estimated in the NFP model: the attraction sensitivity parameter λ and the attraction weight parameter w.

1.4.6 Experience Weighted Attraction Learning (EWA)

EWA introduced by Camerer and Ho (1999) is a hybrid model that encompasses a non-linear combination of reinforcement learning and belief learning. In this way, it is more sophisticated than the models described earlier. The three components of EWA learning model are described below.

- **Initial Attractions:** In the first round, attractions of all strategies are zero. Therefore, $A_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen.

- **Attraction Updating:** Strategy j of agent i has attraction $A_i^j(t)$ in round t:

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

  Where $\mathbb{I}(s_i^j, s_i(t-1))$ is an Indicator function and equals 1 when $s_i^j = s_i(t-1)$. Otherwise it is zero.

  Each attraction is applied an experience weight using:

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

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\lambda A_i^j(t)}}{\sum_{l=1}^{m_i} e^{\lambda A_l^j(t)}}$$

In total EWA has five parameters: $\lambda, \delta, \phi, \kappa, \eta$. λ is the attraction sensitivity parameter and when it approaches zero the decision making approaches a random choice rule. $\phi \in [0,1]$ represents a notion of forgetting. A large ϕ implies that previous attractions are highly discounted and will not have a large effect on future decision making. $\kappa \in [0,1]$ characterizes if the attraction updating is cumulative or averaging. $\eta \in [0,1]$ enters into the computation of initial experience weight $N(0)$. Finally, $\delta \in [0,1]$ captures the relative weight given to the foregone payoffs compared to the actual payoffs in updating attractions. According to Camerer and Ho (1999), δ in EWA helps it to capture both the law of actual effect as in reinforcement learning and law of simulated effect as in belief learning. For different constellations of parameters, EWA takes the form of different well known learning models. When $\delta = 0, \phi = 1, \kappa = 1$ EWA is equivalent to the cumulative reinforcement model (Roth & Erev, 1995; Bush & Mosteller, 1955; Cross et al., 2008; Arthur, 1991). EWA is like the average reinforcement model of Erev and Roth (1998) when $\delta = 0, \phi = 1, \kappa = 0$. When $\delta = 1, \phi = 0, \kappa = 1$ EWA is equivalent to the Cournot learning model. EWA is identical to weighted fictitious play when $\delta = 1, \phi = 1, \kappa = 0$.

1.4.7 Self-Tuning EWA Learning (STEWA)

The STEWA learning was introduced by Ho et al. (2007) as a one-parameter variation of the original EWA model of Camerer and Ho (1999). STEWA fixes the parameters $\{\kappa, N(0)\}$ of the EWA model at $\kappa = 0$, $N(0) = 1$. By fixing $\kappa = 0$, STEWA considers only average updating of attractions. It also replaces the two other parameters ϕ and δ of the EWA model with functionals $\phi(t)$ and $\delta(t)$ respectively. The motivation behind STEWA was to circumvent the problem of overfitting in EWA due to the large number of estimable parameters. Ho et al. (2007) show that STEWA outperformed EWA in a cross-validation tests using data from Mixed strategy, Patent race, Continental divide, Median action, Pot games, p-Beauty Contest and Price matching games. I consider STEWA among other models in this paper because of its success in out-of-sample prediction across wide range of games. The learning model is described below.

- **Initial Attractions:** In the first round, attractions of all strategies are zero. Therefore, $A_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen.

- **Attraction Updating:** Strategy j of agent i has attraction $A_i^j(t)$ in round t:

$$A_i^j(t) = \frac{\phi_i(t).N_i(t-1).A_i^j(t-1) + [\delta_i^j(t) + (1-\delta_i^j(t)).\mathbb{I}(s_i^j, s_i(t-1))].u_i(s_i^j, s_{-i}(t-1))}{N_i(t)}$$

Where $\mathbb{I}(s_i^j, s_i(t-1))$ is an Indicator function and equals 1 when $s_i^j = s_i(t-1)$. Otherwise it is zero.

Each attraction is applied an experience weight using:

$$N_i(t) = 1 + N_i(t-1)\phi(t-1), \text{ with } N_i(0) = 1$$

The change-detector function $\phi(t)$ weights the lagged attractions and characterizes a given player's perception about how quickly the learning environment is changing (Ho et al., 2007). It is defined as,

$$\phi_i(t) = 1 - \frac{1}{2}S_i(t);$$

$$S_i(t) = \sum_{k=1}^{m_{-i}} (h_i^k(t) - r_i^k(t))^2$$

where k is the strategy profile of other players. The cumulative history vector,$h_i^k(t)$, records the historical frequencies of the choices by other players (including the period t).

$$h_i^k(t) = \frac{\sum_{\tau=1}^{t-1} \mathbb{I}(k, s_{-i}(\tau))}{t}$$

where $s_{-i}(\tau)$ is the observed choice of strategies by others in round τ. The immediate history vector, $r_i^k(t)$, is a vector of 1's and 0's. Therefore in the round t, $r_i^k(t) = \mathbb{I}(k, s_{-i}(t)$. S(t) is the quadratic distance between the cumulative history vector $h_i^k(t)$ and the immediate history vector $r_i^k(t)$. It captures the degree of surprise due to change in the observed choices of the other players. It will always lie between 0 and 2 and so $\phi_i(t)$ lies in between 0 and 1.

The last component of the STEWA model, the attention function δ_i generates a weight for foregone payoffs and turns the attention to the strategies that would have led to higher payoffs. $\delta_i^q(t)$ is given as:

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

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\lambda A_i^j(t)}}{\sum_{l=1}^{m_i} e^{\lambda A_l^j(t)}}$$

where λ is a free parameter that determines the attraction sensitivity.

1.4.8 Impulse Matching Learning (IM)

Chmura et al. (2012) have introduced IM by combining the concepts of impulse-balance equilibrium (Selten & Chmura, 2008) and learning direction theory (Selten & Stoecker, 1986). In their study, using the data from Selten and Chmura (2008) on twelve 2×2

games that cover both constant sum and non-constant sum games, Chmura et al. (2012) found that IM explains choices very well compared to a number of other equilibrium and learning models including basic reinforcement learning , self tuning EWA, impulse-balance equilibrium, and action-sampling learning.

In IM, a player receives an impulse in a given round from a strategy that results in a lower payoff to a strategy that results in a higher payoff. In this manner, impulse matching learning encompasses the idea of regret-learning (Marchiori & Warglien, 2008). The impulse sums of strategies determine their propensities to be chosen by a player. In Chmura et al. (2012) impulses were calculated using transformed payoffs since original payoffs involved losses. In the context of public goods games, there are no losses so I compute impulses using original payoffs.

The description of the model is as below. In this learning model, propensities of strategies are described using impulse sums accumulated by strategies over time. Denote impulse sum of strategy j of individual i in round t as $R_i^j(t)$. The description of the model presented here differs slightly from the original version proposed in Chmura et al. (2012). In the original version, impulse sums are translated to choice probabilities via simple normalization, whereas, here choice probabilities are computed from impulse sums using a logit function as described below. The original version without any free parameters did very poorly and so I do not consider it here.

- **Initial Impulse Sums:** In the first round, impulse sums of all strategies are zero. Therefore, $R_i^j(1) = 0$. Thus in the first round, all choices have equal probabilities for being chosen.

- **Impulse Sum Updating:** $R_i^j(t)$ is the sum of all impulses from all other strategies of individual i towards s_i^j experienced up to period $t-1$.

  Impulse from strategy s_i^k towards strategy s_i^j in period $t-1$ for individual i is:

$$r_i^{jk}(t-1) = max[0, u_i(s_i^j, s_{-i}(t-1)) - u_i(s_i^k, s_{-i}(t-1))]$$

and impulse sum for strategy j in period t for individual i is:

$$R_i^j(t) = R_i^j(t-1) + \sum_{k=1}^{m_i} r_i^{jk}(t-1)$$

$u_i(s_i^j, s_{-i}(t-1))$ is the payoff obtained when the agent chooses strategy s_i^j given other have chosen strategies given by $s_i(t-1)$ in round $(t-1)$.

- **Stochastic Choice Rule:** Probability that agent i chooses strategy j in round t:

$$P_i^j(t) = \frac{e^{\lambda R_i^j(t)}}{\sum_{l=1}^{m_i} e^{\lambda R_l^j(t)}}$$

where λ is a free parameter that determines the sensitive to impulse sums.

1.5 Results

1.5.1 Descriptive Success of Learning Models

To begin with I compare descriptive success of learning models by reporting their aggregate fit computed using the pooled data from 18 datasets. Define N_d as the number of individuals and T_d as the number of rounds of the repeated public good game in dataset d. Then, the total number of individuals in the pooled data is $Ntotal = \sum_{d=1}^{18} N_d$. For a given learning model with parameters given by θ, let $P_i^j(t)$ is the predicted probability of individual i choosing a strategy (contribution level) s_i^j. Then, the log-likelihood function for the learning model can be written as:

$$LL(\theta) = \sum_{d=1}^{18} \sum_{i=1}^{N_d} \sum_{t=1}^{T_d} \sum_{j=1}^{m_i} ln\left(\mathbb{I}(s_i^j, s_i(t))P_i^j(t)\right)$$

Where m_i are the number of stage game strategies available to individual i and $\mathbb{I}(s_i^j, s_i(t))$ is the indicator function which is equal to 1 if $s_i^j = s_i(t)$ and 0 otherwise. $s_i(t)$ is the strategy chosen by individual i in round t.

I used Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm with numerical derivatives to maximize the log-likelihood function and estimate the parameters for each learning model. The fit of each learning model is described using four metrics: Log-Likelihood (LL), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and Pseudo-R^2. LL is the maximum log-likelihood achieved by a learning model. The last three metrics are defined as below for a learning model:

$$AIC = LL - k$$

$$BIC = LL - \frac{k}{2} ln(Ntotal)$$

$$Pseudo\text{-}R^2 = \frac{AIC - LLR}{LLR}$$

Where k is the number of parameters of the learning model, $Ntotal$ is the total number of observations, and LLR is the log-likelihood obtained by a random choice model. Since choices of an individual are not independent I consider the number of individuals as the effective size of the sample in the computation of BIC. Therefore, $Ntotal$ is equal to 1201 here. AIC and BIC penalize models for their complexity and therefore are superior metrics than log-likelihood for model comparison.

In evaluating models in terms of their descriptive and predictive success, I also consider two basic benchmarks. The first benchmark is derived from a random choice model. Denote the random choice model as RAND. Since the datasets differ in terms of the number of strategies(contribution levels) available to individuals, it should be noted that RAND model can only be derived at the dataset level. For example, in a dataset where the endowment is 20, e.g. as in DS9, RAND model predicts any given observed choice in that data set with

Table 1.2: Aggregate fit metrics and parameter estimates of learning models - Pooled data of 1201 individuals

Model	Parameter	Estimates	LL	AIC	BIC	Pseudo-R^2
REL	λ	2.2983	-35286	-35288	-35293	0.3080
	$N1$	1.0839				
EMP[†]			-36517			0.2838
EWA	λ	0.0316	-40712	-40717	-40730	0.2015
	δ	0.5880				
	ϕ	0.9688				
	κ	0.9322				
	η	0.8927				
NRL	λ	1.6616	-41086	-41088	-41093	0.1942
	w	0.0147				
RL	λ	1.4882	-42111	-42113	-42118	0.1741
	w	0.0075				
NFP	λ	2.0621	-49422	-49424	-49429	0.0307
	w	0.0149				
SFP	λ	0.4709	-49477	-49479	-49485	0.0297
	w	0.0307				
IM	λ	0.0006	-49475	-49476	-49479	0.0297
STEWA	λ	0.0133	-50235	-50236	-50238	0.0148
RAND[‡]			-50992			0.0000
Total number of individuals			1201			
Total number of decisions			17250			

EMP is the model based on empirical frequencies of strategies within each dataset. Note that the empirical frequencies are derived at the dataset level and not the pooled data level. This is because the number of strategies available to individuals are not identical across datasets due to varying levels of endowment.

RAND is the random choice model derived at the dataset level. This is because the number of strategies available to individuals are not identical across datasets due to varying levels of endowment.

a probability of $\frac{1}{21}$. Instead, if the endowment in a data set is 50 as in DS1, RAND model predicts any given observed choice in that data set with a probability of $\frac{1}{51}$. Obviously, the RAND benchmark is not a serious benchmark and I expect that every learning model should beat it. The second benchmark is derived from a model that is derived from the overall empirical frequencies of choices in each dataset. This model is referred to as EMP in the remainder of the paper. In the context of my analysis, EMP has 18 degrees of freedom since it is derived separately for each dataset. Thus, it is allowed to differ across datasets. A problem in deriving one model for the pooled data of 18 datasets is that, the number of strategies available to individuals varies across datasets. Without some kind of artificial aggregation of strategies (for example, one can put available strategies into three groups as low contribution, medium contribution, high contribution and then derive empirical frequencies of these constructed strategies across the pooled data), one cannot derive a single model of empirical frequencies by combining the data from all datasets. Since EMP is based on observed empirical frequencies of choices and is derived at the dataset level, it has a strong potential to fit the data very well. Thus, EMP is a very challenging benchmark for learning models to beat.If a learning theory successfully describes choice adjustments over time in a data set, it would obtain better fit than EMP since the latter does not depend upon temporal information.

Table 1.2 presents the parameter estimates that maximized the log-likelihood of the pooled data for each learning model. It also presents four aggregate fit metrics computed using the maximum likelihood parameter estimates. The fit achieved by the two benchmarks, RAND and EMP, is also reported for comparison. Using these results, the learning models can be ordered in terms of their performance as REL, EWA, NRL, RL, SFP, NSFP, IM, STEWA. This ordering holds according to any of the four fit metrics that are reported. REL outperforms all other models and the performance gap between REL and the second best performer EWA is quite large. This convinces at least from an aggregate descriptive fit point of view that REL explain the aggregate data convincingly better than all other models that are considered in this paper. As expected, all learning models are able to the perform better than the random choice benchmark. However, except REL all learning models fail to explain the data at least as well as EMP according to all four reported metrics. REL outperforms EMP meaning that REL is able to explain choice adjustment over time.

One of the clear patterns in Table 1.2 is that reinforcement learning models do relatively better in comparison to belief learning or regret learning models in explaining the data. Furthermore, models with adaptive attraction sensitivity, i.e. NRL and NFP, fit better compared to their counterparts, i. e. RL and SFP. Thus, adaptive attraction sensitivity that depends upon observed payoff variability seems to be important in explaining behavior in repeated public goods games. Despite having more degrees of freedom, EWA performs significantly worse than REL. IM and STEWA lie at the bottom of the table and perform marginally (but significantly according to likelihood ratio tests) better than a random choice model.

Another way to look into descriptive success of a learning model is by evaluating how well one set of parameters of the model that is estimated using the pooled data can explain observed choices in different datasets that correspond to different sets of game parameters. This is one way to evaluate the generalizability of a given learning model in terms of its descriptive success. A model that is generalizable should be able to explain behavior across games with varying game parameters (MPCR, endowment, number of rounds, and matching) without the need to adjust its learning parameters based on the parameters of the games. To evaluate models' descriptive generalizability, I compute LL, BIC, and Mean Quadratic Score of each learning model for each of the 18 datasets separately. In doing so, I use the parameter estimates of each learning model computed from the pooled data reported in Table 1.2.

The computation of LL and BIC was introduced earlier. Mean Quadratic Score (MQS) of a learning model is computed for a dataset using the quadratic scoring rule. The quadratic scoring rule was first introduced by Brier (1950), and was axiomatically characterized in Selten (1998). Let, for a subject i in round t of a repeated game, a given learning model predicts each strategy s_j^i with a probability $p_j^i(t)$. If the subject's observed strategy in round t is s_k^i, then the quadratic score of the learning model in round t is given as:

$$q_i(t) = 2p_k^i(t) - \sum_{j=1}^{m_i} (p_j^i(t))^2$$

Thus, in computing the quadratic score, the observed choice is interpreted as a degenerate probability distribution with observed strategy having a probability of 1 and all other strategies having a probability of 0. The quadratic score ranges between [-1, +1].[3] Higher scores indicate the relative success of a model in predicting observed choices.

If there are N_d individuals in dataset d and the public goods game is repeated for T_d rounds, then the MQS achieved by the learning model for the dataset d is:

$$MQS = \sum_{i=1}^{N_d} \left[\sum_{t=1}^{T_d} \frac{q_i(t)}{N_d \times T_d} \right]$$

Table 1.3 compares LL, BIC, and MQS achieved by each learning model at the dataset level. For each learning model, I compute the three fit metrics for each of the 18 datasets. For each pair of learning models, one sided p-value favoring the learning model in the row is reported using the 18 measures of each fit metric computed at the dataset level. The p-values are computed using the Exact Wilcoxon Signed-Rank Test for matched pairs. The results in Table 1.3 lead to similar conclusions as in Table 1.2. REL outperforms all other learning models in terms its descriptive success at the dataset level according to LL, BIC, and MQS metrics. This indicates that the parameter estimates of REL obtained using the estimation with pooled data are quite general that they explain behavior in different public goods games with parameters spanning a broad spectrum. What this indicates is that the parameters of REL need not to be significantly adjusted based on the parameters of the game to describe choice behavior in it. REL outperforms EMP in terms of MQS but performs

[3]If the observed choice is completely in line with prediction of the model then quadratic score achieves +1. This happens if the model puts a probability of one on the observed choice. The quadratic score would be minimum at -1 if the model puts all the probability on a strategy different from that of the observed strategy. All other probability distributions over strategies result in a score strictly between -1 and +1.

Table 1.3: Comparison of descriptive success of learning models. For each pair of learning models, one sided p-value favoring the learning model in the row is reported using the 18 measures of each fit metric computed at the dataset level. The p-values are computed using the Exact Wilcoxon Signed-Rank Test for matched pairs.

	EMP	NRL	EWA	RL	NFP	SFP	IM	STEWA	RAND
REL									
LL	0.3509	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
BIC	0.3830	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
MQS	0.0001	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
EMP									
LL		0.0333	0.0368	0.0104	0.0000	0.0000	0.0000	0.0000	0.0000
BIC		0.0333	0.0269	0.0091	0.0000	0.0000	0.0000	0.0000	0.0000
MQS		0.7387	0.8677	0.3198	0.0000	0.0000	0.0000	0.0000	0.0000
NRL									
LL			0.7659	0.0003	0.0003	0.0003	0.0004	0.0003	0.0006
BIC			0.6647	0.0003	0.0003	0.0003	0.0004	0.0003	0.0006
MQS			0.6006	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
EWA									
LL				0.0014	0.0004	0.0003	0.0004	0.0003	0.0006
BIC				0.0038	0.0004	0.0003	0.0004	0.0003	0.0006
MQS				0.0002	0.0000	0.0000	0.0000	0.0000	0.0000
RL									
LL					0.0005	0.0004	0.0004	0.0004	0.0006
BIC					0.0005	0.0004	0.0004	0.0004	0.0006
MQS					0.0000	0.0000	0.0000	0.0000	0.0000
NFP									
LL						0.4661	0.6006	0.3509	0.0216
BIC						0.4661	0.6491	0.3830	0.0368
MQS						0.2211	0.5339	0.3198	0.0038
SFP									
LL							0.7246	0.1733	0.0152
BIC							0.8036	0.2086	0.0241
MQS							0.5675	0.3047	0.0038
IM									
LL								0.3047	0.0069
BIC								0.3047	0.0091
MQS								0.4325	0.0017
STEWA									
LL									0.0024
BIC									0.0028
MQS									0.0000

Notes:

LL - Log-Likelihood

BIC - Bayesian Information Criterion

MQS - Mean Quadratic Score

EMP is the model based on empirical frequencies of strategies within each dataset. Note that the empirical frequencies are derived at the dataset level and not the pooled data level. This is because the number of strategies available to individuals are not identical across datasets due to varying levels of endowment.

RAND is the random choice model derived at the dataset level. This is because the number of strategies available to individuals are not identical across datasets due to varying levels of endowment.

equally as well as EMP in terms of LL and BIC. However, one need to keep in mind that EMP is a very tough benchmark to beat since the EMP model is derived separately for each data set helping it to fine tune its fit at the data set level. None of the other learning models are able to outperform EMP according to LL and BIC. NRL, EWA, and RL perform as well as EMP in terms of MQS. NRL and EWA jointly take the second place. They both perform equally well in explaining behavior across datasets with one set of parameters according to LL, BIC, and MQS metrics. RL comes third. As observed earlier, belief learning models do not perform well in explaining the data in comparison to the reinforcement models. NSFP, SFP, IM, and STEWA are statistically indistinguishable in terms of their performance. It is interesting to note that the top two models, REL and NRL, are similar in the sense they both use adapt attraction sensitivity based on the observed payoff variability. NRL performs significantly better than RL which is its identical counterpart without the dynamic attraction sensitivity adjustment confirming adaptive attraction sensitivity is an important construct in explaining behavior in repeated public goods games. In addition, REL scales the initial attraction sensitivity parameter using the expected payoff variability due to a random choice. What this scaling does is, it makes the learning model insensitive to payoff magnitudes. This appears to be instrumental in explaining learning behavior across different datasets involving a wide range possible payoffs. Thus, both of these features, the dynamic adjustment of attraction sensitivity based on observed payoff variability and the scaling of initial attraction sensitivity using a baseline payoff variability, appear to be important components of a learning model that does successfully explain behavior in repeated public goods games across a wide range of game parameters.

1.5.2 Predictive Accuracy of Learning Models

A more challenging test of model robustness and generalizability concerns with out-of-sample predictive accuracy of a given model. This line of analysis was pioneered in the context of learning in games in earlier studies by Erev and Roth (1998) and Ho et al. (2007). These studies have compared learning models by first estimating the parameters of the models using data on a set of games and then comparing how well the models predict choices in a new game with the estimated model parameters. My analysis here is different in that I study only one type of games here. But, the parameters of the game vary across datasets. To compare the out-of-sample predictive accuracy of learning models, I predict choices according to each learning model in each of the 18 datasets using the parameters estimated for the learning model using the remaining 17 datasets excluding the data set under consideration. These results test whether a model that is fitted using the data on public goods games with a set of given parameters (e.g. group size, marginal per capita return, number of rounds, matching protocol) can successfully predict choices in a public goods game with a set of game parameters that was not observed previously.

I characterize dataset level predictive accuracy of each learning model using five metrics: LL, AIC, BIC, Pseudo-R^2, and Mean Quadratic Score (MQS). Table 1.4 presents results about the predictive accuracy of models. For each dataset, I report LL, AIC, BIC, Pseudo-R^2, MQS achieved by each learning model with the parameters estimated from the remaining 17 datasets. The performance of RAND and EMP benchmarks is also reported for each dataset. Both of these benchmarks represent models derived at the data set level using the data within the dataset. In that sense, their scores do not reflect their predictive success but their descriptive success. A challenge in reporting predictive accuracy of these models is that data from different datasets cannot be pooled to derive empirical frequencies since the datasets vary in terms of the number of strategies (meaning they involve different levels of endowment) available to individuals. Since EMP reflects within data set empirical frequencies, it is impressive for a learning model if it can do as well as or better than EMP in prediction. In Table 1.4 , I marked the best score achieved for a given data set by any learning model with an asterisk. I marked EMP with a dagger in data sets where none of the learning models performed at least well as EMP in their predictions.

All the learning models other than REL, do not beat EMP in any of the data sets in terms of prediction. REL outperforms EMP in 9 datasets: DS4, DS9, DS11, DS12, DS14, DS15, DS16, DS17, DS18. In other datasets, it is outperformed by EMP. Among the learning models, REL is outperformed only in one instance by EWA in data set DS13 (the difference is very small). Interestingly, REL is the only model that outperforms random benchmark in all data sets. All the remaining learning models fail to perform at least as well as RAND in at least one of the datasets (for a given model, these are the instances where Pseudo-R^2 is negative).

The general picture that emerges here is similar to that of my findings in the descriptive fit results. Reinforcement learning models fare relatively well compared to that of belief learning models. The only regret-based learning model, IM, performs very poorly.

Table 1.4: Predictive accuracy of learning models: Out-of-sample prediction using out-dataset estimates. Fit metrics are computed for each learning model for each data set using the learning parameter estimates from 17 remaining data sets.

Learning Model	LL	AIC	BIC	Pseudo-R^2	MQS
DS1					
RL	-1293.08	-1295.08	-1296.76	0.1765	0.1176
NRL	-1294.75	-1296.75	-1298.44	0.1755	0.1222
REL	-1237.83*	-1239.83*	-1241.52*	0.2117*	0.1282*
SFP	-1541.99	-1543.99	-1545.68	0.0183	0.0243
NFP	-1536.83	-1538.83	-1540.51	0.0216	0.0248
EWA	-1284.80	-1289.80	-1294.02	0.1799	0.1181
STEWA	-1517.21	-1518.21	-1519.06	0.0347	0.0283
IM	-1539.17	-1540.17	-1541.02	0.0207	0.0268
EMP†	-986.64	-986.64	-986.64	0.3727	0.1097
RAND	-1572.73	-1572.73	-1572.73	0.0000	0.0196
DS2					
RL	-929.67	-931.67	-933.07	0.2101	0.1382
NRL	-915.07	-917.07	-918.47	0.2225	0.1553
REL	-872.29*	-874.29*	-875.69*	0.2588*	0.1648*
SFP	-1114.40	-1116.40	-1117.80	0.0535	0.0314
NFP	-1114.97	-1116.97	-1118.37	0.0530	0.0316
EWA	-878.14	-883.14	-886.64	0.2513	0.1649
STEWA	-1112.34	-1113.34	-1114.04	0.0561	0.0319

Continued on next page

Table 1.4 – *Continued from previous page*

Learning Model	LL	AIC	BIC	Pseudo-R^2	MQS
IM	-1092.50	-1093.50	-1094.20	0.0729	0.0394
EMP†	-680.57	-680.57	-680.57	0.4230	0.1557
RAND	-1179.55	-1179.55	-1179.55	0.0000	0.0196
DS3					
RL	-2356.95	-2358.95	-2361.33	0.2827	0.1956
NRL	-2345.58	-2347.58	-2349.96	0.2862	0.2076
REL	-2212.09*	-2214.09*	-2216.47*	0.3268*	0.2268*
SFP	-3184.93	-3186.93	-3189.31	0.0309	0.0250
NFP	-3181.67	-3183.67	-3186.05	0.0319	0.0252
EWA	-2311.96	-2316.96	-2322.92	0.2955	0.2169
STEWA	-3100.64	-3101.64	-3102.83	0.0569	0.0296
IM	-3154.38	-3155.38	-3156.57	0.0405	0.0344
EMP†	-1828.64	-1828.64	-1828.64	0.4440	0.1498
RAND	-3288.70	-3288.70	-3288.70	0.0000	0.0164
DS4					
RL	-578.07	-580.07	-581.47	0.1936	0.2450
NRL	-565.70	-567.70	-569.11	0.2108	0.2674
REL	-548.20*	-550.20*	-551.61*	0.2352*	0.2837*
SFP	-717.96	-719.96	-721.36	-0.0008	0.0919
NFP	-718.09	-720.09	-721.49	-0.0010	0.0919
EWA	-571.92	-576.92	-580.42	0.1980	0.2508
STEWA	-700.65	-701.65	-702.35	0.0246	0.1047
IM	-718.44	-719.44	-720.15	-0.0001	0.0915
EMP	-624.86	-624.86	-624.86	0.1314	0.1625
RAND	-719.37	-719.37	-719.37	0.0000	0.0909
DS5					
RL	-835.56	-837.56	-838.96	0.1431	0.1185
NRL	-820.82	-822.82	-824.22	0.1582	0.1322
REL	-757.89*	-759.89*	-761.29*	0.2226*	0.1736*
SFP	-912.86	-914.86	-916.26	0.0640	0.0587
NFP	-913.89	-915.89	-917.29	0.0630	0.0581
EWA	-778.23	-783.23	-786.74	0.1987	0.1622
STEWA	-929.67	-930.67	-931.37	0.0478	0.0530
IM	-925.95	-926.95	-927.65	0.0516	0.0544
EMP†	-679.64	-679.64	-679.64	0.3047	0.1804
RAND	-977.43	-977.43	-977.43	0.0000	0.0385
DS6					

Continued on next page

Table 1.4 – *Continued from previous page*

Learning Model	LL	AIC	BIC	Pseudo-R^2	MQS
RL	-12101.90	-12103.90	-12106.69	-1.5654	-0.0553
NRL	-8712.81	-8714.81	-8717.60	-0.8471	-0.0129
REL	-3421.43*	-3423.43*	-3426.22*	0.2744*	0.1820*
SFP	-6131.71	-6133.71	-6136.49	-0.3000	0.0024
NFP	-5633.40	-5635.40	-5638.18	-0.1944	0.0082
EWA	-12416.31	-12421.31	-12428.28	-1.6326	-0.0662
STEWA	-12281.91	-12282.91	-12284.30	-1.6033	-0.0032
IM	-5851.64	-5852.64	-5854.03	-0.2404	0.0018
EMP†	-3013.44	-3013.44	-3013.44	0.3613	0.1462
RAND	-4718.19	-4718.19	-4718.19	0.0000	0.0196
DS7					
RL	-992.97	-994.97	-996.66	0.3674	0.2684
NRL	-897.94	-899.94	-901.63	0.4278	0.3489
REL	-752.10*	-754.10*	-755.79*	0.5205*	0.4500*
SFP	-1249.33	-1251.33	-1253.02	0.2044	0.0711
NFP	-1233.11	-1235.11	-1236.80	0.2147	0.0758
EWA	-752.45	-757.45	-761.68	0.5184	0.4308
STEWA	-1416.00	-1417.00	-1417.85	0.0990	0.0397
IM	-1115.11	-1116.11	-1116.96	0.2903	0.1191
EMP†	-568.99	-568.99	-568.99	0.6382	0.4781
RAND	-1572.73	-1572.73	-1572.73	0.0000	0.0196
DS8					
RL	-2625.48	-2627.48	-2630.05	0.1010	0.0973
NRL	-2548.19	-2550.19	-2552.75	0.1275	0.1196
REL	-2261.99*	-2263.99*	-2266.55*	0.2254*	0.2114*
SFP	-2766.53	-2768.53	-2771.09	0.0528	0.0653
NFP	-2756.64	-2758.64	-2761.21	0.0561	0.0667
EWA	-2458.23	-2463.23	-2469.64	0.1572	0.1341
STEWA	-2833.88	-2834.88	-2836.16	0.0301	0.0573
IM	-2821.73	-2822.73	-2824.02	0.0342	0.0587
EMP†	-2202.31	-2202.31	-2202.31	0.2465	0.1864
RAND	-2922.74	-2922.74	-2922.74	0.0000	0.0476
DS9					
RL	-1912.45	-1914.45	-1916.73	0.1266	0.1192
NRL	-1842.03	-1844.03	-1846.31	0.1588	0.1515
REL	-1605.08*	-1607.08*	-1609.36*	0.2669*	0.2703*
SFP	-2078.85	-2080.85	-2083.12	0.0507	0.0643

Continued on next page

Table 1.4 – *Continued from previous page*

Learning Model	LL	AIC	BIC	Pseudo-R^2	MQS
NFP	-2072.32	-2074.32	-2076.60	0.0537	0.0655
EWA	-1797.04	-1802.04	-1807.73	0.1779	0.1575
STEWA	-2118.23	-2119.23	-2120.37	0.0332	0.0585
IM	-2121.45	-2122.45	-2123.58	0.0318	0.0579
EMP	-1638.84	-1638.84	-1638.84	0.2524	0.1866
RAND	-2192.06	-2192.06	-2192.06	0.0000	0.0476
DS10					
RL	-1680.61	-1682.61	-1684.71	0.0789	0.0842
NRL	-1646.93	-1648.93	-1651.03	0.0973	0.0988
REL	-1549.13*	-1551.13*	-1553.22*	0.1509*	0.1332*
SFP	-1787.85	-1789.85	-1791.94	0.0202	0.0550
NFP	-1790.20	-1792.20	-1794.30	0.0189	0.0547
EWA	-1635.25	-1640.25	-1645.49	0.1021	0.0981
STEWA	-1799.10	-1800.10	-1801.15	0.0146	0.0525
IM	-1798.50	-1799.50	-1800.55	0.0149	0.0526
EMP†	-1449.12	-1449.12	-1449.12	0.2067	0.1338
RAND	-1826.71	-1826.71	-1826.71	0.0000	0.0476
DS11					
RL	-5146.47	-5148.47	-5150.40	0.3368	0.3028
NRL	-5381.18	-5383.18	-5385.11	0.3066	0.2726
REL	-4723.23*	-4725.23*	-4727.16*	0.3914*	0.3223*
SFP	-7930.38	-7932.38	-7934.31	-0.0217	0.0985
NFP	-8171.66	-8173.66	-8175.60	-0.0528	0.0913
EWA	-6419.82	-6424.82	-6429.65	0.1724	0.2368
STEWA	-7586.61	-7587.61	-7588.58	0.0227	0.0550
IM	-7211.07	-7212.07	-7213.03	0.0710	0.0829
EMP	-5776.30	-5776.30	-5776.30	0.2560	0.1762
RAND	-7763.53	-7763.53	-7763.53	0.0000	0.0476
DS12					
RL	-2667.58	-2669.58	-2672.55	0.0473	0.1739
NRL	-2621.45	-2623.45	-2626.42	0.0638	0.1867
REL	-2496.56*	-2498.56*	-2501.53*	0.1083*	0.2375*
SFP	-2813.96	-2815.96	-2818.93	-0.0049	0.1407
NFP	-2816.64	-2818.64	-2821.61	-0.0059	0.1402
EWA	-2674.78	-2679.78	-2687.21	0.0437	0.1711
STEWA	-2801.52	-2802.52	-2804.01	-0.0001	0.1430
IM	-2802.45	-2803.45	-2804.94	-0.0005	0.1428

Continued on next page

Table 1.4 – *Continued from previous page*

Learning Model	LL	AIC	BIC	Pseudo-R^2	MQS
EMP	-2735.24	-2735.24	-2735.24	0.0239	0.1583
RAND	-2802.11	-2802.11	-2802.11	0.0000	0.1429
DS13					
RL	-2081.74	-2083.74	-2085.43	0.1445	0.1320
NRL	-2060.82	-2062.82	-2064.51	0.1531	0.1383
REL	-1892.62	-1894.62	-1896.31	0.2221	0.1818
SFP	-2177.31	-2179.31	-2180.99	0.1052	0.0849
NFP	-2182.20	-2184.20	-2185.89	0.1032	0.0845
EWA	-1858.98*	-1863.98*	-1868.20*	0.2347*	0.1897*
STEWA	-2351.47	-2352.47	-2353.31	0.0341	0.0585
IM	-2252.47	-2253.47	-2254.32	0.0748	0.0737
EMP†	-1853.36	-1853.36	-1853.36	0.2391	0.1541
RAND	-2435.62	-2435.62	-2435.62	0.0000	0.0476
DS14					
RL	-1626.38	-1628.38	-1629.96	0.2571	0.2374
NRL	-1584.17	-1586.17	-1587.75	0.2764	0.2587
REL	-1480.27*	-1482.27*	-1483.85*	0.3238*	0.2858*
SFP	-2274.53	-2276.53	-2278.11	-0.0385	0.0390
NFP	-2319.41	-2321.41	-2323.00	-0.0590	0.0348
EWA	-1612.05	-1617.05	-1621.01	0.2623	0.2449
STEWA	-2202.48	-2203.48	-2204.28	-0.0052	0.0468
IM	-2209.20	-2210.20	-2210.99	-0.0083	0.0457
EMP	-1842.77	-1842.77	-1842.77	0.1593	0.1256
RAND	-2192.06	-2192.06	-2192.06	0.0000	0.0476
DS15					
RL	-1408.14	-1410.14	-1412.17	0.1729	0.1564
NRL	-1373.61	-1375.61	-1377.63	0.1932	0.1730
REL	-1297.78*	-1299.78*	-1301.80*	0.2376*	0.2097*
SFP	-1724.40	-1726.40	-1728.43	-0.0126	0.0445
NFP	-1726.30	-1728.30	-1730.32	-0.0137	0.0442
EWA	-1412.01	-1417.01	-1422.07	0.1689	0.1563
STEWA	-1715.50	-1716.50	-1717.52	-0.0068	0.0465
IM	-1710.65	-1711.65	-1712.66	-0.0039	0.0467
EMP	-1401.02	-1401.02	-1401.02	0.1783	0.1256
RAND	-1704.93	-1704.93	-1704.93	0.0000	0.0476
DS16					
RL	-2790.79	-2792.79	-2795.54	0.2092	0.1938

Continued on next page

Table 1.4 – *Continued from previous page*

Learning Model	LL	AIC	BIC	Pseudo-R^2	MQS
NRL	-2685.64	-2687.64	-2690.39	0.2390	0.2254
REL	-2513.88*	-2515.88*	-2518.63*	0.2876*	0.2766*
SFP	-3558.55	-3560.55	-3563.30	-0.0082	0.0457
NFP	-3560.73	-3562.73	-3565.48	-0.0088	0.0456
EWA	-2799.87	-2804.87	-2811.75	0.2058	0.1919
STEWA	-3498.62	-3499.62	-3500.99	0.0091	0.0523
IM	-3534.05	-3535.05	-3536.43	-0.0010	0.0475
EMP	-2704.55	-2704.55	-2704.55	0.2342	0.1620
RAND	-3531.65	-3531.65	-3531.65	0.0000	0.0476
DS17					
RL	-1933.20	-1935.20	-1936.89	0.1930	0.2340
NRL	-1861.05	-1863.05	-1864.74	0.2230	0.2673
REL	-1644.38*	-1646.38*	-1648.07*	0.3134*	0.3372*
SFP	-2371.23	-2373.23	-2374.91	0.0103	0.0994
NFP	-2377.79	-2379.79	-2381.48	0.0076	0.1016
EWA	-1913.91	-1918.91	-1923.13	0.1998	0.2370
STEWA	-2370.28	-2371.28	-2372.12	0.0111	0.0963
IM	-2372.80	-2373.80	-2374.65	0.0100	0.0959
EMP	-1960.87	-1960.87	-1960.87	0.1823	0.1907
RAND	-2397.90	-2397.90	-2397.90	0.0000	0.0909
DS18					
RL	-5586.16	-5588.16	-5590.95	0.2232	0.2543
NRL	-5226.96	-5228.96	-5231.75	0.2731	0.3044
REL	-4093.37*	-4095.37*	-4098.16*	0.4307*	0.4393*
SFP	-6670.77	-6672.77	-6675.56	0.0724	0.1258
NFP	-6651.66	-6653.66	-6656.44	0.0751	0.1275
EWA	-5190.02	-5195.02	-5201.99	0.2778	0.2964
STEWA	-6992.92	-6993.92	-6995.31	0.0278	0.1037
IM	-6968.91	-6969.91	-6971.31	0.0311	0.1052
EMP	-4570.10	-4570.10	-4570.10	0.3647	0.3653
RAND	-7193.69	-7193.69	-7193.69	0.0000	0.0909

† EMP outperforms all learning models in this data set.

* Best among all other learning models for this data set. EMP is not considered in this comparison.

Table 1.5 compares LL, BIC, and MQS achieved by each learning model in predicting choices at the dataset level. For each pair of learning models, one sided p-value favoring the learning model in the row is reported using the 18 measures of each fit metric computed at the dataset level. The p-values are computed using the Exact Wilcoxon Signed-Rank Test for matched pairs. REL outperforms all other learning models in terms its predictive success at the dataset level according to LL, BIC, and MQS metrics. REL, NRL, RL, and EWA comfortably outperform the random choice benchmark. It is very impressive that REL outperforms EMP in its predictive success in terms of MQS given that EMP is derived within the sample. In terms of LL and BIC, REL performs as well as EMP. Again, one need to keep in mind that EMP is a very tough benchmark to beat since the EMP model is derived separately for each data set helping it to fine tune its fit at the dataset level. None of the other learning models are able to outperform EMP according to LL and BIC. NRL, EWA, and RL perform as well as EMP in terms of MQS. NRL and EWA jointly take the second place. RL comes third.

Belief learning models (SFP and NFP), regret learning model (IM), and the hybrid model of STEWA are indistinguishable from each other. They barely outperform the RAND benchmark. Thus, I reach a similar conclusion that learning in public goods games is more in line with reinforcement learning than belief learning or regret based learning.

1.6 Discussion

1.6.1 EWA Overfits and REL Generalizes Well

In this subsection I discuss the factors that contribute to REL's success in its descriptive success and predictive accuracy. First, I compare it side by side with EWA which does not do as as well. Among the learning models that are considered in this paper, EWA is the most flexible model since it has free parameters that help it to adjust the rate of discounting previous attractions using ϕ, type of learning (i.e. reinforcement or belief based

Table 1.5: Comparison of predictive accuracy of learning models. For each pair of learning models, one sided p-value favoring the learning model in the row is reported using the 18 measures of each fit metric computed at the dataset level. The p-values are computed using the Exact Wilcoxon Signed-Rank Test for matched pairs.

	EMP	NRL	EWA	RL	NFP	SFP	IM	STEWA	RAND
REL									
LL	0.3669	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
BIC	0.3994	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
MQS	0.0002	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
EMP									
LL		0.0269	0.0028	0.0091	0.0000	0.0000	0.0000	0.0000	0.0000
BIC		0.0269	0.0028	0.0091	0.0000	0.0000	0.0000	0.0000	0.0000
MQS		0.6170	0.6953	0.2899	0.0000	0.0000	0.0000	0.0000	0.0000
NRL									
LL			0.5169	0.0008	0.0010	0.0010	0.0010	0.0000	0.0010
BIC			0.3994	0.0008	0.0010	0.0010	0.0010	0.0000	0.0010
MQS			0.2341	0.0003	0.0000	0.0000	0.0000	0.0000	0.0000
EWA									
LL				0.0333	0.0010	0.0010	0.0010	0.0000	0.0010
BIC				0.0708	0.0010	0.0010	0.0010	0.0000	0.0010
MQS				0.0104	0.0000	0.0000	0.0000	0.0000	0.0000
RL									
LL					0.0010	0.0010	0.0010	0.0000	0.0010
BIC					0.0010	0.0010	0.0010	0.0000	0.0010
MQS					0.0001	0.0001	0.0000	0.0000	0.0000
NFP									
LL						0.4493	0.5841	0.3994	0.1144
BIC						0.4493	0.6331	0.4325	0.1519
MQS						0.2899	0.3353	0.0982	0.0060
SFP									
LL							0.7914	0.3830	0.0770
BIC							0.8376	0.3994	0.0982
MQS							0.3509	0.1061	0.0069
IM									
LL								0.1733	0.0080
BIC								0.1733	0.0104
MQS								0.0837	0.0033
STEWA									
LL									0.0024
BIC									0.0028
MQS									0.0024

Notes:

LL - Log-Likelihood

BIC - Bayesian Information Criterion

MQS - Mean Quadratic Score

EMP is the model based on empirical frequencies of strategies within each dataset. Note that the empirical frequencies are derived at the dataset level and not the pooled data level. This is because the number of strategies available to individuals are not identical across datasets due to varying levels of endowment.

RAND is the random choice model derived at the dataset level. This is because the number of strategies available to individuals are not identical across datasets due to varying levels of endowment.

or a mix of the two) using δ, the type of attraction updating (i.e. cumulative or averaging) using the parameter κ. Despite all the flexibility it has, it does not fare as well as REL. To gain insight into what is behind the superior performance of REL compared to EWA in its generalizability, I fit these two models at the data set level. Under Calibration in Table 1.6, the parameter values that maximized log-likelihood of a given dataset are reported for the REL Model. Under Prediction in Table 1.6, the parameter values that maximized log-likelihood across the 17 datasets other than a given dataset are reported. These are the parameters that were used to predict the choices in the dataset in Section 1.5.2. MQS values achieved in calibration and prediction are also reported. Table 1.7 presents parameter estimates and MQS achieved for both calibration and prediction for EWA model. The parameters in the calibration of EWA help it to fit the choices in each dataset separately giving it a much better chance in fitting the data due to its high degrees of freedom. Indeed, when we compare the MQSs achieved by EWA and REL at the dataset level calibration, EWA outperforms REL (Exact Wilcoxon Signed-Rank Test: $p < 0.001$). However, as we noted in the Results section 1.5.2, in prediction REL outperforms EWA (Exact Wilcoxon Sign-Rank Test: $p < 0.001$). This highlights the problem with EWA: it overfits the data.

The parameter values for REL are very close in both calibration and prediction for all datasets helping it to achieve very close performance in both calibration and prediction. Note that one cannot achieve a better fit than the fit achieved in calibration. However, in the context of EWA, there are significant differences in the estimated parameter values in calibration and prediction. Since the parameters used for prediction are way far away from the calibration parameters that lead to the maximum possible likelihood using the learning model (the log-likelihood in calibration), the achieved prediction MQS values are quite low.

The most important parameters in EWA learning model are the rate of forgetting parameter ϕ, the aspiration level parameter δ that determines the type of learning, and κ that determines whether attractions are updated cumulatively or by averaging. Table 1.7 shows that large discrepancies are occurring in terms of κ and δ between the calibrated parameters and prediction parameters. κ's used for prediction are closer to 1 meaning attraction updating is cumulative, but the calibrated parameter values are significantly smaller indicating that learning that achieves the best fit is actually closer to an averaging attraction updating model. Similarly, the prediction δ values lean towards more of belief learning ($\delta = 1$ indicates pure belief learning and $\delta = 0$ indicates pure reinforcement learning). However, for many data sets (for example DS6, DS15, DS16), the calibrated δ is very close to 0 indicating the learning that explains the data in those datasets is much closer to that of reinforcement learning. The parameters that are used for prediction are very close the parameters that maximized log-likelihood for the pooled data reported in Table 1.2, which makes sense because these parameters were estimated with 17 of the 18 datasets. The take away from the comparison of calibration and prediction parameters of EWA is that the model parameters need to be adjusted significantly based on the game parameters to explain behavior and it has a hard time finding one set of parameters that can describe well the behavior across public goods games with a wide range of parameters. Such discrepancies between the parameters estimated using pooled data and the parameters estimated at the individual dataset level are also reported in Erev and Haruvy (2005). These results demonstrate how easy it is to run into overfitting problems when studying learning in games. For example, if we take any of the 18 datasets in isolation, EWA achieves significantly better performance than REL in fitting the data, therefore one may have to conclude learning is more in line with that of EWA. We are only able to identify the overfitting problem of EWA by studying its performance in describing and predicting choices in multiple datasets that have data on public goods games with a wide ranging parameters.

There are two components in REL that set it apart from other models in the reinforcement family: adaptive attraction sensitivity that decreases with increasing payoff variability and its insensitivity to payoff magnitude. By comparing RL and NRL, one can conclude that adaptive attraction sensitivity that decreases with increasing payoff variability is important in explaining choice behavior in public goods environment. NRL is identical to RL but has an adaptive attraction sensitivity and achieves a significantly better performance in both descriptive fit and predictive accuracy. Apart from very minor specification details, both REL and NRL are identical in that they both use average choice reinforcement and both employ an adaptive attraction sensitivity that decreases with increasing payoff variability. Here, I assess if REL's insensitivity to payoff magnitude puts it ahead of NRL. REL achieves insensitivity to payoff magnitude by scaling the attraction sensitivity parameter (λ) by baseline payoff variability that is computed based on a random choice in the first round. Therefore, first round's effective attraction sensitivity is $\lambda/PV1$. For example, in DS1 where E = 50 and M = 0.5, PV1 is 26.07. Whereas in DS4 where E = 10 and M = 0.75, PV1 is 18.47. First round effective attraction sensitivity of the REL learning model will be different for individuals in these two data sets. In contrast, NRL's first round payoff variability is initialized as $PV1 = \lambda$. This makes the effective attraction sensitivity, $\lambda/PV1$, equal to 1 in any dataset irrespective of E or M. Thus, in NRL learning will be sensitive to payoff magnitudes in the game which depend upon the game's E and M. To evaluate how much REL's payoff magnitude insensitivity contribute to its success , I estimate a variation of REL model where I initialize $PV1 = \lambda$ making the first round effective attraction sensitivity equal to 1 as in NRL. The resulting model indeed performs significantly worse than the original REL (Likelihood Ratio Test, $\chi^2 = 842, p < 0.0001$). I conclude that both adaptive attraction sensitivity that scales with payoff variability and leaning that is insensitive to payoff magnitudes are important factors that underlie the success of REL in explaining choice behavior in public goods games.

Table 1.6: Fit and parameter estimates for calibration and validation of REL at the dataset level

	Characteristics				Calibration			Prediction		
Data Set	T	E	M	MT	MQS	λ	$N1$	MQS	λ	$N1$
DS1	10	50	0.5	S	**0.127**	2.275	0.762	**0.128**	2.299	1.100
DS2	10	50	0.5	P	**0.182**	2.795	0.063	**0.165**	2.295	1.116
DS3	10	60	0.5	S	**0.248**	3.008	0.082	**0.227**	2.292	1.250
DS4	10	10	0.75	P	**0.288**	2.759	1.861	**0.284**	2.296	1.082
DS5	10	25	0.3	P	**0.179**	2.503	0.205	**0.174**	2.297	1.113
DS6	10	50	0.3	P	**0.220**	2.936	0.496	**0.182**	2.279	1.344
DS7	10	50	0.03	P	**0.565**	3.788	1.029	**0.450**	2.271	1.147
DS8	10	20	0.4	P	**0.214**	2.235	2.125	**0.211**	2.311	1.058
DS9	10	20	0.4	S	**0.272**	2.408	2.444	**0.270**	2.296	1.027
DS10	10	20	0.5	P	**0.160**	1.909	2.555	**0.133**	2.326	1.090
DS11	50	20	0.5	P	**0.336**	2.334	4.363	**0.322**	2.293	0.984
DS12	10	6	0.5	P	**0.257**	1.920	3.728	**0.238**	2.323	0.975
DS13	20	20	0.4	P	**0.190**	2.111	2.475	**0.182**	2.319	1.102
DS14	20	20	0.65	P	**0.289**	2.146	1.642	**0.286**	2.305	1.071
DS15	10	20	0.8	S	**0.217**	2.058	0.549	**0.210**	2.316	1.130
DS16	10	20	0.8	PS	**0.276**	2.331	1.975	**0.277**	2.299	1.035
DS17	25	10	0.5	P	**0.340**	2.117	1.700	**0.337**	2.303	1.022
DS18	25	10	0.5	S	**0.441**	2.674	2.549	**0.439**	2.271	1.039

T - Number of Rounds
E - Endowment
M - MPCR
MT - Matching: P-Partners, S - Strangers, PS-Perfect Strangers
MQS - Mean Quadratic Score of REL at the data set level

Table 1.7: Fit and parameter estimates for calibration and validation of EWA at the dataset level

	Characteristics				Calibration						Prediction					
Data Set	T	E	M	MT	MQS	η	ϕ	κ	δ	λ	MQS	η	ϕ	κ	δ	λ
DS1	10	50	0.5	S	**0.134**	0.353	0.901	0.298	0.327	0.070	**0.118**	0.920	0.969	0.958	0.596	0.031
DS2	10	50	0.5	P	**0.187**	0.487	0.898	0.114	0.493	0.168	**0.165**	0.871	0.970	0.913	0.589	0.032
DS3	10	60	0.5	S	**0.258**	0.267	0.828	0.305	0.368	0.071	**0.217**	0.992	0.970	0.999	0.607	0.030
DS4	10	10	0.75	P	**0.290**	0.136	0.933	0.000	0.732	0.461	**0.251**	0.884	0.970	0.926	0.588	0.032
DS5	10	25	0.3	P	**0.208**	0.000	0.825	0.461	0.513	0.109	**0.162**	0.885	0.970	0.925	0.586	0.031
DS6	10	50	0.3	P	**0.228**	0.180	0.781	0.244	0.000	0.011	**-0.066**	0.208	0.894	0.520	0.461	0.099
DS7	10	50	0.03	P	**0.589**	0.756	0.662	0.000	0.461	0.244	**0.431**	0.002	0.975	0.973	0.576	0.028
DS8	10	20	0.4	P	**0.252**	0.234	0.827	0.371	0.506	0.203	**0.134**	0.001	0.974	0.953	0.576	0.028
DS9	10	20	0.4	S	**0.289**	0.227	0.890	0.095	0.443	0.313	**0.157**	0.000	0.973	0.962	0.581	0.029
DS10	10	20	0.5	P	**0.176**	0.769	0.815	0.000	0.386	0.447	**0.098**	0.001	0.971	0.968	0.587	0.030
DS11	50	20	0.5	P	**0.358**	0.417	0.862	0.018	0.250	0.337	**0.237**	0.271	1.000	0.754	0.699	0.043
DS12	10	6	0.5	P	**0.261**	0.359	0.845	0.000	0.118	0.442	**0.171**	0.005	0.970	0.974	0.592	0.030
DS13	20	20	0.4	P	**0.237**	0.188	0.846	0.256	0.542	0.223	**0.190**	0.000	0.971	0.961	0.577	0.029
DS14	20	20	0.65	P	**0.302**	0.385	0.804	0.000	0.300	0.212	**0.245**	0.001	0.971	0.964	0.595	0.030
DS15	10	20	0.8	S	**0.229**	0.045	0.775	0.255	0.000	0.065	**0.156**	0.001	0.971	0.981	0.592	0.029
DS16	10	20	0.8	PS	**0.287**	0.104	0.908	0.139	0.000	0.086	**0.192**	0.000	0.973	0.968	0.598	0.029
DS17	20	10	0.5	P	**0.358**	0.346	0.770	0.038	0.374	0.433	**0.237**	0.002	0.968	0.964	0.598	0.030
DS18	20	10	0.5	S	**0.498**	0.000	0.825	0.214	0.502	0.404	**0.296**	0.504	0.970	0.741	0.570	0.033

T - Number of Rounds
E - Endowment
M - MPCR
MT - Matching: P-Partners, S - Strangers, PS-Perfect Strangers
MQS - Mean Quadratic Score of EWA at the data set level

1.6.2 Comparison of REL with Individual Evolutionary Learning

There is one thing that is common among the learning models that are considered so far in this chapter. All of them are based on interpretable cognitive or psychological mechanisms that were put forward by decades of experimental research on humans and animals. For example, the idea of reinforcing actions that led to higher payoffs in the previous rounds of play with higher probabilities is one of the simplest and oldest of learning models and its origins can be traced to the psychological learning theories of Luce (2005), Thurstone (1930) and Bush and Mosteller (1955). Thus, learning models considered until now rank high in terms of cognitive plausibility. On the other end of the spectrum are learning models based on the template of evolution. These models are often hard to interpret in terms of underlying cognitive or neural mechanism substrates. A number of papers have considered if evolutionary learning models can provide a successful account of the outcomes in linear public goods games. For example, see Clemens and Riechmann (2006); Miller and Andreoni (1991); Arifovic and Ledyard (2012).

In this subsection, I compare the Individual Evolutionary Learning (IEL) model put forward by Arifovic and Ledyard (2012) with REL model in its ability to explain contribution patterns observed in the data considered in this Chapter Arifovic and Ledyard (2012) have shown that IEL in conjunction with their social preference model can explain many observed regularities in repeated public goods games. To put IEL and REL on the same footing, in this paper I do not consider social preferences in conjunction with IEL. A brief description of IEL learning model is provided in Appendix 1.8. There are two free parameters in IEL model: memory size (J), experimentation probability (ρ). J denotes the size of the set of strategies an individual remembers and ρ determines the probability with which each strategy in the memory of an individual gets replaced by a random strategy that is picked from all available strategies.

Standard econometric methods like maximum likelihood estimation are difficult to use with IEL since it induces a Markov process on the sets of actions rather than on individual actions. Therefore, I use simulation methods to compare IEL with REL in terms of how closely they fit the distribution of choices in the data.

Since the decision making process in the learning models is stochastic, I conduct 100 simulations per dataset and derive a measure of how close the simulated contributions are to the observed contributions. One way to evaluate the closeness of observed contributions and simulated contributions is to compare the average contributions in simulations with the average contributions observed in each dataset over the rounds. First average contribution in each round is computed in each simulation run to derive the path of average contributions. These simulated paths of average contributions over 100 simulation are averaged to obtain the representative path of average contributions. One can then compute the distance between the computed path of average contributions with the actual path of average contributions in the dataset. This distance measure average over 18 data sets would give a measure of how successful a learning model is in generating the average contributions that are closer to the observed average contributions. This is the approach taken by a number of previous papers in evaluating learning models (Deadman, 1999; Janssen & Ahn, 2006; Castillo & Saysel, 2005). However, contributions in public good experiments are consistently observed to be multi modal with generally three peaks at zero, half endowment, and full endowment(Ledyard, 1995). While contributions decrease over time, multi modal nature of them persists. Thus a model that produces a very close pattern of average contributions over time might not necessarily capture the distributional aspect of the contributions well. In light of this, a good model should be the one that generates a distribution of contributions in each round that is close to the distribution of observed contributions in that round. I characterize the similarity between the distribution of choices simulated by a learning model in a round with the distribution of the observed choices using two distributional distance measures. I use two measures, since it was observed in previous literature that different choices of distance measures can lead to different conclusions in model selection (Feltovich, 2000). To guard the analysis from such an occurrence here, I use two distance measures to compare REL and IEL.

To compare how well IEL and REL generate contributions that are closer to the observed contributions, I use two distributional distance measures. A number of distance metrics are proposed to measure the distance between two probability distributions. The Kullback-Leibler (KL) divergence and Pearson's-χ^2 divergence are two of the well-known such metrics. A particular concern in using these widely used metrics is that they are asymmetric. To circumvent this problem, I use symmetric versions of these divergence measures: Jensen-Shannon divergence (Lin, 1991) is a symmetric extension of Kullback-Leibler (KL) divergence and Probabilistic Symmetric$-\chi^2$ (Deza & Deza, 2006) is a symmetric variation of the Pearson's-χ^2. Say the empirical discrete distribution of the observed contributions in a given round is P and the empirical discrete distribution of the simulated choices in the round is Q. If E is the endowment, there are $E+1$ strategies (discrete contribution levels from 0 to E.). The two distance measures can then be defined as:

1. Probabilistic Symmetric$-\chi^2$ distance:

$$d_{\chi^2} = 2\sum_{j=0}^{E} \frac{(P_j - Q_j)^2}{P_j + Q_j}$$

2. Jensen-Shannon divergence:

$$d_{JS} = \frac{1}{2}\left[\sum_{j=0}^{E} P_j ln\left(\frac{2P_j}{P_j + Q_j}\right) + \sum_{j=0}^{E} Q_j ln\left(\frac{2Q_j}{P_j + Q_j}\right)\right]$$

For REL and IEL, I compute the distance measures for each of the 18 datasets using simulation. I take an average of the dataset level distance measures as the overall measure of distance between the simulated and observed contributions. I conduct a grid search to find the values of learning parameters that minimize the distances for each model. For REL, I varied λ from 0 to 4 in increments of 0.1 and varied $N1$ from 0 to 10 in increments of 1. For IEL, I varied J from 50 to 200 in increments of 10 and varied ρ from 0 to 1 in

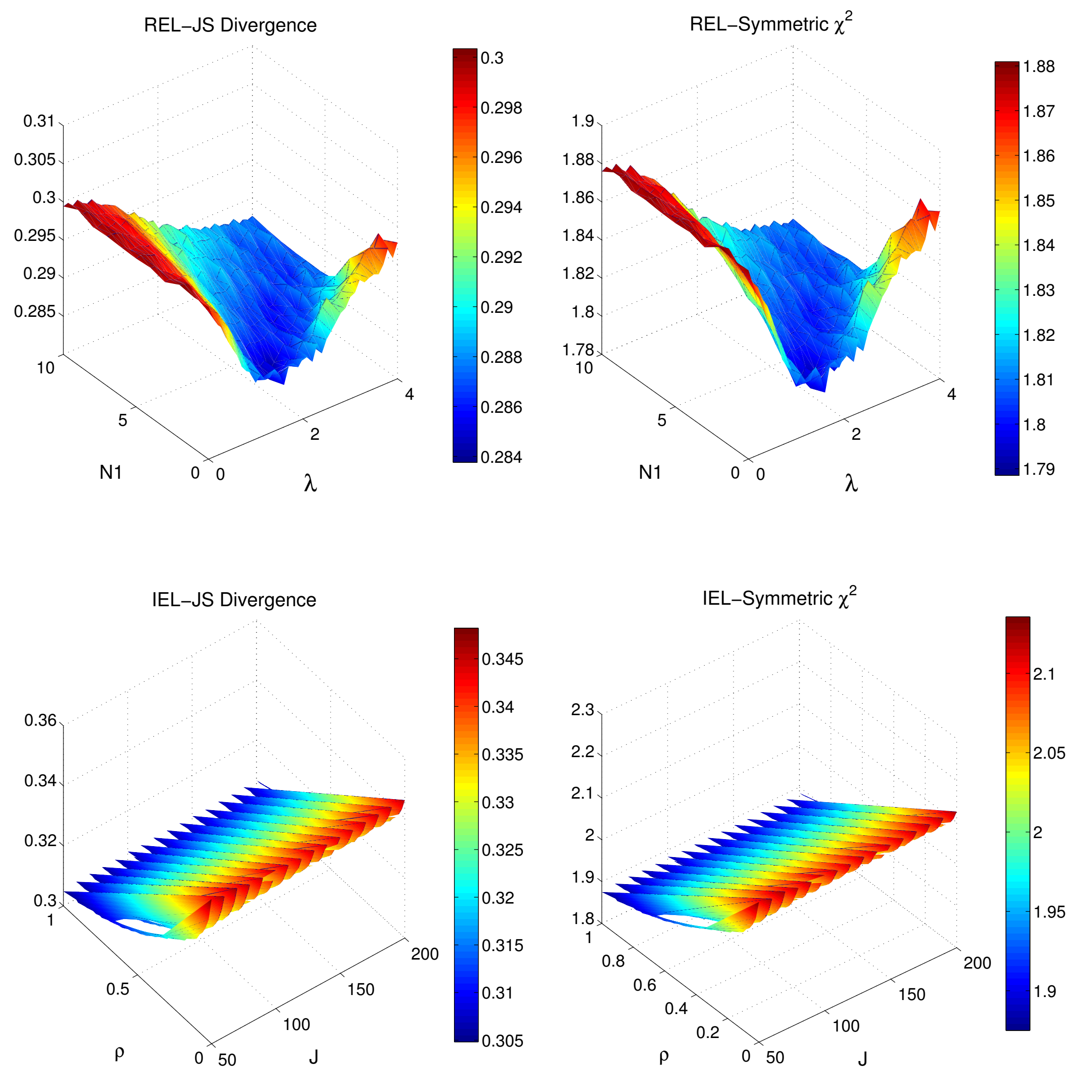

Figure 1.1: Evaluation of the distribution of the population's contributions for REL and IEL models. JS divergence and Symmetric-χ^2 distances measure (symmetrically) the difference between the true data and the model output. Lower JS divergence and Symmetric-χ^2 distances represent higher accuracy. The plots show computed distance measures at each pair of parameter values in the grid search.

Table 1.8: Evaluation of the distribution of the population's contributions for REL and IEL models. JS divergence and Symmetric-χ^2 measure (symmetrically) the distance between the true data and the model output. Lower JS divergence and Symmetric-χ^2 distances represent higher accuracy. Reported are the minimum distances achieved by each model.

	JS-Divergence		Symmetric-χ^2	
Model	Parameters	d_{JS}	Parameters	d_{χ^2}
REL	$\lambda = 1.9, N1 = 2$	0.2838	$\lambda = 1.9, N1 = 2$	1.7885
IEL	$J = 50, \rho = 1$	0.3048	$J = 50, \rho = 1$	1.8750

increments of 0.02. Achieved distance measures over the corresponding grids are shown in Figure 1.1 for both models. The smallest distances achieved for two models are presented in Table 1.8. REL outperforms IEL in predicting the full distribution of contributions over rounds. From Figure 1.1 it can be noted that IEL performs better for high values of ρ. The best score achieved by IEL is for $\rho = 1$ which means that every strategy in the memory is replaced by a random strategy with a probability of 1 in each round. REL does better than a random choice model since the best score REL achieves is better than what it achieves when $\lambda = 0$ (when $\lambda = 0$, REL is equivalent to the random choice model). I conclude that REL outperforms IEL in predicting full distribution of contributions.

1.7 Conclusions

In this Chapter, the models of reinforcement learning (RL), normalized reinforcement learning (NRL), reinforcement average model with loss aversion strategy (REL), stochastic fictitious play (SFP), normalized stochastic fictitious play (NFP), experience weighted attraction learning (EWA), self-tuning EWA (STEWA), and impulse matching learning (IM) were applied and tested in the environment of public goods games with a wide ranging game parameters. Experimental data from 12 published studies was used to construct 18 datasets each of which is different from the remaining in at least one the public goods game parameters: marginal per capita return, group size, number of rounds, endowment, and matching protocol. Learning models were compared on the basis of their descriptive fit and prediction accuracy.

The main finding in this study is that REL outperforms all other learning models in terms of both descriptive fit and prediction accuracy. It is the only model to outperform dataset level overall empirical frequencies in both descriptive fit and predictive accuracy when model performance is compared using Mean Quadratic Score (MQS). In terms of log-likelihood (LL) and BIC, REL performs as well as the dataset level overall empirical frequencies in both comparisons. None of the remaining models out-perform the dataset level overall empirical frequencies in both comparisons according to any of the metrics. EWA and NRL jointly take the second place. They are outperformed by the dataset level overall empirical frequencies in terms of LL and BIC. But, they do as well as the dataset level overall empirical frequencies in terms of MQS. The ordering of learning models in both comparisons is identical: REL, NRL, EWA, RL, NSFP, SFP, IM, STEWA (the ordering is not strict in all cases. REL performs strictly better than all other models). REL also outperforms individual evolutionary learning (IEL) in predicting full distribution of contributions.

The results in this chapter suggest that belief learning models and regret-based learning models perform poorly in comparison to reinforcement models in public goods environments. Learning in public goods games appears to be more in line with reinforcement learning. Average reinforcement learning that is adaptive to observed payoff variability and insensitive to the payoff magnitude underlies the success of REL. These results stand in stark contrast to the results from studies with 2×2 games. For example, Chmura et al. (2012) finds that regret-based models and belief learning models out-perform reinforcement learning models in explaining data from twelve 2×2 games that cover both constant sum and non-constant sum games. One potential reason for the success of reinforcement learning models in public goods games could be that individuals perceive public goods environments being too complex and thus resort to simple experience based reinforcement learning model rather than much more complex belief learning or regret-based learning.

There could be two additional potential reasons behind the failure of belief based and regret-based learning models in linear public goods environments. First, in the basic public goods game setting, these models imply a faster learning of the Nash equilibrium contribution which is zero. Thus contributions drop very fast to near zero. This also happens with IEL model without social preferences. But in the data, the movement towards zero is slow and contributions in the later rounds of a repeated game are multi-modal with a non-trivial probability mass at non-zero contribution levels. Second, it is possible that individuals could be forming beliefs in a way very different from that of what is captured in the belief learning models like SFP and NFP. In Chapter 2, I evaluate if belief learning models can perform better at explaining public goods game data when they are combined with a model of social preferences. Social preferences help the learning models not to drop to zero contributions since multiple Nash equilibria are possible and not all of them involve zero contributions. I also study if elicited beliefs capture the learning much more accurately than the beliefs modeled by prominent belief learning models. To do so, I conduct experiments where I elicit beliefs of individuals in each round of a repeated public goods game in an incentive compatible manner.

1.8 Appendix

1.8.1 Individual Evolutionary Learning (IEL)

Denote a given repeated public goods game as (G, T). Where G is the stage game and T is the number of rounds. The stage game $G = (N, X, u, r)$ consists of the number of players in the group N, the strategy space $X^i = [0, E]$ for $i \in N$, the payoff function $u^i = u^i(x^1, x^2, ..., x^N)$, and the information reported to player i at the end of a round, $r^i(x)$. IEL posits that individual i remembers a finite set of strategies A^i_t in round t and $A^i_t \subset X^i$. The size of the set A^i_t, which can be interpreted as memory size, is bounded and is a parameter of the learning model. It is assumed that the size $|A^i_j| = J$ is identical for all rounds $t \in T$ and across all agents. The probability that a strategy $a^i_{j,t} \in A^i_t$ is chosen by i is computed from a probability measure on A^i_t.

IEL has four components as described below:

- **Initilization:**

  Following, Arifovic and Ledyard (2012) I assume things begin randomly. This also puts things on the same footing compared to the other learning models I considered earlier. A^i_1 is population randomly with J random draws from X^i. This is a reasonable assumption from an empirical standpoint as it would explain why contributions start around average endowment in the first round.

- **Experimentation:**

  For each $j = 1, ..., J$, with experimentation probability ρ, a random strategy a^i_k is selected from X^i and replaces a^i_j. a^i_k is a random draw from the normal distribution whose mean is a^i_j and standard deviation is σ which is a parameter of the learning model. Arifovic and Ledyard (2012) set $\sigma = \frac{E}{10}$, where E is the endowment, in their paper and I follow the same.

- **Replication:**

  The idea of replication is to reinforce strategies that could lead to higher payoffs in the previous rounds. Thus replication of strategies is done using foregone payoffs. Say, the foregone utility of strategy j is denoted by $v^i(a^i_{jt}|r_i(x_t))$, where $r_i(x_t)$ is the information given to the individual i after round i. The way replication works is described as follows. For $j = 1, ...J$, pick two members from A^i_t randomly (giving equal probability to each member) with replacement. Let these two members are $a^i_{m,t}$ and $a^i_{n,t}$. Then the j^{th} member of A^i_{t+1} is chosen as:

$$a^i_{j,t+1} = \begin{cases} a^i_{m,t} & \text{if } v^i(a^i_{m,t}|r_i(x_t)) > v^i(a^i_{n,t}|r_i(x_t)) \\ a^i_{n,t} & \text{Otherwise} \end{cases}$$

- **Selection:**

  Each strategy $a^i_{k,t+1} \in A^i_{t+1}$ is selected with probability

$$P^i_{k,t+1} = \frac{u^i(a_{k,t+1}|r^i(x_t)) - \varepsilon^i_{t+1}}{\sum_{j=1}^{J}\left[u^i(a_{j,t+1}|r^i(x_t)) - \varepsilon^i_{t+1}\right]}$$

  for al $i \in 1, ..., N$ and all $k \in 1, ..., J$, where

$$\varepsilon^i_{t+1} = \min_{a \in A^i_{t+1}} \{0, u^i(a|r^i(x_t))\}$$

Chapter 2: Social Preferences and Learning in Public Goods Games: An Experimental Investigation

Abstract

I examine social preferences and learning in linear public goods games. I conduct economic experiments in which I collect decisions on conditional contributions in one-shot games and contributions in finitely (and indefinitely) repeated games at different relative costs of cooperation (i.e. the marginal per capita return (MPCR) to the public good) for each participant. Data from conditional contribution choices is used to estimate a structural model of social preferences. The estimated distribution of social preference parameters is reasonably successful in explaining the aggregate changes in the distribution of cooperative types (free riders, conditional cooperators, and full cooperators) when prices change. Choices in the repeated games are best described by social preferences of participants affecting the choice of first round contributions and then subsequent contributions based on payoff-based reinforcement learning. I find that learning is identical across finitely and indefinitely repeated public goods games. I deploy the behavioral specification thus obtained in an agent-based model. Simulations using the agent-based model demonstrate a novel result that reducing the price of cooperation in the first round alone is sufficient to sustain significantly higher contributions over the later rounds of a repeated game. The quantitative predictions of the empirical agent-based model are successfully reproduced using follow-up behavioral experiments substantiating the predictive value of agents behavioral specification.

2.1 Introduction

In this chapter, I study the role of social preferences and learning in repeated (linear) public goods games. The econometric investigations in Chapter 1 have compared a broad set of models from the families of reinforcement learning, belief learning, regret learning, and evolutionary learning and concluded that learning in public goods games is more in line with an average reinforcement learning model. Several research papers have argued that individuals' predisposition towards cooperation is also important alongside with learning in explaining contributions in repeated public goods games (Janssen & Ahn, 2006; Wendel & Oppenheimer, 2010; Arifovic & Ledyard, 2012). These arguments were put forward based on the experimental evidence that individuals vary in their predisposition towards cooperation and can be categorized into cooperative types (Fischbacher, Gächter, & Fehr, 2001; Fischbacher & Gächter, 2010; Kurzban & Houser, 2001, 2005; Burlando & Guala, 2005; Duffy & Ochs, 2009; Kocher, Cherry, Kroll, Netzer, & Sutter, 2008; Herrmann & Thöni, 2009; Muller, Sefton, Steinberg, & Vesterlund, 2008; Bardsley & Moffatt, 2007). Some individuals give everything or act as full cooperators (FC), some individuals give nothing or act as free riders (FR), and some individuals act as conditional cooperators (CC) meaning they match the average contribution of others in public goods games. In their seminal study, Fischbacher and Gächter (2010) devised an experimental design where they elicited the cooperative type information of individuals in a one-shot game and then observed the same individuals in a repeated game. In the repeated game, both contribution decisions and beliefs about the average contribution of others were elicited. Their empirical analysis showed that belief formation is identical across cooperative types. However, an individual's contribution choices were significantly informed by his cooperative type and average contributions in the repeated game were significantly different across types. While my investigation of learning models in Chapter 1 is extensive in that it considers a broad range of learning models and data from a number of experimental studies, it is necessarily incomplete. All the learning models I have considered in Chapter 1 do not incorporate elements that can explain why contributions will be different across cooperative types.

An individual's predisposition towards cooperation or his cooperative type is argued to be arising from the so-called social preferences (Fehr & Schmidt, 1999; Bolton & Ockenfels, 2000; Charness & Rabin, 2002; Cox & Sadiraj, 2007; Cox, Friedman, & Sadiraj, 2008; Arifovic & Ledyard, 2012). According to the models of social preferences, an individual's cooperative type is something that emerges in the equilibrium due to the interaction between his social preferences and the parameters of the public goods game. There is empirical evidence that an individual's cooperative type can change when the price of cooperation changes (Cartwright & Lovett, 2014). In theory, social preference models can account for this change in the cooperative type of an individual when prices change. An individual with given social preference parameters can be a free rider when the price of cooperation is sufficiently high and can behave as a full cooperator when the price of cooperation is sufficiently low and vice versa. Other transitions are also possible. By allowing type transitions, models of social preferences provide quantitative predictions about how the proportions of the cooperative types change when prices change. In this Chapter, I empirically test if a model of social preferences can indeed explain the observed changes in the distribution of types when the price of cooperation changes using the data generated from an experiment.

The findings from Fischbacher and Gächter (2010) imply that a possible reason for why belief learning models performed poorly in comparison with reinforcement learning models in my earlier investigations in Chapter 1 could be that they were not considered alongside with social preferences. Without social preferences, belief learning models converge to zero contribution very fast since it is the only Nash equilibrium. But in the data, the movement towards zero is slow and contributions in the later rounds of a repeated game are multi-modal with a non-trivial probability mass at non-zero contribution levels. Directional learning as in belief learning models using social preferences can explain this since non-zero contributions can be sustained in equilibrium. In light of this, I re-evaluate representative models of learning by merging them with social preferences. I consider the reinforcement average model (REL) that was found to be the most successful in the previous chapter as the representative model of reinforcement learning. I consider Normalized Stochastic Fictitious Play (NFP) as the representative model of belief learning. For each of these models, I consider two variations: the first variation does not consider social preferences whereas the second variation merges learning with social preferences. By comparing the variations of learning models that consider social preferences with their counterpart variations that do not consider social preferences, I study the role of social preferences in explaining the contribution patterns. Another potential reason for the poor performance of belief learning models could be that individuals may be forming beliefs in a manner that is very different from that of how beliefs are formed in the mainstream belief learning models. By comparing the belief learning model, NFP, with a model that considers elicited beliefs (in an experiment) of individuals, I examine if the belief learning model performs poorly because it does not model the actual belief formation of individuals accurately.

To investigate if social preferences are able to organize changes in the distribution of cooperative types when prices change and to examine the performance of learning models in conjunction with social preferences, I use data generated from controlled laboratory experiments. The experiments are based on Fischbacher and Gächter (2010) and designed to observe the behavior of a given participant in four linear public goods game settings. These are a one-shot conditional contribution game (labeled the P-task) and a repeated game with partners (labeled the R-task) at two different costs to cooperate (i.e. marginal per capita return (MPCR) levels), one low and one high, that preserve the social dilemma. The one-shot game uses the strategy method to elicit an unconditional contribution and conditional contributions based on several possible average contributions of the other group members. Each one-shot game is then followed by a repeated game at the same MPCR level. In this game, for each participant along with a contribution decision in each round, I also elicited a belief about the average contribution of the other members of his group. The elicited beliefs allow for an examination of belief learning over rounds. The one-shot and repeated game sequence is then conducted again at a different MPCR level. In addition, across sessions but not within sessions, I varied whether the repeated games were finite or indefinite. The data from indefinitely repeated games in conjunction with data from finitely repeated games allows me to study if individuals learn differently across these two environments. Finally, the order of which MPCR is seen first is also reversed across sessions to control for experience effects.

The novelty of my approach here is two-fold. First, using the conditional responses in P-tasks elicited at two different prices the experimental design affords the possibility of testing if a single distribution of social preferences can explain the conditional choices of the same set of individuals at two different prices. In other words, the design allows me to test if one distribution of social preferences can explain changes in the distributions of cooperative types when prices change. A further innovation in this work is the structural estimation of social preferences separate from decision error. Much of the research on social preferences either assumes no error or does not explicitly model it at the individual level. If individuals do make errors, then not accounting for these in the estimations would produce biased results. The bias could be particularly pronounced in environments such as linear public goods games where both pure self-interested (free riding) behavior and full cooperation would be underestimated due to errors because any deviation from these two extremes (zero and full contribution) would be interpreted as an evidence against those behaviors. I also incorporate heterogeneity in the tendency to make random optimization errors in my estimations, and this allows for cleaner estimates of social preferences. Second, separate estimation of social preferences of participants from P-tasks allows me the opportunity to disentangle the role of social preferences and learning in explaining their choices in repeated games. One of the challenges in previous research in disentangling the roles of social preferences and learning was that they estimated both social preferences and learning parameters together using only the repeated game data. I overcome this challenge by measuring social preferences separately using the data from P-tasks.

This Chapter has a number of results. First, the structural estimation of social preferences using the data from P-tasks shows that the distribution of social preference parameters needs to be shifted significantly to explain the distributions of cooperative types under different prices. However, when the effect of experience is accounted for, the model is able to predict the direction and quantity of change in the proportion of different types when prices change reasonably well. The effect of experience on social preferences is well documented in the earlier literature and I also find similar results here that with experience individuals move towards selfish behavior (Brosig, Riechmann, & Weimann, 2007; Sass & Weimann, 2012). Once the effect of experience is accounted for, the social preference model is able to explain the general direction of change in the distribution of cooperative types when prices change. However, there are some minor discrepancies in the quantities of type transitions. Second, when learning models are merged with social preferences they perform significantly well with the exception of reinforcement learning. The gains in the explanatory power of the belief learning model, NFP, are significant when it is merged with social preferences. A model of elicited beliefs and social preferences perform significantly better than NFP indicating that elicited beliefs capture belief formation more accurately compared to that of NFP. Reinforcement learning model, REL, performs at least as well as the model with elicited beliefs and social preferences and outperforms NFP. Interestingly, merging the learning model with social preferences does not improve the fit of REL significantly. In other words, it does not make a difference in terms of the fit if the attractions of the strategies in REL model are updated using pure payoffs in the game or utilities computed with social preferences. Third, I show that first round choices in the repeated games are significantly different across cooperative types. In other words, they are not random but are informed by social preferences. I find that choices in the repeated games are best described by social preferences affecting the choice of first round contributions and then subsequent contributions based on payoff-based reinforcement learning. Differences in the first round contributions and high inertia of reinforcement learning explain the differences in average contributions among different cooperative types. Fourth, I find that learning is identical in both finitely and indefinitely repeated games. This finding explains identical contribution patterns across finitely and indefinitely repeated public goods games as found in earlier studies (Tan & Wei, 2014; Lugovskyy et al., 2015). Finally, using a dialogue between agent-based models and behavioral experiments I show that subsidizing cooperation in the first round alone could be sufficient to sustain significantly higher average levels of cooperation in the later rounds.

My results contribute to the ongoing discussion of the relative importance of social preferences and learning in explaining contributions in repeated public goods games. Learning in repeated public goods games is hypothesized to be individual level directional learning based on social preferences (Anderson, Goeree, & Holt, 2004; Wendel & Oppenheimer, 2010; Cooper & Stockman, 2002; Janssen & Ahn, 2006; Arifovic & Ledyard, 2012). Recent experimental results, however, have found that social preferences may be unnecessary to explain contributions in repeated games and that individuals learn based on pure payoffs in the game (Burton-Chellew & West, 2013; Burton-Chellew, Nax, & West, 2015). My analysis shows that the importance of social preferences may be overestimated in the former and underestimated in the latter. The results in this Chapter show that social preferences matter to the extent that they determine the first round contributions, thereafter, individuals' behavior is explained by pure payoff-based reinforcement learning.

The Chapter proceeds as follows. In Section 2.2, I discuss the experimental design. Section 2.3 presents a descriptive analysis of the data. Section 2.4 describes the social preference and learning models, and section 2.5 outlines the econometric model that is used in the estimations. In Section 2.6, I present results. I close with conclusions in Section 2.7.

2.2 Experimental Design & Procedures

The experiment uses a linear public goods game. Each individual i is assigned to a group with N members and is given an endowment w that can be invested in a public good account and a private account. The sum of contributions by all group members to the public good account is multiplied by an enhancement factor, E, where $1 < E < N$, and the resulting amount is redistributed equally among all the members of the group. Player i's contribution to the public good account is $c_i \in [0, w]$, and i's payoff is then

$$\pi_i = w - c_i + M \sum_{j=1}^{N} c_j$$

where $M = \frac{E}{N}$ is the marginal per capita return (MPCR) from contributing one unit from the endowment to the public good account. The social dilemma is evident. For any given level of contribution by i's group members, a payoff-maximizing individual's best response is to contribute nothing. Thus Nash equilibrium predicts a zero contribution from everybody in the group. However, a Pareto optimum is achieved when everybody contributes the entire endowment to the public good. In all decisions, there are three individuals in each group, and the endowment is 20 tokens. All decisions were made with tokens, and converted into a monetary payoff at a rate of 20 tokens = $1.

The experiment is based on the Fischbacher and Gächter (2010) design. Each participant is asked to complete four tasks in the following order: P1, R1, P2, R2. In the $P - task$, the participant completes a conditional contribution table in which he decides how much to contribute to the public good account for 21 possible average contribution amounts of his group members (e.g. 0, 1, 2, ..., 20). He also makes an unconditional contribution decision by choosing how many of the 20 tokens to contribute to the public good. Payoffs from this task are determined as follows. Once all decisions are completed by the group members, two group members are chosen at random. The unconditional contributions of those two group members are averaged together and rounded up or down to the nearest integer k. The contribution of the third group member is determined by the amount specified in the conditional contribution table for the average contribution amount of k by the other two group members. The total amount then contributed to the public good account is the sum of the unconditional contributions of the first two group members and the conditional contribution of the third. Participant payoffs are based on this. Everyone knows these procedures before making their decisions.

In the $R - task$, participants make an unconditional decision of how many tokens to put in the public good account. This decision is repeated over several rounds, and the members of a group are fixed for all rounds in an R-task. This is a partners matching protocol. After deciding how much to contribute to the public good account, a participant is asked to state his belief of the average contribution of the other two group members in the current round. Participants were paid for the accuracy of this stated belief.[1] At the end of a round, a participant is informed of the exact contribution of each group member, the average contribution, and his payoff for that round.

In sessions with finitely repeated games, the number of rounds was fixed at seven. In the sessions with indefinitely repeated games, there was at least one round and then after that the probability of a subsequent round was 0.85. The continuation probability of 0.85 yields, on average, seven rounds of play so the finitely repeated games was fixed at seven rounds to make the games comparable in the expected number of rounds.

Each session had exactly 15 participants and thus 5 groups in each of the tasks. Participants were randomly reshuffled across groups before each of the tasks and subjects were aware of this before making decisions. Tasks P1 and R1 use the same MPCR level, and tasks P2 and R2 also use the same MPCR level. The MPCR used for P1 and R1 is different than that used for P2 and R2. The two MPCR levels are 0.4 (Low) and 0.8 (High). In five sessions, the low MPCR was used for P1 and R1 and the high MPCR was used for P2 and R2. To control for order effects, in five sessions, this was reversed. Each participant completed all four tasks. Participants knew that there would be four tasks, and the instructions for each task were distributed and read out loud prior to the start of each task.[2] To make

[1]Specifically, the individual was asked to guess the average contribution of the two other group members rounded to the nearest integer. If the guess was exactly equal to the rounded average contribution of the group members, an individual was paid three tokens ($0.15) in addition to the other experimental earnings. If the guess deviated by only one point, payment was 2 tokens, and if the guess deviated by two points, payment was 1 token. If the guess was off by three points, no tokens were paid. The financial incentives to elicit beliefs were small to avoid hedging.

[2]Participant instructions for sessions with indefinitely repeated games are in Appendix 2.8.1. These include screen shots of the decision screens participants used to make decisions during the experiment. The instructions for the sessions with finitely repeated games are identical except that they mention the repeated games last exactly seven rounds. The instructions for finitely repeated games are available online at *http://www.chennacotla.org/research*

sure participants understood the decisions they are asked to make and how to calculate payoffs, a short quiz was administered prior to the start of Task P1. Answers to the quiz were explained before proceeding to the experiment. Participants were paid their earnings for all tasks and all rounds within a task. Earnings were paid privately and in cash at the end of the session.

Participants made their decision on a computer, using a web-based software. The experiments were run in the Interdisciplinary Center for Economic Science (ICES) at George Mason University during September and October of 2014. Ten sessions were run, and there were a total of 150 participants. Participants were randomly assigned to cubicles and made their decisions privately and anonymously. Six of the ten sessions involved indefinitely repeated games and four sessions involved finitely repeated games. No one participated in more than one session. Participants were recruited via email from a pool of George Mason University students who had all previously registered to receive invitations for experiments. Each experimental session lasted for approximately 1.5 hours. Average participant earnings were \$26.36 (s.d. \$8.67).

2.3 Data

In this section, I describe the conditional contribution decisions from the P-tasks and the contribution and belief decisions from the R-tasks.

To examine the conditional contributions in the P-tasks, I classified participants into cooperative types using the statistical classification algorithm of Kurzban and Houser (2005). This algorithm uses a linear conditional-contribution profile (LCP) to determine a given participant's type. The LCP is the result of an ordinary least squares regression of a participant's conditional contribution in the P-task on each of the 21 possible average contributions of the other group members. If the estimated LCP is strictly below half the endowment everywhere, then the participant is classified as a Free Rider (FR). A participant is classified as a Full Cooperator (FC) if the LCP lies at or above half the endowment

Table 2.1: Distribution of cooperative types in strategy games computed using the LCP method.

Type	P1	P2	Low MPCR (0.4)	High MPCR (0.8)
Free Riders	36 (24%)	54 (36%)	54 (36%)	36 (24%)
Conditional Cooperators	89 (59%)	69 (46%)	71(47%)	87 (58%)
Full Contributors	17 (11%)	20 (13%)	18 (12%)	19 (13%)
Noisy Contributors	8 (5%)	7 (5%)	7 (5%)	8 (5%)
Total	150 (100%)	150 (100%)	150(100%)	150 (100%)

everywhere. If the LCP of a participant has a positive slope and lies both above and below half the endowment then he is a Conditional Cooperator (CC). Any participant who did not fall into one of these three categories is classified as a Noisy Contributors (NC).[3]

The distribution of types identified with the classification algorithm is presented in Table 2.1. The first two columns present the distributions observed in the P1 and P2 tasks. The third and fourth columns present the distributions observed in the low MPCR and high MPCR P-tasks respectively. Consistent with previous studies, most participants are classified as conditional cooperators, and roughly one quarter are classified as free riders.[4] These two types account for more than 80% of participants in the combined P1 and P2 tasks and also when separated by the high and low MPCR treatments. Full cooperators and noisy contributors are less frequent. From the P1 to P2 task, as experience increases, there is a significant change in the distribution of types, with an increase in the proportion of free riders and a decrease in the proportion of conditional cooperators (Chi-Squared Test, $\chi^2(3) = 12.39; p = 0.006$). From the low to high MPCR, as the cost to cooperate declines, there is also a significant shift in the distribution, with the proportion of full cooperators rising (Chi-Squared Test, $\chi^2(3) = 12.12; p = 0.007$). A social preference model should be able to explain the transition of the distribution of types when prices change. And, I test

[3] In Kurzban and Houser (2005), the authors classify individuals into only the first three groups and exclude analysis on three participants who did not fall into these categories.

[4] Previous studies include Fischbacher et al. (2001); Fischbacher and Gächter (2010); Kurzban and Houser (2001, 2005); Burlando and Guala (2005); Duffy and Ochs (2009); Kocher et al. (2008); Herrmann and Thöni (2009); Muller et al. (2008); Bardsley and Moffatt (2007)

this in the later parts of this chapter.

Looking at the contribution behavior of each type, in P1, the average contribution across all conditional contributions is 2.39 tokens for free riders, 9.66 for conditional cooperators, 18.10 for full contributors and 10.45 for noisy contributors. In P2, the average contribution is 2.47 for free riders, 9.69 for conditional cooperators, 17.53 for full contributors, and 9.80 for noisy contributors. In both P1 and P2, average contributions across types are significantly different (Kruskal-Wallis tests, $p < 0.0001$).

Turning now to the R-task, four of my ten experimental sessions involved finitely repeated games of seven rounds each. The remaining six sessions used indefinitely repeated games. The probabilistic continuation rule produced rounds of the following length for R1 and R2 across the six sessions: {{6,7}, {3,25}, {7,12}, {5,2}, {9,5}, {6,5}}. While the finite and indefinitely repeated games had different lengths, I find no significant difference in average contributions across these games.[5] The high MPCR results in higher contributions in both the finitely and indefinitely repeated games, and this is consistent with previous research (Isaac, Walker, & Thomas, 1984; Lugovskyy et al., 2015).[6] Consistent with previous studies on repeated public goods games, contributions decline over rounds in both finitely and indefinitely repeated games. When contributions are regressed over round number the coefficient on the round number is negative and significantly different from zero as shown in Table 2.2. The Table also shows that the coefficients on Finitely Repeated dummy and its interaction with round number are not significant.

I also find that elicited beliefs are identical across finitely and indefinitely repeated games.[7] The round level contributions and beliefs are positively and significantly correlated in all ten sessions, as also found in previous studies (Fischbacher & Gächter, 2010; Croson,

[5] This was also found by Lugovskyy et al. (2015). I test for the effect of indefinite repetition by clustering at the session level and find no effect on contributions ($F(1,9) = 0.24, p = 0.634$). Identical conclusions are reached using clustering at the participant level. Since the participants are shuffled across 3-person groups at the start of each of the strategy and repeated games, I favor the session level clustering to report the results

[6] All sessions: $F(1,9) = 16.78, p = 0.0027$; Finitely Repeated: $F(1,3) = 14.96, p = 0.0306$; Indefinitely Repeated: $F(1,5) = 13.31, p = 0.0148$

[7] I test for the effect of indefinite repetition by clustering at the session level and find no effect on beliefs ($F(1,9) = 0.37, p = 0.557$)

Table 2.2: Contributions over rounds

	Contribution	
	(1)	(2)
Constant	11.79***	11.58***
	(0.92)	(0.78)
Round	-0.17***	-0.16**
	(0.049)	(0.051)
Finitely Repeated		0.22
		(1.40)
Finitely Repeated × Round		0.056
		(0.13)
N	2220	2220
R^2	0.014	0.015
F	11.92	20.91
$Prob > F$	0.0072	0.0002

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Robust standard errors clustered at the session level are reported. There are in total 10 sessions.

2007; Weimann, 1994).[8] Figures 2.1 and 2.2 show session level average contributions and beliefs in finitely and indefinitely repeated games. Appendix 2.8.2 provide plots of average contributions and beliefs across groups.

2.4 Social Preferences and Learning Models

In this section, I describe the utility specification and learning models that forms the basis of the econometric analysis of social preferences and learning. I use the utility function of Arifovic and Ledyard (2012) in the basic specification because it accommodates social preferences that include altruism and concern with one-sided fairness. I then proceed to a description of learning models that I use to explain the repeated game data.

[8]The correlations in each of the ten sessions are: 0.61, 0.81, 0.73, 0.50, 0.76, 0.82, 0.41, 0.62, 0.62, and 0.72. Spearman rank correlation tests, $p < 0.0001$.

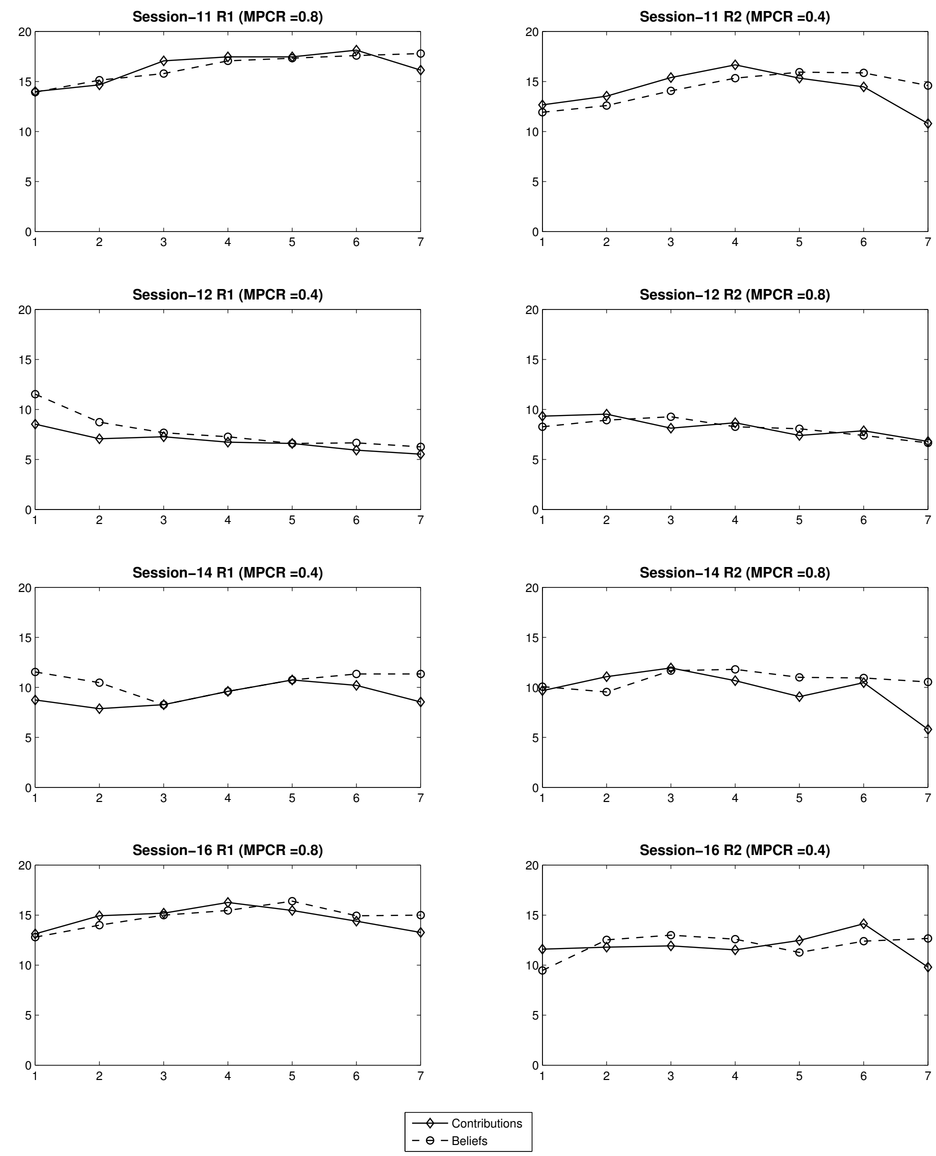

Figure 2.1: Contributions and beliefs in finitely repeated sessions. Average contributions and beliefs are reported for each round in a session.

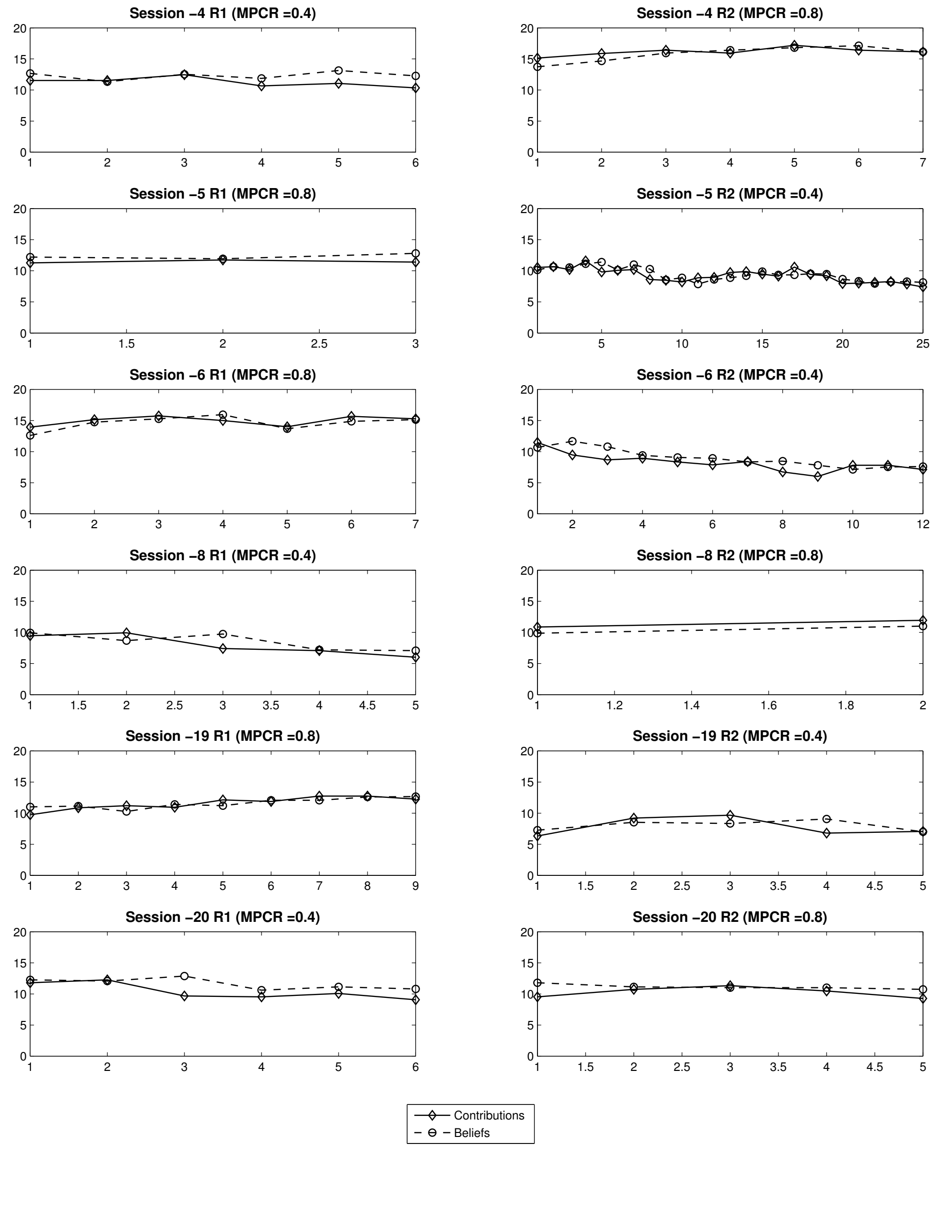

Figure 2.2: Contributions and beliefs in indefinitely repeated sessions. Average contributions and beliefs are reported for each round in a session.

2.4.1 Arifovic-Ledyard Model of Social Preferences

I use the linear utility specification with social preferences as developed in Arifovic and Ledyard (2012) for the context of public goods games. Consider a group of size N and an MPCR level M. Each individual $i \in \{1, 2, ...N\}$ is endowed with w. The payoff an individual i receives by contributing c^i when others in his/her group contribute on average o can be written as $\pi^i(c^i, o)) = w - c^i + M(c^i + (N-1)o)$. Similarly, the average payoff of the group can be written as $\bar{\pi}(c^i, o) = w - \bar{c} + MN\bar{c}$, where $\bar{c} = \frac{c^i+(N-1)o}{N}$. The utility derived by the individual i:

$$u^i(c^i, o) = \pi^i(c^i, o) + \beta^i \bar{\pi}(c^i, o) - \gamma^i max\{0, \bar{\pi}(c^i, o) - \pi^i(c^i, o)\} \tag{2.1}$$

Where $\beta^i \geq 0; \gamma^i \geq 0$ are social preference parameters. $\beta^i > 0$ implies that individual i has a preference for a higher average payoff to all agents in the group and thus higher welfare of group members. In other words, β^i characterizes an individual's altruistic preference. $\gamma^i > 0$ implies that individual i obtains a disutility when his/her payoff is smaller than the average payoff of the group, i.e. when $\bar{\pi}(c, o) > \pi^i(c, o)$. γ^i captures the discomfort the individual i face when being taken advantage of by the group.

In the equilibrium individual i would choose a contribution c^i as follows:

$$c^i = \begin{cases} 0 & \text{if } 0 \geq \left(M - \frac{1}{N}\right)\beta^i + M - 1 \\ \bar{c} & \text{if } \gamma^i\left(\frac{N-1}{N}\right) \geq \left(M - \frac{1}{N}\right)\beta^i + M - 1 \geq 0 \\ w & \text{if } \gamma^i\left(\frac{N-1}{N}\right) \leq \left(M - \frac{1}{N}\right)\beta^i + M - 1 \end{cases} \tag{2.2}$$

The social preference parameters (β^i, γ^i) along with the parameters of the public goods game (N, M) determine his cooperative type, i.e. if individual i behaves as a free rider or a conditional cooperator or acts as a pure altruist (full cooperator).[9] Figure 2.3 illustrates

[9]While social preference models of Fehr and Schmidt (1999) and Charness and Rabin (2002) also capture

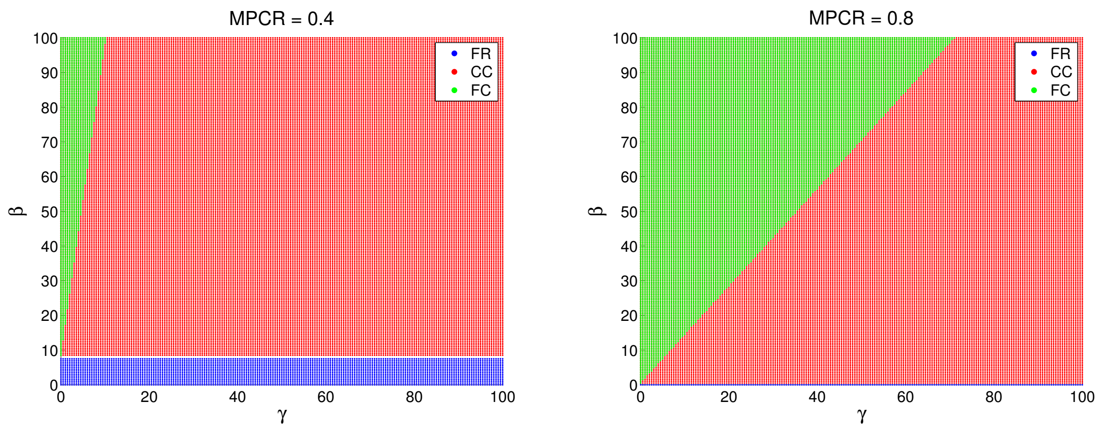

Figure 2.3: Distribution of cooperative types across two different MPCR levels. Group size is 3. β, γ are distributed according to a uniform distribution over $[0, 100] \times [0, 100]$. Left panel corresponds to the case when MPCR $= 0.4$. Right panel corresponds to the case when MPCR $= 0.8$.

how social preference parameter values determine cooperative types under two different MPCR levels $M = 0.4$ and $M = 0.8$ that are considered in this paper when the group size is $N = 3$. In the figure, I consider β and γ distributed independently according to a uniform distribution over [0, 100] for illustration purposes. The left panel of the figure corresponds to the case when M $=$ 0.4 and right panel of the figure corresponds to the case when M $=$ 0.8. In both panels, the area colored in green refers to the social preference parameter pairs that correspond to full cooperator behavior. The area colored in red refers to the preference parameter pairs that correspond to conditional cooperator behavior. The area colored blue refers to the preference parameter pairs that correspond to free rider behavior.

It is easy to see that for some pairs of social preference parameters, there is a switch in the behavior when M changes. For example, consider a participant with $\beta = 2$ and $\gamma = 1$ in our experiments where $N = 3$, and $M \in \{0.4, 0.8\}$. In the equilibrium this participant will give nothing when $M = 0.4$ and thus acts as a free rider. However, when $M = 0.8$

the effect of MPCR (M) on the cooperative type of an individual, they suggest behavior is independent of group size. Thus, they do not pick up the variation in observed contributions across small and larger groups as identified in Isaac et al. (1994); Isaac and Walker (1988). Other regarding preference model of Arifovic and Ledyard (2012) can explain this variation.

he will contribute all of the endowment and thus behaves as a full cooperator. Thus, the social preference model provides a theoretical foundation for the observed changes in conditional contribution preferences. The theory does not allow all possible type switches when MPCR changes. Appendix 2.8.3 presents the ranges of β, γ for different types at different levels of MPCR in our experiments. It also presents which type switches are possible and which are not when MPCR changes. Possible switches of types due to changes in prices can also be obtained by superimposing the left and right panels. The intersection of the areas corresponding to three types would indicate the possible type switches under the two MPCRs. Since the the parameters are assumed to be independently distributed according to uniform distribution, it can be noted that the colored areas also correspond to the proportions of the corresponding types. Thus the social preference model makes testable quantitative predictions of the distribution of cooperative types at different levels of MPCR. The econometric model setup in the next section tests the predictions of the social preferences model using data from our experiments.

2.4.2 Learning Models

I consider three learning models to model the learning mechanism of the participants in the repeated games. The three models are: Reinforcement Average Model with Loss Aversion Strategy (REL), Normalized Stochastic Fictitious Play (NFP), and Social Preference Based Empirical (Elicited) Belief Learning Model (AL). The description of REL and NFP is presented in Section 3 of Chapter 1. For each of these learning models, I estimate four variations: payoff based learning with no parameter heterogeneity, social preference based learning with no parameter heterogeneity, payoff based learning with parameter heterogeneity, and social preference based learning with parameter heterogeneity. In the payoff based version of a model, attractions of the strategies (contributions) in the learning model are updated using payoffs in the game. In the utility based model, attractions of the strategies are updated using utilities computed with social preferences. I use the distribution of social preferences estimated from strategy experiments (P-tasks) in computing the utilities. In

Chapter 1, all learning models were estimated from the standpoint of a representative agent meaning learning parameters were constrained to be identical for all individuals. However, it would be more realistic to assume that learning parameters are heterogeneous across individuals. Previous research has shown that allowing for heterogeneity in learning parameters improves the explanatory power of learning models significantly in public goods games (Wunder et al., 2013). By comparing the variations of the learning models that allow for parameter heterogeneity with the variations that do not allow it, I examine if parameter heterogeneity leads to significant gain in the explanatory power of learning models.

The new model introduced here, AL, assumes that the participants are stochastically best responding to their elicited beliefs according to the utility function in Equation 3.3. Thus, the learning mechanism is captured by the updating of beliefs about the average contribution of the other group members over time in the repeated game and best responding to the beliefs according to one's social preferences. By comparing Comparing AL with NFP, I examine if NFP perform poorly because they do not model actual belief formation of individuals accurately.

2.5 Econometric Framework

In this section, I formulate structural econometric models of discrete choice that can be estimated by maximum likelihood to estimate the social preference parameters and learning parameters allowing for heterogeneity. This is the appropriate approach when using the data generated in our experiments because choices are made on a discrete scale.

2.5.1 Estimation of Social Preferences from the Strategy Games

Based on the specification of the utility function that includes social preferences in Equation 3.3, I formulate structural econometric models of discrete choice that can be estimated by maximum likelihood. Let T_i is the number of decision situations an individual i has faced in the strategy games. Let $C_i = \{c_{it} | t \in \{1, 2, ..., T_i\}\}$ be the vector of observed contributions of the individual i and $O_i = \{o_{it} | t \in \{1, 2, ..., T_i\}\}$ be the vector of the average contributions of i's group members (excluding i itself). The average contributions of other group members are stated explicitly in the P-tasks.

I begin developing the econometric model by assuming the participants' decisions reflect maximization of utility function with social preferences specified in Equation 3.3. In the absence of any errors in decision making, for a given level of average contribution of others in his/her group, a participant chooses a contribution that maximizes his/her utility.

As a first step to allow for stochastic decision making, I add a standard extreme value distributed error term to the utility derived from each level of contribution. Assuming that these errors are independent of other parameters of the model and regressors, I obtain the logit probability of choosing a contribution level c_{it} as:

$$l_{it} = \frac{e^{U^{it}(c_{it}, o_{it})}}{\sum_{j=0}^{20} e^{U^{it}(j, o_{it})}}$$

The errors from a standard extreme value distribution that are added to the utility capture the idea that a participant's computation of subjective utility may be subject to some variability (Loomes, 2005). In addition to these errors, a number of experimental studies involving public goods games also found an evidence for the so-called "trembles" or "random choice errors" (Bardsley & Moffatt, 2007; Moffatt & Peters, 2001). These "trembles" account for a participant' failure to understand the decision problem or attention lapses during decision making. I model the propensity of a participant to chose randomly in any given task using a "trembling hand" parameter. A participants' tendency to make a random choice in a decision situation is given by a parameter ω_i.[10] Since each participant is endowed with 20 tokens, the probability of choosing a given contribution level via random choice is $\frac{1}{21}$.

[10]Alternatively, one could also model error heterogeneity at the individual level via individual level parameter $\tau_i \in R^+$ that scales the variance of the extreme value distributed error term. However, Hess and Rose (2012) show that scale heterogeneity and parameter heterogeneity are hard to identify separately when utility is linear in parameters

Then, for the participant i, the probability of observed level of contribution c_{it} for an average contribution o_{it} of his/her group members can be written as:

$$l_{it}(c_{it}, o_{it}, \beta_i, \gamma_i, \omega_i) = (1-\omega_i)\frac{e^{U^{it}(c_{it}, o_{it})}}{\sum_{j=0}^{20} e^{U^{it}(j, o_{it})}} + \frac{\omega_i}{21} \tag{2.3}$$

A random coefficient model is employed to estimate the distribution of the individual-specific structural parameters β_i, γ_i and ω_i in the population. This has a better justification than doing a separate estimation for each participant since the number of observed choices will be rather small. The random coefficient model proves to be more useful in studying treatment effects on preference parameters, since the numbers of observed choices per participant and experiment (strategy or repeated) combination are much smaller.

Since β, γ are constrained to be positive I modeled them using log-normal distributions.[11,12] To bound ω between 0 and 1, I modeled it as a logistic-normal distribution over [0 1]. For a concise notation define,

$$\eta_i = g_\eta(X_i^\eta \delta^\eta + \xi_i^\eta), \eta_i \in \{\beta_i, \gamma_i, \omega_i\}$$

η_i denotes one of the three individual specific parameters, X_i^η are $1 \times K^\eta$ vectors of regressors, δ^η are $K^\eta \times 1$ parameter vectors, and ξ_i^η are the unobserved heterogeneity components of the parameters. The first element of each X_i^η contains 1. The functions $g_\eta(.)$ impose theoretical restrictions on the individual specific parameters. For β, γ it is the exponential function ensuring that they are positive. For ω, it is the logistic distribution function ensuring that ω is always between zero and one. $g(X_i\delta + \xi_i)$ stands for a vector of

[11] Arifovic & Ledyard (2012) consider a mixed distribution of social preference parameters. They assumed that the parameters are distributed identically and independently according to a distribution where the probability that $(\beta, \gamma) = (0, 0)$ is P and the rest of the distribution is $F(\beta, \gamma) = U([0, B]) \times U([0, G])$. Where where $U([0, D])$ is the uniform density on the interval [0, D]. I did not follow this specification here for two reasons. It is not intuitive to specify treatment effects on parameters using a uniform distribution specification. Furthermore, the log normal distribution being flexible than a uniform distribution can account for the presence of selfish individuals by putting mass close to zero for both parameters.

[12] Alternatively, (β, γ) can be modeled using a bounded logistic-normal distribution. However, this specification found to be difficult to empirically identify as shown in Appendix 2.8.6

the three functions.

I assume that $\xi_i = (\xi_i^\beta, \xi_i^\gamma, \xi_i^\omega)'$ follows a jointly normal distribution with a diagonal covariance matrix Σ independent of the regressors. The regressor matrix contains only ones in the minimal estimation. In the full estimation case, it contains a dummy for high MPCR, a dummy for repeated game treatment, and a dummy for experience which was one for the second strategy game (P2) and the second repeated game (R2).

The likelihood contribution of participant i can be written as:

$$l_i = \int_{\mathbb{R}^3} \left[\prod_{t=1}^{T_i} l_{it}(c_{it}, o_{it}, g(X_i\delta + \xi)) \right] \phi(\xi) d\xi \tag{2.4}$$

where l_{it} is the probability given in Equation 2.3 and $\phi(.)$ denotes the density of multivariate normal ξ. The above integral does not have a closed form solution. I approximate it using $R = 100$ Halton draws from ξ to obtain simulated likelihood (Train, 2009; Bhat, 2001). The simulated likelihood contribution of participant i is:

$$sl_i = \sum_{r=1}^{R} \frac{l_i(\xi_r)}{R}$$

The (simulated) log-likelihood is given by the sum of the logarithms of sl_i over all respondents in the sample. I maximized the log-likelihood function of entire sample using a two-step hybrid approach a multiple number of times as discussed in Liu and Mahmassani (2000) to avoid local maxima.[13] The variance-covariance matrix of the parameter estimates is computed using the *sandwich estimator* (Wooldridge, 2010). Standard errors are calculated using the sandwich estimator and treating all of each participant's choices as a single

[13]In the first step, I have employed a genetic algorithm to find parameters that maximize log-likelihood of the sample. Genetic algorithms are very effective in searching many peaks of likelihood function based on a rich "population" of solutions and thus reduce the probability of trapped into a local maximum. Since they do not require gradients to be computed, they are computationally very efficient for a global search of the parameters. In the second step, I used the solution of genetic algorithm as a starting point to Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm with numerical derivatives to maximize the log-likelihood function.

super-observation, that is, using degrees of freedom equal to the number of participants rather than the number of participants times the number of choices made. Standard errors for transformed parameters are calculated using the delta method.

2.5.2 Estimation of Learning Models

Estimation of the belief learning model with social preferences, AL, is straightforward. In R-tasks, beliefs about the average contribution of other group members are elicited from the participants. AL assumes that the participants stochastically best respond (subject to noise in decision making) to their beliefs in each round. The estimated distribution of social preferences obtained from P-tasks is used to specify social preferences of the same individuals in R-tasks. The estimation methodology is then identical to the estimation of random coefficients model described in the context of strategy game data except that I only estimate the distribution of the random choice propensity parameter. I left random choice choice propensity parameter as a free parameter since the repeated game environment being naturally more complex compared to the one-shot game in P-tasks could possibly involve a higher level of decision error. Allowing the random choice parameter to be different in the repeated environment captures this possibility.

Estimation of learning models, REL and NFP, without parameter heterogeneity follows the description in Section 4 of Chapter 1. Here, I describe in detail the estimation method used in the context of REL with parameter heterogeneity. The method can then be easily extended to other learning models when parameter heterogeneity is considered. The REL model has two structural parameters: the attraction sensitivity parameter λ_i and the parameter that defines the initial strength of attractions $N_i(1)$. I model both of them as random coefficients allowing for individual level heterogeneity. Using the notation in the previous subsection:

$$\eta_i = g_\eta(X_i^\eta \delta^\eta + \xi_i^\eta), \eta_i \in \{\lambda_i, N_i(1)\}$$

η_i denotes one of the two individual specific parameters, X^{η_i} are $1 \times K^{\eta}$ vectors of regressors, δ^{η} are $K^{\eta} \times 1$ parameter vectors, and ξ_i^{η} are the unobserved heterogeneity components of the parameters. The first element of each X_i^{η} contains 1. Since both $\lambda_i, N_i(1)$ are positive, I used the exponential function for g_{η} for both parameters. Assuming that $\xi_i = (\xi_i^{\lambda}, \xi_i^{N(1)})'$ follows a jointly normal distribution with a diagonal covariance matrix Σ independent of the regressors, the likelihood contribution of participant i can be written as:

$$l_i = \int_{\mathbb{R}^2} \left[\prod_{t=1}^{T_i} \left(\sum_{l=0}^{20} \mathbb{I}(c_{it}, l) P_i^l(t) \right) \right] \phi(\xi) d\xi \tag{2.5}$$

Where

$$P_i^l(t) = \frac{e^{\frac{\lambda_i}{PV_i(t)} A_i^l(t)}}{\sum_{k=0}^{20} e^{\frac{\lambda_i}{PV_i(t)} A_i^k(t)}}$$

is the probability of choosing contribution level l in round t. c_{it} is the observed contribution in t. $\mathbb{I}(c_{it}, l) = 1$ if $c_{it} = l$, 0 otherwise. T_i is the total number of rounds in the repeated games R1 and R2 that individual i has participated in (it should be noted that attractions, payoff variability, and accumulated payoff average in REL model will be reinitialized at the start of the R2). The integral in Equation 2.5 is computed using simulation and the total log-likelihood of the sample is computed as the sum of the logarithms of simulated individual level likelihoods of all respondents. I maximized the log-likelihood function of the entire sample using a two-step hybrid approach (described in the previous subsection) a multiple number of times to avoid local maxima. The variance-covariance matrix of the parameter estimates is computed using the *sandwich estimator*.

The NFP model has two structural parameters: the attraction sensitivity parameter $\lambda_i \in \mathbb{R}^+$ and the experience weighting parameter $w_i \in [0, 1]$. I model both of them as random coefficients. λ_i is modeled as a log-normal distribution and w_i is modeled as a logistic-normal distribution. Both of these distributions are assumed to be independent of each other and also are independent of any regressors. The estimation procedure is identical to the simulated maximum likelihood method described in the context of REL model.

2.6 Results

I present the results in four stages. First, using the structural econometric approach outlined in the previous section, I examine how estimated social preferences from decisions in a one-shot environment, using data from the P-tasks, are affected by the cost to cooperate and experience. Second, I examine the choices solely in the repeated game environment by comparing variations of representative learning models that differ in terms of whether they allow for parameter heterogeneity and whether they consider social preferences. Third, I conduct a study in which I assess the roles of social preferences and learning in explaining choices in the repeated game environment. I conclude by studying if learning is influenced by the type of repetition.

2.6.1 Cooperative Types and Social Preferences in the One-Shot Environment

Conditional choices in the P-tasks provide data to directly test if the social preferences model that is outlined earlier can organize changes in cooperative disposition due to changes in the MPCR. The decisions made in this task do not require computations of beliefs about others' contributions since others' contributions are explicitly presented. Choices should then reflect best responses given social preferences as accurately as possible. Of course, decisions errors are possible, and the estimation approach takes this into account.

Table 2.3 reports the estimates for the three parameter vector $\delta^{\eta}, \eta \in \{\beta, \gamma, \omega\}$ for both minimal and full estimation cases. The minimal model contains only intercept terms for the parameters. The full model contains an intercept and also treatment dummies for experience and high MPCR for each parameter. The values in the table are in the original parameters'

Table 2.3: Distribution of social preferences (β, γ) and random choice propensity (ω) parameter estimated using data from strategy games.

	β		γ		ω	
	Min	Full	Min	Full	Min	Full
Constant	10.23***	31.30***	11.45***	8.89***	0.12***	0.0042
	(0.80)	(3.62)	(0.91)	(0.79)	(0.017)	(0.0048)
Experience		-22.21***		0.11		-0.0042
		(2.75)		(1.20)		(0.0048)
High MPCR		-25.89***		-6.26***		0.086
		(3.15)		(0.76)		(0.14)
σ	3.09***	2.21***	1.79***	2.06***	3.34***	13.99***
	(0.029)	(0.056)	(0.016)	(0.050)	(0.24)	(2.13)
N	150	150	150	150	150	150
LL	-13625	-12902	-13625	-12902	-13625	-12902

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Coefficients are transformed back to the original scale. A treatment effect is the (partial) effect of moving from the baseline treatment to the corresponding treatment on the median parameter value. Standard errors are calculated using the *sandwich estimator* and treating all of each subjects choices as a single *super observation*. The entries for σ are the standard deviations of the untransformed normal distributions of the random coefficients.

scale. In the full estimation, the constant term is $g_\eta(\delta_1^\eta)$ and represent the median parameter value for the baseline treatment, which is the P1 task with a low MPCR. Treatment effects reported in the table are the calculated partial effects of moving from the baseline treatment to the treatment under consideration using the median parameter value. For example, the value shown in the table for the high MPCR dummy is $g^\eta(\delta_1^\eta + \delta_{HighMPCR}^\eta) - g_\eta(\delta_1^\eta)$ and represents the partial effect of moving from the baseline to the high MPCR treatment on the parameter $\eta \in \{\beta, \gamma, \omega\}$.[14]

According to the minimal model the medians of the estimated β, γ are $10.23, 11.45$. Using the categorization outlined in Appendix 2.8.3, a median participant is a conditional cooperator at both MPCR levels M = 0.4 and M = 0.8. This is in line with the classification results using the LCP method presented in Section 2.1. The median random choice

[14]For illustration, in the estimations I have, $\delta_1^\beta = 3.44$ and $\delta_{HighMPCR}^\beta = -1.75$. Note that g^β is the exponential function. The corresponding treatment effect of the high MPCR on the median value of β is computed as $e^{(3.44-1.75)} - e^{3.44} = -25.89$ and is reported in the table.

Table 2.4: Distribution of cooperative types computed from the estimated distribution of social preferences from strategy games. Full refers to the case where the proportions of types are computed from the distribution of social preferences derived from the full estimation. Min refers to the case where the proportions of types are computed from the distribution of social preferences derived from the minimal estimation.

	P1-Low		P1-High		P2-Low		P2-High	
	Full	Min	Full	Min	Full	Min	Full	Min
Free Riders	19%	46%	4 %	4%	49%	46%	19%	4%
Conditional Cooperators	54%	42%	42%	56%	41%	42%	52%	56%
Full Contributors	27%	12%	54%	40%	10%	12%	29%	40%

propensity parameter is 12%, which means 12% of the choices in P-tasks can be categorized as being made randomly. According to the estimates for the full model, there are significant effects of MPCR and experience on the median values of social preference parameters. This indicates that the distribution of social preferences need to be shifted significantly to explain the changes in the proportion of cooperative types (free riders, conditional cooperators, and full cooperators) due to experience and due to changes in prices. Or, in other words, a single distribution of social preferences estimated in the minimal model is not able to account for the changes in the distribution of cooperative types due to the changes in experience and prices. There is also a strong effect of MPCR on random choice propensity parameter.

To understand the estimates of treatment effects on the social preference parameters, it is important to keep in mind that the distribution of cooperative types is determined by the joint distribution of β, γ conditional on the group size (N) and the MPCR (M).[15] Table 2.4 shows the proportions of types by MPCR and experience level using the estimates in Table 2.3.[16]

[15]The social preference parameter ranges that lead to different type behaviors are reported in Appendix 2.8.3.

[16]These proportions are determined in a given treatment using the estimated joint distribution of β and γ in that treatment. The joint distribution of β and γ is identical across treatments when results from the minimal model are used. When the estimates from the full model are considered, the distributions of β and γ in a treatment are adjusted based on the effects of experience and MPCR. The proportions of different types are computed using Appendix 2.8.3. For example, the proportion of free riders in P1-Low is computed as the probability that $\beta \leq 8.57$, the proportion of conditional cooperators in P1-Low is computed as the probability that $\beta > 8.57$ and $\beta \leq 8.57 + 9.43\gamma$, and finally, the proportion of full contributors in P1-Low is computed as the probability that $\beta \geq 8.57 + 9.43\gamma$ given the joint distribution of β and γ for P1-Low.

Table 2.4 shows that the estimation results from the full and minimal models predict that the proportion of full contributors rises with an increase in the MPCR level. This is consistent with changes in the distribution of types based on the LCP classification algorithm reported in Table 2.1. However, the exact proportions of types predicted by the full and minimal model are different across treatments. The full model estimates show that the proportion of free riders substantially rises with experience, moving from 19% in P1-Low to 49% in P2-Low and from 4% in P1-High to 19% in P2-High. Thus, we observe an increase in selfish behavior with the experience similar to the findings of Brosig et al. (2007) and Sass and Weimann (2012). The minimal model cannot explain this effect of experience.

One of my primary objectives in this chapter is to test if the social preference model is able to explain the changes in the distribution of cooperative types when prices change. The estimates from the full model show that this is not the case since they show a significant effect of MPCR on the social preference parameters even after controlling for experience effects. In other words, the distribution of social preferences has to be shifted significantly to explain the changes in the distribution of cooperative types when the prices change. A potential candidate of explanation for the significant effect of MPCR on the social preference parameters can be the misspecification of the distribution of social preferences. I chose a bivariate log-normal distribution to model the distribution of social preferences in my estimations. There are two explanations for this choice. First, in my investigations with a bounded logistic-normal distribution specification for social preferences parameters I faced empirical identification issues. Appendix 2.8.6 illustrates the empirical identification issues with logistic-normal specification for social preference parameters. Second, when I considered a uniform distribution for social preference parameters, the achieved fit was significantly worse compared to that of the case when I considered the log-normal distribution specification. Thus, I preferred the log-normal specification over these two alternatives. But, a numerous other specifications can be used to model the distribution of social preference parameters. Evaluating all possible specifications and identifying the least misspecified specification for the distribution of social preferences is not a primary objective of this chapter and so I do not pursue that line of investigation here. Instead, given the log-normal specification assumption, I evaluate how worse off the minimal model is in predicting the changes in the distribution of cooperative types across different levels of MPCR compared to that of the full model that incorporates the effects of MPCR. Despite the significant effects of MPCR on the median values of the parameters, if the minimal model is able to capture the general direction of changes in the distribution of cooperative types one could use it as a practical unifying model of organizing different distributions of types under varying prices.

Table 2.5 presents the changes in the percentages of each cooperative type when one moves from low MPCR (0.4) to high MPCR (0.8). The table reports the changes in the percentages of types using the minimal model and full model. The changes in the percentages of types is reported for P1 and P2 tasks separately. This helps to understand separately how successful is the minimal model in explaining the changes in the distribution of types due to the changes in MPCR when participants do and do not have experience. In both P1 and P2, the full and minimal models predict that percentage of free riders decreases and the percentage of full cooperators increases when one moves from low MPCR to high MPCR. In P1, there is disagreement between the full model and minimal model in terms of change in the percentage of conditional cooperators. There is a substantial difference in terms of the magnitude of change in the percentages of free riders and conditional cooperators between the full and minimal models in P1. However, both models provide a very close prediction of changes in percentages of types in P2. Thus, if one considers that with experience participants' responses are less noisy, the minimal model is quite good at describing how the proportions of cooperative types change when the MPCR changes.

Another way to evaluate the success of social preference model is to see how many individual level type transitions are consistent with the model. A type transition at the individual level is consistent if one can find at least one pair of social preference parameters according to which the transition is possible. Appendix 2.8.3 describes which individual level transitions are possible and which are not possible according to the social preference

Table 2.5: Changes in the percentages of cooperative types predicted by the full model and minimal model in P1 and P2 tasks across MPCR levels. Reported are the differences in percentages obtained by subtracting the percentage of a type at low MPCR from the percentage of a type at high MPCR.

	P1		P2	
	Full	Min	Full	Min
Free Riders	-15%	-42%	-30 %	-42%
Conditional Cooperators	-12%	14%	11%	14%
Full Contributors	27%	28%	19%	28%

model. The only individual level transitions that are not possible are the transitions from CC to FR, from FC to FR, and from FC to CC when MPCR changes from 0.4 to 0.8 (or the transitions in the opposite direction when MPCR changes from 0.8 to 0.4). A small number of these transitions in the data would support that social preference model is relatively successful. Table 2.6 shows the individual level types at the two MPCR levels in the experiment. Individual level type information is computed using the LCP method as described in Section 2.3. I have dropped the cases where the transitions involved a noisy type. In total, the table reports the type transition for 137 individuals in the data. There are only 12 cases where type transitions are not consistent with the social preference model. Five of them involve a transition from CC to FR when MPCR changes from 0.4 to 0.8, two of them involve a transition from FC to FR when MPCR changes from 0.4 to 0.8 and five of them involve transition from FC to CC when MPCR changes from 0.4 to 0.8 . Thus, only 8.7% of individual level transitions cannot be accounted by the social preference model indicating that the model is reasonably successful in organizing changes in behavior when prices change.

2.6.2 Comparing Models of Learning with Repeated Game Data

In this section, I conduct an investigation to build on the results in Chapter 1. In Chapter 1, all learning models were estimated from the standpoint of a representative agent. The

Table 2.6: Cross tabulation of types at two MPCR levels: $MPCR = 0.4$ and $MPCR = 0.8$. Types are computed using Linear Contribution Profile (LCP) method.

		$MPCR = 0.8$		
		FR	CC	FC
	Free Riders (FR)	43	22	5
$MPCR = 0.4$	Conditional Cooperators (CC)	5	59	5
	Full Cooperators (FC)	2	5	10

results indicated that REL outperforms all other learning models and the learning is more in line with choice reinforcement in public goods games. It remains an open question if the ordering of learning models changes if one allows heterogeneity in the learning model parameters. Furthermore, a number of previous studies have argued that social preferences are important alongside with learning in explaining contribution patterns in repeated games. It could be possible that belief learning models perform well when they are merged with social preferences. Belief learning models without social preferences imply a faster learning of the only Nash equilibrium contribution which is zero. Thus, contributions drop very fast to near zero. In the actual data, the movement is relatively slow. Social preferences will help the belief learning models not to drop to zero contributions since multiple Nash equilibria are possible as can be seen in Equation 3.4 and not all of them involve zero contributions. Finally, it is possible that individuals could be forming beliefs in a way different from that of what is captured in the belief learning models. I evaluate this possibility using empirical beliefs elicited in the data from the experiments in this Chapter.

I consider REL as the representative model of reinforcement learning and NFP as the representative example for belief learning. These models can help me to examine whether belief learning or reinforcement learning best explains the choices when parameter heterogeneity is allowed and when merged with social preferences. For each of these learning models, I estimate four variations: payoff based learning with no parameter heterogeneity, social preference based learning with no parameter heterogeneity, payoff based learning with parameter heterogeneity, and social preference based learning with parameter heterogeneity. In the payoff based version of a model, attractions of the strategies in the learning model are updated using the payoffs in the game. In the utility based model, attractions of the strategies are updated using the utilities computed with social preferences. I use the distribution of social preferences estimated from strategy experiments (P-tasks) in computing the utilities. I consider estimates from the minimal estimation in Table 2.3. The payoff based variation of a learning models is suffixed with PAYOFF and the social preferences based variation of a learning model is suffixed with UTIL. Finally, AL models individuals as best responding to their elicited beliefs and evaluates how well the elicited beliefs capture learning of individuals. Here, the best responses are computed by integrating over the distribution of social preferences estimated from P-tasks. AL has two free parameters which characterize the distribution of its random choice propensity parameter.

All the learning models are estimated using the aggregated repeated game data. The estimation methodology follows the description in Section 2.5. To keep all models on the same footing initial attractions of all strategies in the first round are set to zero. Thus in the first round, all choices have equal probabilities for being chosen according to all these models. I revisit this assumption in the next subsection. Table 2.7 reports the parameter estimates of each of the models and the corresponding fit metrics. I report the maximum log-likelihood (LL) achieved, AIC, BIC, and Pseudo-R^2 for each learning model.[17] The performance of a random choice model, RAND, is also reported for comparison purposes.

First, I compare the variations of REL and NFP that allow for parameter heterogeneity with the ones that do not allow for parameter heterogeneity. In all cases, the variations

---

[17]

$$AIC = LL - k$$

$$BIC = LL - \frac{k}{2} ln(Ntotal)$$

$$Pseudo\text{-}R^2 = \frac{AIC - LLR}{LLR}$$

Where k is the number of the parameters of the learning model, $Ntotal$ is the total number of observations, and LLR is the log-likelihood obtained by a random choice model. Since choices for an individual are not independent I consider the number of individuals as the effective size of the sample in the computation of BIC. Therefore, $Ntotal$ is equal to 150 here. AIC and BIC penalize models for their complexity and therefore are superior metrics than log-likelihood for model comparison.

Table 2.7: Comparison of learning models. UTIL means utilities computed using social preferences are used for attraction updating and PAYOFF means pure payoffs in the repeated game are used for attraction updating in the learning model.

Model	Parameter Estimates				LL	AIC	BIC	Pseudo-R^2
Models with Parameter Heterogeneity†								
AL	ν_ω	0.3377	σ_ω	1.3875	-5017	-5019	-5022	0.2574
REL-UTIL	ν_λ	1.3958	σ_λ	0.3449	-5099	-5103	-5109	0.2450
	ν_{N1}	0.0148	σ_{N1}	9.0662				
REL-PAYOFF	ν_λ	1.3850	σ_λ	0.3099	-5103	-5107	-5113	0.2444
	ν_{N1}	0.0966	σ_{N1}	5.9423				
NFP-UTIL	ν_λ	1.2110	σ_λ	2.1653	-5496	-5500	-5506	0.1863
	ν_w	0.0252	σ_w	2.1807				
NFP-PAYOFF	ν_λ	4.4674	σ_λ	156.10	-6305	-6309	-6315	0.0666
	ν_w	0.0010	σ_w	1.4160				
Models Without Parameter Heterogeneity								
REL-UTIL	λ	1.7866	$N1$	0.0435	-5222	-5225	-5228	0.2269
REL-PAYOFF	λ	0.9202	$N1$	10.8527	-5246	-5248	-5251	0.2234
NFP-UTIL	λ	1.9370	w	0.0114	-5546	-5547	-5551	0.1792
NFP-PAYOFF	λ	153448	w	0.0028	-6754	-6755	-6759	0.0004
Random Choice Benchmark								
RAND‡					-6759	-6759	-6759	0.0000

† ν stands for the median of the distribution of a parameter. Therefore, ν_λ refers to the median value of λ. σ stands for the standard deviation of the untransformed normal distribution of a random coefficient. Therefore, σ_λ is the standard deviation of the normal distribution that was exponentiated to obtain the log-normal distribution of λ.

‡ RAND is the random choice model. According to this model any choice has a probability of $\frac{1}{21}$ of being chosen.

Table 2.8: Comparison of learning models. UTIL means utilities computed using social preferences are used for attraction updating and PAYOFF means pure payoffs in the game are used for attraction updating in the learning model. For each pair of learning models, a two sided p-value favoring the learning model in the row is reported. The two-sided p-values are computed using the Vuong Test for Non-Nested Models.

	REL-UTIL	REL-PAYOFF	NFP-UTIL	NFP-PAYOFF	RAND
AL	0.2830	0.2550	0.0000	0.0000	0.0000
REL-UTIL		0.7467	0.0000	0.0002	0.0000
REL-PAYOFF			0.0002	0.0000	0.0000
NFP-UTIL				0.0000	0.0000
NFP-PAYOFF					0.0011

with parameter heterogeneity outperform the variations without parameter heterogeneity (Likelihood Ratio Tests, $p < 0.0000$).[18] Therefore, in the following discussion I consider only the variations of the models that allow for heterogeneity in learning parameters.

Comparing the variations of the models (that allow for parameter heterogeneity) that consider payoffs in the learning and utilities in the learning, it can be seen that considering utilities appear to improve the fit. The gains are very prominent in the case of NFP. AL achieves the best fit in terms of all four fit metrics. However, one needs to keep in mind that AL uses twice as much of the data as all other learning models use since AL uses elicited beliefs in each round along with choices. Since, all the models considered here, AL, REL-PAYOFF, REL-UTIL, NFP-PAYOFF, and NFP-UTIL are non-nested, I assess their performance in pairs using the Vuong's test for non-nested models (Vuong, 1989) as a model selection criterion. Table 2.8 presents p-values computed using the Vuong test for each pair of learning models. A two sided p-value favoring the learning model in the row is reported. AL, REL-UTIL, REL-PAYOFF perform equally well. This is significant since it indicates

[18] For each of the models, REL-UTIL, REL-PAYOFF, NFP-UTIL, and NFP-PAYOFF, a Likelihood ratio test is done between the two variations: the one that allows parameter heterogeneity and the one that does not. Note that the variation that does not allow parameter heterogeneity is nested within the model that allows parameter heterogeneity.

that the REL variations are able to perform as well as AL though AL uses twice the amount of the data compared to that of REL variations. The p-values computed for AL and REL provide an unfair advantage to AL since the Vuong test does not take into account the fact that AL uses twice the amount of data compared to that of REL. Interestingly, for REL it does not seem to matter whether attractions are updated using the payoffs or utilities. This is in contrast with the results from NFP. NFP-UTIL which uses the utilities computed with social preferences to update attractions perform significantly better. NFP-PAYOFF barely outperforms the random choice benchmark.

The results on this section reinforce my findings in Chapter 1. REL provides the best account of observed choices in repeated games. NFP performs significantly better when it is merged with social preferences. However, it does not beat the both variations of REL. For REL, the payoff based attraction updating works as well as the utility based attraction updating.

2.6.3 Explaining the Repeated Game Choices: Disentangling the Roles of Social Preferences and Learning

My analysis thus far has found that the payoff based and utility based REL models provide the best account of the repeated game data. Both of the REL's variations, REL-PAYOFF and REL-UTIL, are identical in terms of their explanatory power. Thus, the behavior in repeated games is more in line with reinforcement learning than that of belief learning.

An equally important observation from the repeated game data is that there are substantial differences in contributions across cooperative types. For example, using the LCP classification of types in P1-task as in Section 2.3, the average contribution of free riders, conditional cooperators, and full cooperators are 7.53, 12.45, and16.14 respectively in R1. The contributions are significantly different across types (Kruskall-Wallis Test: $p < 0.0001$), and there is a significant increasing trend in contributions across the groups (Jonckheere-Terpstra Test, $p < 0.0001$). Similarly, using the type information from the P2-task, the average contribution of free riders, conditional cooperators, and full cooperators are 6.30, 12.28, and 15.28 respectively in R2. The contributions are significantly different across types (Kruskall-Wallis Test: $p < 0.0001$), and there is a significant increasing trend in contributions (Jonckheere-Terpstra Test, $p < 0.0001$). This indicates that social preferences elicited in P-tasks have a nontrivial relevance in explaining the average contributions in the repeated games.

Earlier studies have argued that differences in contributions of different cooperative types occur because these types learn their equilibrium contribution as the game progresses (Cooper & Stockman, 2002; Janssen & Ahn, 2006; Wendel & Oppenheimer, 2010; Arifovic & Ledyard, 2012). This assumption would argue that REL-UTIL should perform better than REL-PAYOFF. Note that both REL-UTIL and REL-PAYOFF start with random contributions in the first round. But, my results in the previous subsection indicate REL-UTIL does not explain the data better than REL-PAYOFF. REL-PAYOFF makes no distinction among cooperative types and models all of them as learning based on payoffs. Thus, all cooperative types would be learning the only Nash equilibrium that is zero contribution. In this manner, my results are at odds with the assumption made in the earlier research.

It should be noted that REL involves high inertia and slow learning. The learning is much slower in public goods environments since there are no negative payoffs involved. One potential way in which the aggregate differences would be emerging across the different cooperative types could be due to the fact that different cooperative types may be starting with different first round contributions. This combined with high inertia of REL can account for aggregate differences in the contributions across the types. To investigate if this is indeed the case, I look into the first round contributions in repeated games across cooperative types computed using the LCP method. To begin, I compare the first round contributions in repeated games and the unconditional contributions in strategy games and find they are not significantly different from each other (Paired Hotelling's T^2 test, $p = 0.118$). In addition, the first round contributions in repeated games and the unconditional contributions in strategy games are significantly different across types (Kruskal Wallis

Tests: $p < 0.0001$ in all cases). Figure 2.4 presents the average contribution by cooperative types in the first round of a repeated game and the average unconditional contribution by cooperative types in the corresponding strategy game. Cooperative type information for P1 and R1 is computed using the conditional responses in P1 and the cooperative type information for P2 and R2 is computed using the conditional responses in P2. There is also a significant increasing trend in the first round contributions in the repeated games across free riders, conditional cooperators and full cooperators (R1: Jonckheere-Terpstra Test, $p < 0.0001$; R2: Jonckheere-Terpstra Test, $p < 0.0001$). The trend is also significant for the unconditional contributions in the strategy games across these three types (P1: Jonckheere-Terpstra Test, $p < 0.0001$; P2: Jonckheere-Terpstra Test, $p < 0.0001$). Thus, the first round contributions in the repeated game are similar to the unconditional contributions in the corresponding strategy game and are informed by individual level social preferences derived from the strategy game. In the first round of a repeated game, free riders tend to start with contributions that are significantly smaller on average (though not exactly zero) and full contributors start with significantly higher contributions on average (though not exactly full endowment). The remaining participants start around half of the endowment.

Thus, it is possible that aggregate differences across types are arising simply because they are starting at different levels of initial contributions and their learning in later rounds is characterized by high inertia like that of REL's. This stands in contrast with the assumption in earlier literature that the aggregate differences across types arise because types start with random contributions and learn their equilibrium contributions over time. It is also possible that aggregate differences are emerging because individuals start with different initial contributions and then they also proceed to learn equilibrium contributions. To be able to systematically disentangle the roles of social preferences and learning in repeated games, I consider four models that involve different assumptions about learning and social preferences.

- **RAND-PAYOFF**: REL learning based on payoffs in the game with random first

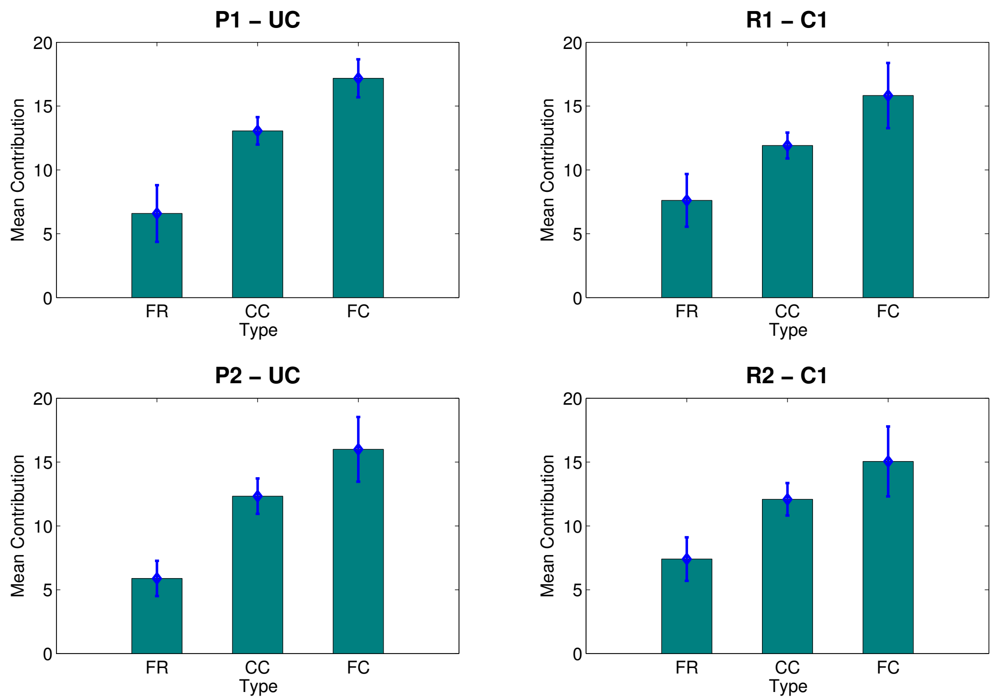

Figure 2.4: Unconditional contributions across types in P-tasks and the first round contributions across types in R-tasks. P1-UC refers to the unconditional contributions in P1 and P2-UC refers to the unconditional contributions in P2. R1-C1 refers to the contributions in the first round of R1 and R2-C1 refers to the contributions in the first round of R2. Cooperative types in P1-UC and R1-C1 are computed using the conditional choices in P1 and types in P2-UC and R2-C1 are computed using the conditional choices in P2. Error bars represent 95% confidence intervals.

round contributions. This model does not use any information about the social preferences of the participants. In the learning phase, individuals use the payoffs in each round to update the attractions of strategies in the REL learning model.

- **RAND-UTIL**: REL learning based on utilities in the game with random first round contributions. In the learning phase, individuals use the utilities computed with estimated distribution of social preferences from the P-tasks to update the attractions of strategies in each round in the REL learning model.

- **SP-PAYOFF**: REL learning based on payoffs in the game with the first round contributions determined from the social preference parameters of the participants. The first round contribution decisions of individuals are modeled as best responses to their elicited beliefs in the first round given the distribution of social preferences estimated from P-tasks. In the learning phase, individuals use payoffs in each round to update attractions of strategies in reinforcement learning model.

- **SP-UTIL**: REL learning based on utilities in the game with the first round contributions determined from the social preference parameters of the participants. The first round contribution decisions of individuals are modeled as best responses to their elicited beliefs given the distribution of social preferences estimated from the P-tasks. In the learning phase, individuals use the utilities computed with estimated distribution of social preferences from the P-tasks to update the attractions of strategies in each round in the REL learning model.

The parameter estimates and fit of the four models are reported in Table 2.9. Models whose first round contributions are determined by social preferences achieve higher likelihood compared to the models that use random contributions in the first round. It does not seem to matter much whether payoffs or utilities are used in updating the attractions in the REL learning model. Table 2.10 reports two-sided p-values computed using the Vuong's test for each pair of the four models. SP-UTIL and SP-PAYOFF perform equally well in explaining contributions. They outperform RAND-UTIL and RAND-PAYOFF models. RAND-UTIL and RAND-PAYOFF are indistinguishable from each other in terms of their performance. The results highlight the fact that initial contributions determined by social preferences combined with high inertia associated with REL learning model explain the aggregate differences in contributions across cooperative types. This result stands in contrast to the assumption in the existing literature that individuals start with random contributions and then move towards their equilibrium contributions based on learning with utilities derived from social preferences (Cooper & Stockman, 2002; Janssen & Ahn, 2006; Wendel & Oppenheimer, 2010; Arifovic & Ledyard, 2012). No significant difference between the

Table 2.9: Disentangling the role of social preferences and learning. REL is used to model learning. SP means social preferences determine the first round choices and RAND means the first round choices are drawn from a uniform distribution. UTIL means utilities computed using social preferences are used for attraction updating and PAYOFF means pure payoffs in the game are used for attraction updating in the learning model.

Model	Parameter Estimates				LL	AIC	BIC	Pseudo-R^2
SP-UTIL	μ_λ	1.4873	σ_λ	0.3132	-5025	-5029	-5035	0.2559
	μ_{N1}	0.0753	σ_{N1}	7.9524				
SP-PAYOFF	μ_λ	1.3959	σ_λ	0.3286	-5032	-5036	-5042	0.2549
	μ_{N1}	0.2624	σ_{N1}	5.5052				
RAND-UTIL	μ_λ	1.3958	σ_λ	0.3449	-5099	-5103	-5109	0.2450
	μ_{N1}	0.0148	σ_{N1}	9.0662				
RAND-PAYOFF	μ_λ	1.3850	σ_λ	0.3099	-5103	-5107	-5113	0.2444
	μ_{N1}	0.0966	σ_{N1}	5.9423				
RAND‡					-6759	-6759	-6759	0.0000

‡ RAND is the random choice model. According to this model any choice has a probability of $\frac{1}{21}$ of being chosen.

fit of SP-PAYOFF and SP-UTIL indicates that once initial round contributions are determined by the social preferences, there is no additional information to be gained about the path of the contributions by using social preferences based utilities to update the attractions of strategies rather than using simple payoffs within the a repeated game. In other words, social preferences matter insofar they determine the first round contributions and then individuals behave solely based on payoff-based learning.

My analysis shows that choices in the repeated games are best described by social preferences affecting the choice of first round contributions and then subsequent contributions based on payoff-based learning. Free riders contribute on average less than other cooperative types in the repeated games, but their contributions are significantly larger than zero. An interesting candidate explanation for this observation is that free riders are behaving strategically (Muller et al., 2008; Ambrus & Pathak, 2011). In an experiment where the type information of the group one belongs to is common knowledge, Ambrus and Pathak (2011)

Table 2.10: Disentangling the role of social preferences and learning. REL is used to model learning. SP means social preferences determine first round choices and RAND means first round choices are drawn from a uniform distribution. UTIL means utilities computed using social preferences are used for attraction updating and PAYOFF means pure payoffs in the game are used for attraction updating in the learning model. For each pair of learning models, a two sided p-value favoring the learning model in the row is reported. The two-sided p-values are computed using the Vuong Test for Non-Nested Models.

	SP-PAYOFF	RAND-UTIL	RAND-PAYOFF	RAND
SP-UTIL	0.4625	0.0003	0.0003	0.0000
SP-PAYOFF		0.0011	0.0000	0.0000
RAND-UTIL			0.7467	0.0000
RAND-PAYOFF				0.0000

find that free riders reduce their contributions significantly earlier than that of others, thus providing evidence for strategic motives. I test this proposition using the data from the repeated games. I compute the number of the round in which the drop in contributions is the largest and test if the round number is significantly smaller for free riders compared to others. I do not find a significant difference in either R1 or R2 games (R1: Mann-Whitney test, $p = 0.299$; R2: Mann-Whitney test, $p = 0.951$) or when I consider games with finite and indefinite repetition separately (finitely repeated: Mann-Whitney test, $p = 0.487$; indefinitely repeated: Mann-Whitney test, $p = 0.796$). Thus, I do not find evidence that free riders contribute larger than zero in the first round because they are strategic. In fact, I do not find difference across all types in either R1 or R2 games (R1: Kruskal-Wallis test, $p = 0.390$; R2: Kruskal-Wallis test, $p = 0.962$) or when I consider games with finite and indefinite repetition separately (finitely repeated: Kruskal-Wallis test, $p = 0.329$; indefinitely repeated: Kruskal-Wallis test, $p = 0.792$). Therefore, behavior in repeated games is more in line with learning based on payoffs rather than strategic motives. It is also equally important to note that while full contributors contribute significantly higher amounts in the first round of a repeated game compared to other types, their contributions are significantly

Table 2.11: Effects of the type of repetition on the parameters of REL learning model

	λ	$N1$
Constant	1.46***	0.090
	(0.12)	(0.067)
Infinitely Repeated	0.23	-0.0054
	(0.16)	(0.12)
σ	0.31***	6.67***
	(0.10)	(0.47)
N	150	
LL	-5022	

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Coefficients are transformed back to the original scale. A treatment effect is the (partial) effect of moving from the baseline treatment to the corresponding treatment on the median parameter value. Standard errors are calculated using the *sandwich estimator* and treating all of each subjects choices as a single *super observation*. The entries for σ are the standard deviations of the untransformed normal distributions of the random coefficients.

smaller on average than the full endowment. The underlying reason for these patterns of first round contributions of the free riders and full contributors could be potentially their resistance to use extreme strategies (Kümmerli, Burton-Chellew, Ross-Gillespie, & West, 2010).

2.6.4 Effect of Finite and Indefinite Repetition on Learning

My analysis thus far has shown that choices in the repeated games are best described by social preferences affecting the choice of first round contributions and then subsequent contributions based on payoff-based reinforcement learning. In this subsection, I investigate if the type of repetition has any effect on the learning parameters of REL model.

Table 2.11 reports estimation results using only repeated game data from the R1 and R2-tasks. It contains the estimates for the two parameters of REL model, λ and $N1$ for the model that contains an intercept and a dummy for the finitely repeated game. All values are on the original scale. There are no significant effects of finite repetition on the estimated parameters of the REL learning model. Thus learning is identical in both finitely and indefinitely repeated games. This result explains the similar observed contribution patterns across finitely and indefinitely repeated games.

2.6.5 Predictive Power of Behavioral Specification: Subsidizing Cooperation in the First Round Can Sustain Higher Levels of Cooperation

In Section 2.6.3, I obtained a behavioral specification for individuals that can successfully explain the observed regularities in repeated public goods games. In this subsection, I investigate if this behavioral has predictive power. In other words, I test if the quantitative predictions made using this behavioral specification can be reproduced in lab experimentally. In Section 2.6.3, I showed that first round contributions are in line with that of social preferences and in later rounds decision making is in line with reinforcement learning. This implies that by changing price of cooperation in the first around alone relative to the price of cooperation in later rounds one can manipulate average contributions in the later rounds of a repeated game. By reducing the cost of cooperation in the first round alone, one can make individual contributions higher in the later rounds. This is because, at low prices the proportion of free riders decreases and full cooperators increases. Given my finding that individuals learn using a high inertia learning rule like that of reinforcement, the higher level of contributions in the first round would mean that higher levels of contributions will be sustained in later rounds too. Analogously, by increasing the cost of cooperation in the first round alone, one can make individual contributions lower in the later rounds.

I examine these predictive implications of the behavioral specification using a dialogue between agent-based models and lab experiments. I construct an empirical agent-based model with the behavioral specification obtained in previous subsections to make quantitative predictions about the mean contribution levels when cooperation is subsidized in the first round and when it is made costly in the first round relative to the later rounds. Then, I conduct behavioral experiments with treatments that are identical to the treatments in agent-based simulations. The data from lab experiments is used to test the quantitative

Table 2.12: Means and standard deviations of the underlying normals of the distributions of social preference and learning parameters

	β	γ	λ	$N1$
μ	2.3257	2.4382	0.3336	-1.3380
σ	3.0894	1.7865	0.3286	5.5052

predictions obtained from the agent-based models. If lab experiments can reproduce the results from the agent-based model, it would confirm that the behavioral specification obtained in this chapter not only explains the observed regularities in repeated public goods games but also carries predictive value so that quantitative predictions made using it can be reproduced with human subject experiments.

The empirical agent-based model is formulated after the estimation results in Section 2.6.3. I consider an agent population of size 76. For each agent, the social preference parameters and the learning parameters are drawn from the corresponding estimated distributions. The minimal model estimated in Table 2.3 is used to specify social preferences of individuals. The learning parameters are specified using the estimates for SP-PAYOFF model in Table 2.9. Note that social preference parameters and learning parameters were specified and estimated as log-normal distributions. Table 2.12 presents the means and standard deviations of the underlying normals of the log-normal distributions of the parameters. For each agent, each parameter value is assigned using a random draw taken from the corresponding log-normal distribution. In the simulations of repeated games, in the first round each agent makes a contribution decision based on its social preferences and a random belief. In the later rounds, agents use REL learning model with their corresponding agent-specific parameters to make contribution decisions.

The simulations are carried out for two different treatments. In both treatments, initial endowment is 20 tokens. The agent population is split into groups of size 4 in both treatments and the repeated public goods game lasts for 10 rounds. In the first treatment,

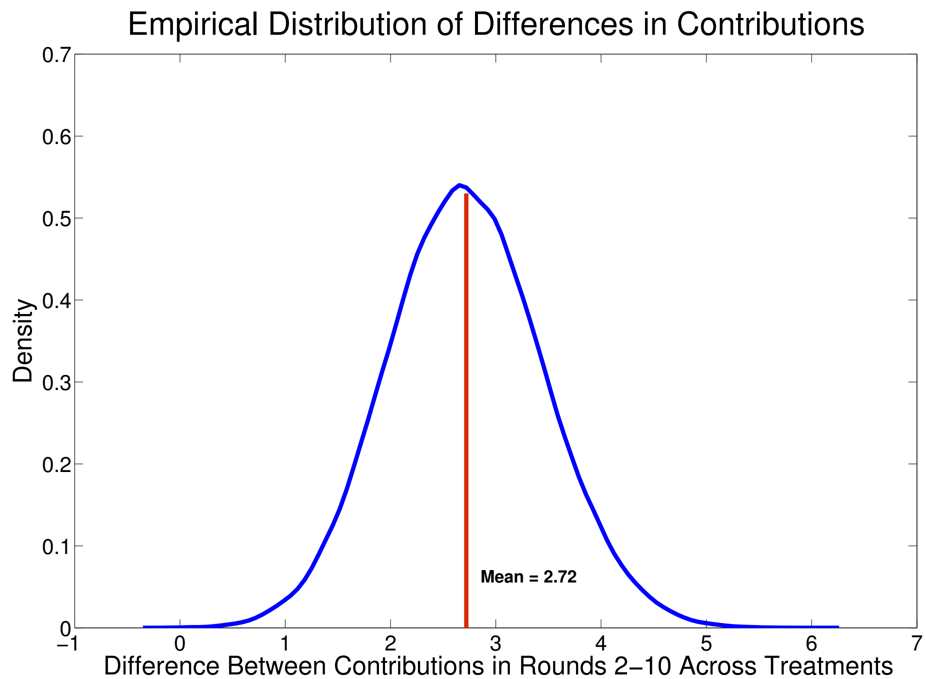

Figure 2.5: The empirical distribution of difference in average contributions in rounds 2-10 across high first round MPCR treatment and low first round MPCR treatment from agent-based simulations. The empirical distribution is obtained from 100,000 agent-based simulations of two treatments.

I fix the first round MPCR at 0.26 and the MPCR in later rounds is 0.5. In the second treatment, I fix the first round MPCR at 0.99 and in the later rounds the MPCR is 0.5. I then simulate a repeated public goods game for 10 rounds. The composition of each group remains fixed for all 10 rounds (the Partners' matching). In each simulation, I run the agent-based model for each treatment and record the difference between average contributions in rounds 2-10 across the treatments. In total, I conducted 100,000 simulations. The empirical distribution of difference between high first round MPCR treatment and low first round MPCR treatment is shown in Figure 2.5. Mean difference in contributions in rounds 2-10 across two treatments is 2.72 (s.d 0.74) and is significantly different from zero (Empirical p-value < 0.0001). As hypothesized from the behavioral specification, average contributions in later rounds significantly depend on the MPCR in the first round relative

Table 2.13: Validation experiment treatments

PG Game Parameters: Group Size =4, No. of Rounds = 10, $MPCR_{2-10} = 0.5$, Partners' Matching	
Treatment 1	Treatment 2
First Round MPCR = 0.99 4 sessions (20, 20, 20,16) = 76 subjects	First Round MPCR = 0.26 5 sessions (20,16, 16, 8,12) = 72 subjects

to the MPCR in the later rounds. Average contribution in rounds 2-10 when the first round MPCR is 0.99 is 10.31 (s.d 0.56). Average contribution in rounds 2-10 when the first round MPCR is 0.26 is 7.59 (s.d 0.51).

I followed up the results from the empirical agent-based model using a set of behavioral experiments. I conducted experiments involving two treatments which are identical to the simulation treatments. Details of the treatments and conducted sessions per each treatment are presented in Table 2.13. The experiments were run in the Interdisciplinary Center for Economic Science (ICES) at George Mason University between February and April of 2016. Nine sessions were run, and there were a total of 148 participants. Participants were randomly assigned to cubicles and made their decisions privately and anonymously. Five of the nine sessions involved low first round MPCR of 0.26. There were in total 72 participants in these sessions. Four of the nine sessions involved high first round MPCR of 0.99. There were in total 76 participants in these sessions. Both treatments involved an MPCR of 0.5 in rounds 2-10 of the repeated game and employed Partners' matching with groups of size 4. Each experimental session lasted for approximately 1.5 hours. Average participant earnings were $22.27. Experimental instructions for the high first round MPCR treatment are provided in Appendix 2.8.9 [19]. These include screen shots of the decision screens participants used to make decisions during the experiment.

In the high first round MPCR treatment, I find that average contribution in rounds 2-10 is 7.93 tokens. In the low first round MPCR treatment, I find that average contribution in

[19] The instructions for the low first round MPCR treatment are available online at http://www.chennacotla.org/research

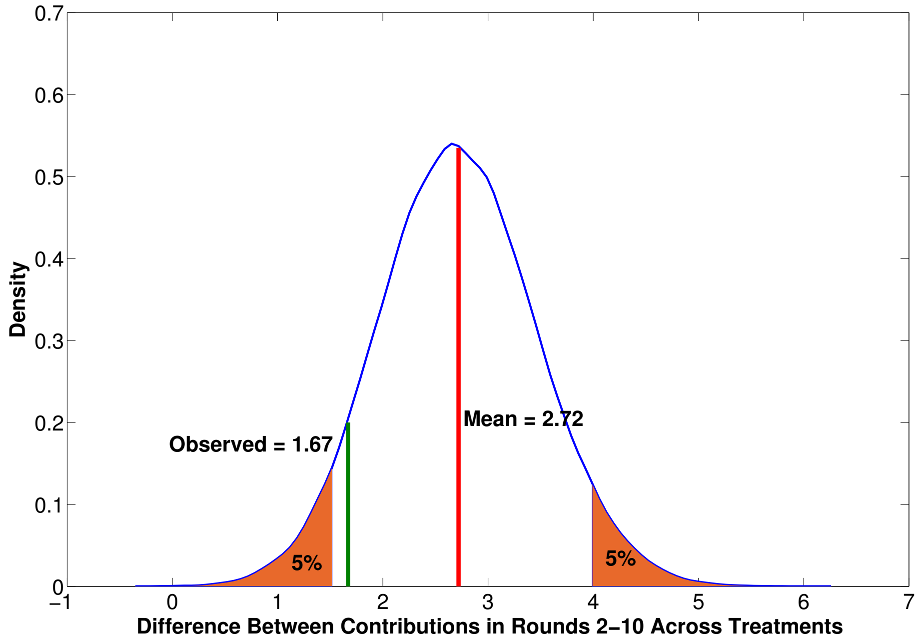

Figure 2.6: Observed difference in average contributions in rounds 2-10 across high first round MPCR treatment and low first round MPCR treatment of human subject experiments is shown in green. 10% critical region of the empirical distribution of difference in average contributions in rounds 2-10 from agent-based simulations is shown in orange.

rounds 2-10 is 6.25. The difference between two treatments is 1.67 tokens and significantly different from zero (Contributions in 2-10 rounds are regressed on high first round MPCR dummy. Standard errors are clustered at the subject level. P-value value on high first round MPCR is 0.069). These findings confirm that just by making cooperation less or more costly in the first round, one could significantly influence average contributions in later rounds. More importantly, the difference between average contributions in rounds 2-10 across two treatments is not significantly different from the average difference between contributions across two treatments from the empirical agent-based model (empirical two sided p-value = 0.152). Figure 2.6 shows 10% critical region in the empirical distribution of the difference between contributions in rounds 2-10 across two treatments obtained from the agent-based simulations. It can be seen that the observed difference between contributions in rounds

2-10 across two treatments in behavioral experiments cannot be rejected at 10% significance level as being different from the mean difference between contributions in rounds 2-10 across two treatments obtained from the empirical agent-based model. This validates that the behavioral specification obtained in this chapter not only explains the previously observed behavioral regularities but also carries predictive accuracy in making novel predictions that can be reproduced in the lab with human subjects.

2.7 Conclusions

This paper investigates if the framework of social preferences can organize varying distributions of conditional contribution decisions when prices change. I use laboratory experiments and structural estimation of preference parameters to investigate this. Estimated distribution of social preference parameters is reasonably successful in explaining the aggregate changes in the distribution of cooperative types (free riders, conditional cooperators, and full cooperators) when prices change. The model is able to predict an increasing proportion of full cooperators when the price of cooperation gets smaller. However, there is a significant effect of experience on the distribution of cooperative types. There are more free riders than predicted by the social preference model in the treatment where subjects had prior experience.

A number of previous studies have recognized that both learning and social preferences are important in explaining the contribution patterns in repeated public goods games (Cooper & Stockman, 2002; Janssen & Ahn, 2006; Wendel & Oppenheimer, 2010; Arifovic & Ledyard, 2012). These studies assume individuals start with random contributions and move towards the equilibrium contributions determined by their social preferences. This is inconsistent with choices we observe in our repeated game data. First round contributions are not random and do correlate with individual social preferences. A hybrid model in which first round contributions are determined by social preferences and after that participants use a payoff-based reinforcement learning to make contribution decisions best reconciles choices across the one-shot and repeated environments.

Voluntary cooperation is inherently fragile. Mechanisms like costly punishment have been shown to sustain cooperation (Fehr & Gächter, 2000; Sefton et al., 2007), however, such mechanisms in addition to being costly to implement can also lead to undesirable outcomes of antisocial behavior (Herrmann, Thöni, & Gächter, 2008). My findings show that individuals are identical in the manner in which they learn, no matter what the underlying social preferences are, but differ in terms of the contributions that they start with. High inertia learning like that of reinforcement that individuals use in public goods games implies that higher first round contributions would lead to higher sustained levels of cooperation over time. Higher first round contributions can be induced by reducing the cost of cooperation in the first round since individuals are observed to make their first round choices based on their social preferences. Thus, subsidizing cooperation in the first round alone could be sufficient to sustain higher average levels of cooperation in the later rounds. I test and confirm this novel finding using a dialogue between agent-based simulations and behavioral experiments. This result can have attractive policy implications since subsidizing cooperation in the first round is straightforward to implement and will not invoke undesirable outcomes of antisocial behavior like mechanisms similar to that of punishment.

2.8 Appendix

2.8.1 Experiment Instructions

.

Instructions

Welcome and thank you for participating in todays economic experiment. Please put away all your belongings and turn off your cell phones. You are not allowed to talk to any other participant during the experiment. If you have any questions, please raise your hand. We will come to you and answer your questions in private. The experiment will be run entirely on the computer and all interactions between yourself and others will take place via the computer terminal.

You have earned $5 just for showing up on time. This is yours to keep. In addition, depending upon the decisions you make, the decisions others make and random choice, you can earn more money. These instructions describe in detail the experiment and tasks you are asked to complete.

During the experiment, your earnings will be described in terms of tokens. At the end of the experiment, the total number of tokens you have earned will be converted to money at the following rate:

20 tokens = $1 (1 token = 5 cents)

The experiment consists of four tasks. You will receive instructions for each task prior to making decisions for that task. Your total earnings from the experiment will be the sum of your earnings in each task. At the end of the session, the total number of tokens earned across all four tasks will be converted to money and paid to you privately in cash, along with the $5 show-up fee.

Your Neighborhood

You will be placed in a network with 14 other participants as shown in the Figure below. The placement of participants in the network is random, and participants do not know who is connected to whom.

Each participant is placed at one position on the network, shown in the Figure below, and is connected to exactly two other participants. This placement and connection are fixed throughout each of the four tasks. You and the two other participants that are connected to you in the network define your neighborhood. In the Figure, for example, if you are placed in the position of the circle that is highlighted in pink then your neighbors are highlighted in yellow. In the network, there are three connected participants in each neighborhood and five neighborhoods.

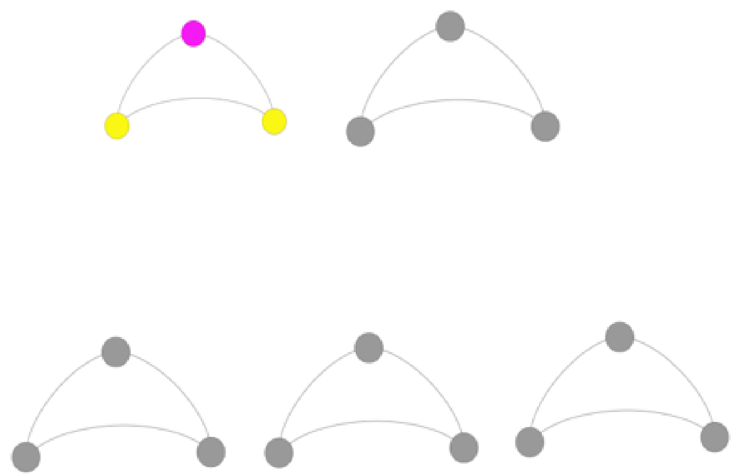

Figure 2.7: Groups in the experiment

The Decision Situation

Each participant is provided with 20 tokens and must decide how to allocate the tokens between a private account and a group project. You can choose to put none, all or some of your tokens into the group project. The tokens you choose not to contribute to the group project will remain in your private account. Everyone makes the same decision.

For each token put in the private account you earn exactly one token. You are the only one who earns tokens from your private account.

What you earn from the group project depends on the total number of tokens that you and your neighbors contribute to the group project. The more each member of the neighborhood contributes to the group project, the more each member earns. Remember that your neighborhood includes you and two other participants.

Your earnings from the group project are best explained by a number of examples.

Example 1: Suppose that you decided to contribute no tokens to the group project but the 2 other members of your neighborhood contribute a total of 36 tokens. Then your earnings from the group project would be 36 tokens $\times$ 0.4 = 14.4 tokens. Everyone else in your group would also earn 14.4 tokens.

Example 2: Suppose that you contribute 15 tokens to the group project and the 2 other members of your neighborhood invest a total of 36 tokens. This makes a group total of 51 tokens. Your earnings from the group project would be 51 tokens $\times$ 0.4 = 20.4 tokens. The other 2 members of the group would also earn 20.4 tokens.

Example 3: Suppose that you contribute 20 tokens in the group project but the other 2 members in your neighborhood invest nothing. Then you, and everyone else in the group, would earn from the group project 8 tokens (20 tokens $\times$ 0.4 = 8 tokens).

As you can see, every token contributed to the group project earns 0.4 tokens for every member of the neighborhood, not just the participant who puts it there. It does not matter who contributes tokens to the group project. Everyone will get a return from every token contributed therewhether they contributed tokens in the group project or not.

Your total earnings from the private account and group project will be:

Your total earnings = 20 - your tokens contributed to the group project + 0.4 $\times$ sum of tokens contributed to the group project by all members of your neighborhood

You will now complete some questions to make sure everyone understands how earnings are calculated.

Questions

Subject number: _______________

Please answer the following questions. These will help you understand how earnings are calculated. Your payoff is not affected by your answers to these questions.

Each token in the private account earns 1 token. Each token in the group project earns 0.4 tokens for each participant in the neighborhood.

1. Each participant has 20 tokens. Suppose you contribute 12 tokens to the group project and the other two participants contribute 18 tokens in total.

What are your earnings from your private account? _______________

What are your earnings from the group project? _______________

What are your neighbors earnings from the group project? _______________

What are your total earnings? _______________

2. Each participant has 20 tokens. Suppose you contribute 20 tokens to the group project and the other two participants contribute 38 tokens in total.

What are your earnings from your private account? _______________

What are your earnings from the group project? _______________

What are your neighbors earnings from the group project? _______________

What are your total earnings? _______________

3. Each participant has 20 tokens. Suppose you contribute 2 tokens to the group project and the other two participants contribute 38 tokens in total.

What are your earnings from your private account? _______________

What are your earnings from the group project? _______________

What are your neighbors earnings from the group project? _______________

What are your total earnings? _______________

Instructions for Task 1 – A-Task

In the A-task, you will make a choice for the decision situation described earlier. You will have 20 tokens and must decide how many to put into your private account and the group project. You will be randomly assigned to a 3-participant neighborhood.

Each participant has two decisions in this task: make an unconditional contribution and complete a contribution table. Details about these two decisions are as follows.

Unconditional Contribution: In this decision, you must decide how many of the 20 tokens you would like to put in the group project. You will make your decision on a screen such as the following.

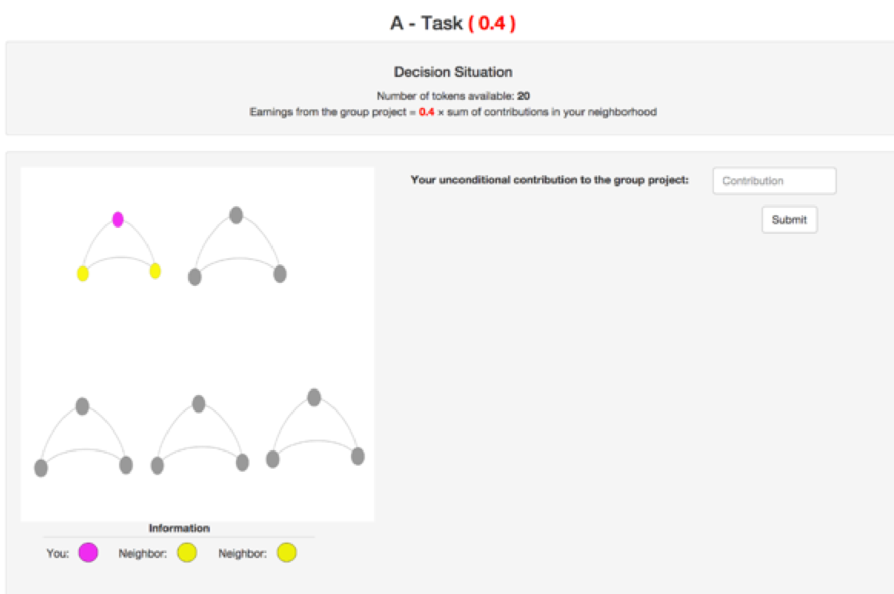

Figure 2.8: Decision screen for unconditional contribution choice in strategy games

Contribution Table: In this decision, you must decide how many tokens you would like contribute to the group project for each possible average contribution of your neighbors (e.g. 0, 1, 2,..., 20). That is, if your neighbors contributed 0 tokens on average, how much would you contribute? If they contributed 1 token on average, how much would you contribute? If they contributed 2 tokens on average, how much would you contribute? And so on, up to 20 tokens on average.

This means that in total you have to give 21 responses. You will make your decisions on a screen such as the following

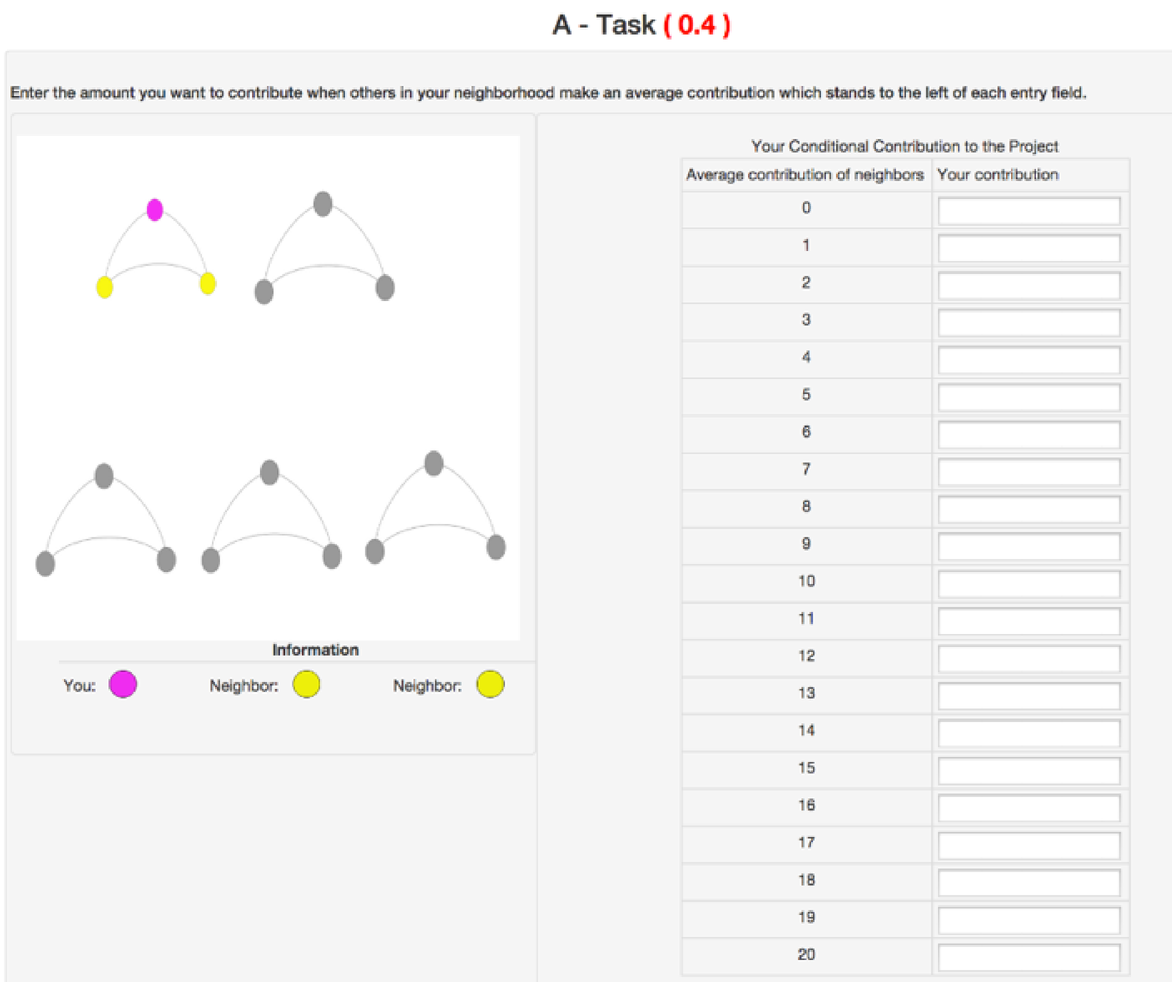

Figure 2.9: Decision screen for conditional contribution table in strategy games

Once each participant has made the unconditional decision and completed the contribution table, the computer will randomly determine if the unconditional contribution or the contribution table will be used to determine your earnings. In each neighborhood, one of the three participants is randomly chosen to have the contribution table count to calculate earnings. For the other two participants in the neighborhood, the unconditional contribution counts to calculate earnings. How is this done? If the participant is chosen to have his contribution table count for earnings, first the unconditional contributions to the group project of his neighbors are averaged and rounded to the nearest integer (e.g. 0,1,2,...,20). Then, the contribution table of the participant is used to determine how many tokens the participant contributes to the group project. The number of tokens contributed is the amount he specified for the average contribution of his neighbors.

So, if the average contribution of his neighbors is 16 tokens, and he specified 10 tokens if the average contribution of his neighbors is 16, the total contributed to the group project by everyone in the neighborhood would be 42 ($16 \times 2 + 10$) tokens.

You will not know if advance which decision, the unconditional contribution or the contribution table, will count to determine your earnings, so you should make each decision as though it will count for your earnings.

The following examples should help make this procedure clear.

Example 1: Suppose the contribution table was randomly chosen to count for you. This means that the decisions you made in the contribution table determine your earnings. For the other two neighbors their unconditional contributions determine their earnings. Suppose that the total contributions of the other two neighbors are 26 tokens, and the average contribution 13 tokens (26 tokens/2). In your contribution table, suppose you chose to contribute 4 tokens if the average contribution of your neighbors is 13, then your earnings for Task 1 would be: $20 - 4 + 0.4 \times (4 + 26) = 28$. If instead you chose to contribute 14 if the average contribution of neighbors is 13, your payoff would be: $20 - 14 + 0.4 \times (14 + 26) = 22$.

Example 2: Suppose the unconditional contribution was randomly chosen to count for your earnings. Also, suppose that the unconditional contribution of the neighbor who was not selected for the contribution table to count is 12. If your unconditional contribution is 20, then the average unconditional contribution is 16 tokens ($(20 + 12)/2$). If the neighbor selected to have his contribution table count chose 18 tokens if the average contribution of his neighbors is 16, then your earnings are: $20 - 20 + 0.4 \times (20 + 12 + 18) = 20$.

Are there any questions before we begin?

Instructions for Task 2 – B-Task

In the B-Task, you will be randomly assigned to a 3-participant neighborhood as described earlier. Your neighbors in Task 2 may be different from your neighbors in Task 1, however, you will remain with the same neighbors for all decisions you make in Task 2.

The B-Task lasts for several rounds. The number of rounds is randomly determined. In each round, you face the basic decision situation described at the beginning of the experiment. After each round, there is an 85% probability that there will be one more round. So, for instance, if you are in round 2, the probability there will be a third round is 85% and if you are in round 9, the probability there will be another round is also 85%. How this works is as follows. After each round, the computer will randomly draw a number between 1 and 100 (e.g. 1, 2, 3,..., 100), where each number is equally likely to be chosen. If the chosen number is 85 or lower, there will be another round. If the chosen number is 86 or above, there will be no additional rounds, and the task will end. You will know there is another round if you see the decision screen again and are asked to make a decision. If the task ends, you will get a message saying the task is done. You will not know ahead of time for how many rounds you will make decisions.

In each round, you will be given 20 tokens and must decide how many tokens you would like to contribute to the group project and how many you would like to put in your private account. You will receive earnings only from the group project that involves participants in your neighborhood. Your earnings from your contribution decision in a given round are determined as:

Your total earnings in a round = 20 - your tokens contributed to the group project + 0.4 $\times$ sum of tokens contributed to the group project by all members of your neighborhood

You will participate in the decision situation repeatedly with the same neighbors, until it is randomly determined that there are no more rounds.

You will make decisions on a screen such as the following:

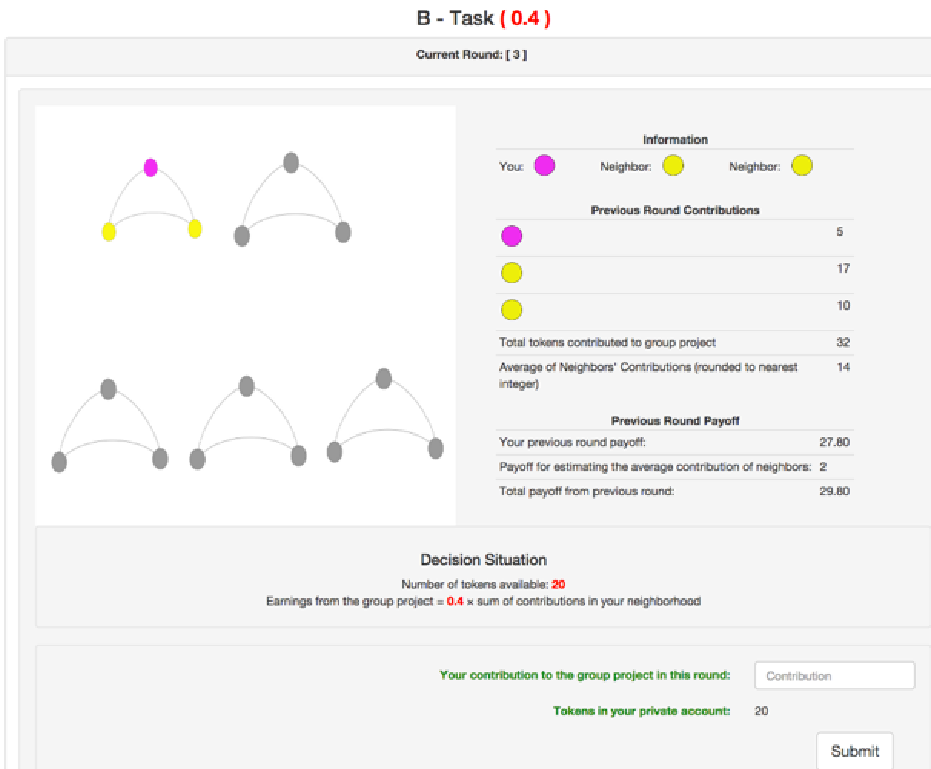

Figure 2.10: Decision Screen for contribution choice in repeated game

Here is an example to explain how earnings are calculated in each round:

Example 1: Suppose you chose to contribute 12 tokens and your neighbors chose to contribute 30 tokens in total. Your earnings in that round would be: 20 - 12 + 0.4 × (12 + 30) = 24.8 tokens.

In each round, after you decide how much to contribute to the group project, you will be asked to guess the average contribution to the project (rounded to the nearest integer) of your two neighbors. You will receive tokens for the accuracy of your estimate. If your guess is exactly equal to the average contribution of your neighbors you will receive 3 tokens in addition to your earnings for that round. If your guess was off by 1 token, you will get 2 additional tokens. If your guess was off by 2 tokens, you will get 1 additional token. And, if your guess was off by 3 or more tokens, you will get 0 additional tokens.

When everybody in your neighborhood has completed the two decisions, you will be shown each of their contributions, the total contributions to the group project, and the average contribution. You will only be informed of the contributions of those in your neighborhood. You will not be informed of contributions of participants in other neighborhoods. You will also be informed of your earnings for the current round.

Once all subjects in the experiment have completed the two decisions and are told their earnings and the contributions of their neighbors in the current round, the computer will randomly draw a number between 1 and 100 to see if everyone plays another round. If there is not another round, the task is done.

Are there any questions before we begin?

Instructions for Task 3 – A-Task

In Task 3, you will make two decisions again as you did in the Task 1 A-Task. The difference between this task and Task 1 is that for each token contributed to the group project you, and the other two neighbors, will get 0.8 tokens back.

Things to remember:

1. You will be randomly placed on the network at the beginning of the task and assigned to a 3-participant neighborhood. Your neighbors in this Task 3 may be different than in the previous two tasks.
2. You will make two decisions.
3. The first decision, the unconditional contribution, is how many of your 20 tokens you want to contribute to the group project.
4. The second decision, completing the contribution table, is how many of your 20 tokens you want to contribute to the group project for each possible average contribution of your neighbors (e.g. 0, 1, 2,..., 20).
5. Each token contributed to the group project will earn 0.8 tokens for each participant in the neighborhood.
6. One of the two decisions, the unconditional contribution or the contribution table, will be randomly chosen to determine earnings. You will not know ahead of time which decision will count.

Are there any questions before we begin?

Instructions for Task 4 – B-Task

In Task 4, you will make decisions again as you did in the Task 2 B-Task. The difference between this task and Task 2 is that for each token contributed to the group project you, and the other two neighbors, will get 0.8 tokens back.

Things to remember:

1. You will be randomly placed on the network at the beginning of the task and assigned to a 3-participant neighborhood. Your neighbors in Task 4 may be different than in the previous tasks, however, you will remain with the same neighbors for all rounds in this task.

2. You will face the same decision situation for several rounds. You must decide how many of your 20 tokens you want to contribute to the group project.

3. The number of rounds is randomly determined. After each round, there is an 85

4. If there is another round, you will see a decision screen to make another choice. If there is not another round, you will be a message saying the task is over.

Are there any questions before we begin?

2.8.2 Group Level Contributions and Beliefs

In experiments with Strangers' matching, which randomly shuffle participants across groups in each round, the unit of interaction is all participants in a session in a repeated game. In contrast, the unit of interaction in our repeated games is a 3-person group since our design uses fixed groups in repeated games (the Partners' matching). Therefore, contribution and beliefs patterns at a group level rather than the session level are more informative in understanding the contribution dynamics in the repeated games. Figures 2.11, 2.12, 2.13 and 2.14 provide plots of average contributions and beliefs across groups in the experiments. We have in total 50 distinct groups in R1 and R2 each. 25 of the groups involve high MCPR and 25 of the groups involve low MPCR in each case. We have split the groups according to the level of MPCR in the figures. Each group level plot contains information about group number, whether it is from R1 or R2, the type of repetition (I - Indefinitely repeated, F-finitely repeated), MPCR (0.4 or 0.8) and composition of the group in format of (no. of FR - no. of CC - no. of FC - no. of NC) in its title. The type information is computed using the classification algorithm using LCP (Kurzban & Houser, 2005). The x-axis label of each plot contains information about the individual level first round contribution and corresponding types. Types are coded as FR-1, CC-2, FC-3, NC-4.

2.8.3 Arifovic-Ledyard Model of Social Preferences and Type Switches

Monetary payoff from contributing c

$$\pi^i = e^i - c^i + \alpha \sum_{j=1}^{N_g} c^j$$

Arifovic-Ledyard (2012) specification: Utility derived for agent i from contributing

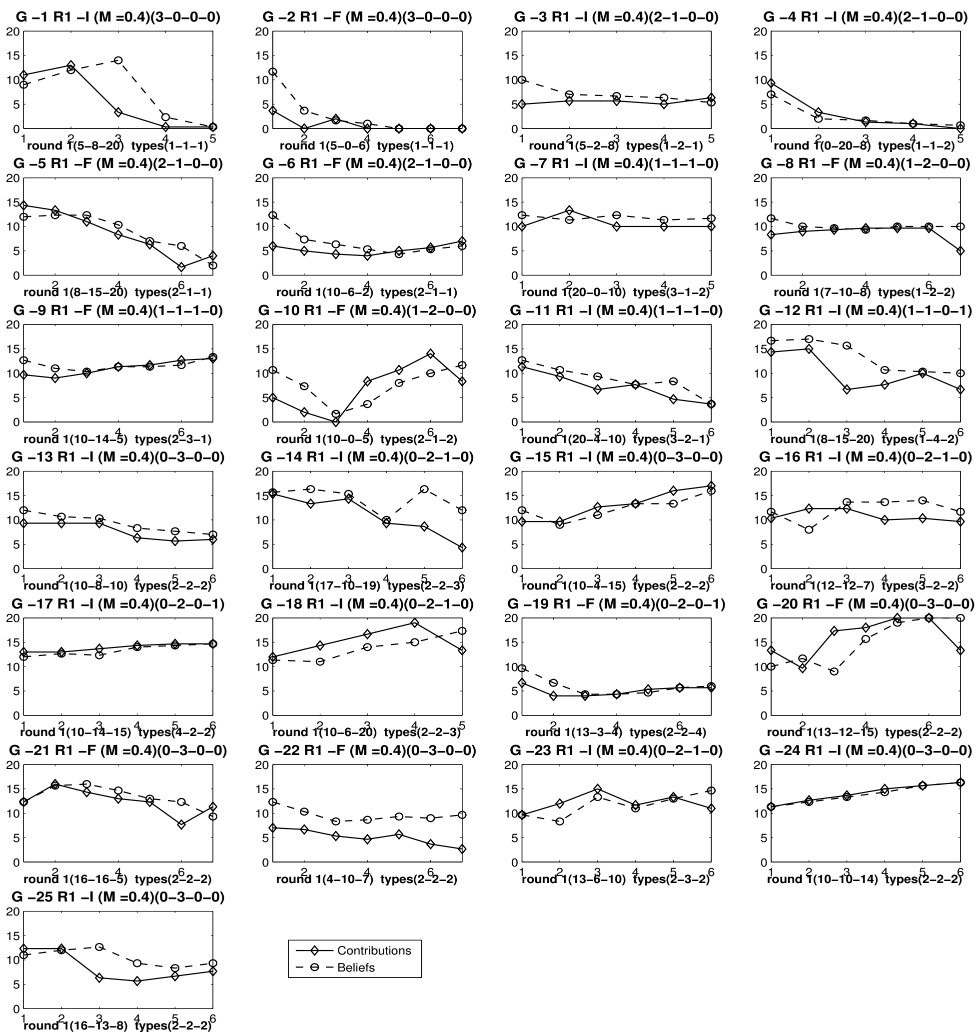

Figure 2.11: Group level contributions and beliefs in R1 - Low MPCR (0.4)

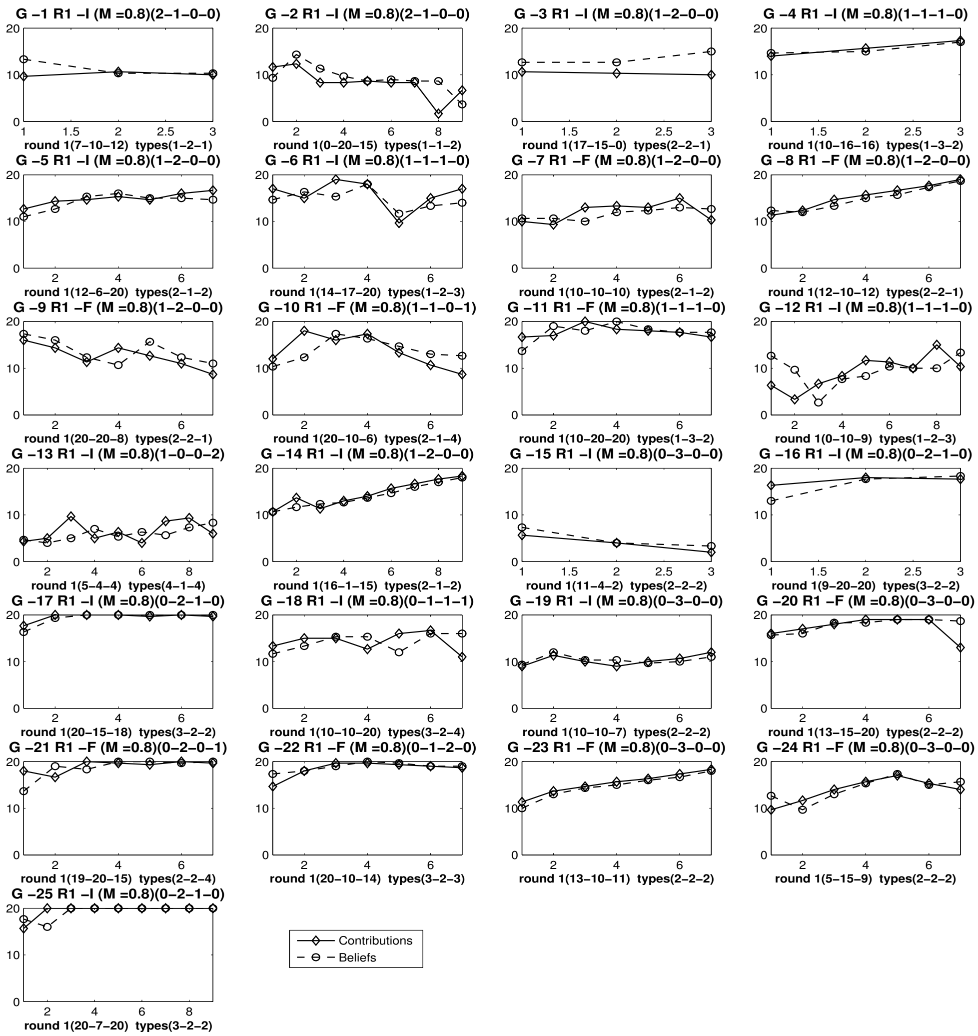

Figure 2.12: Group level contributions and beliefs in R1 - High MPCR (0.8)

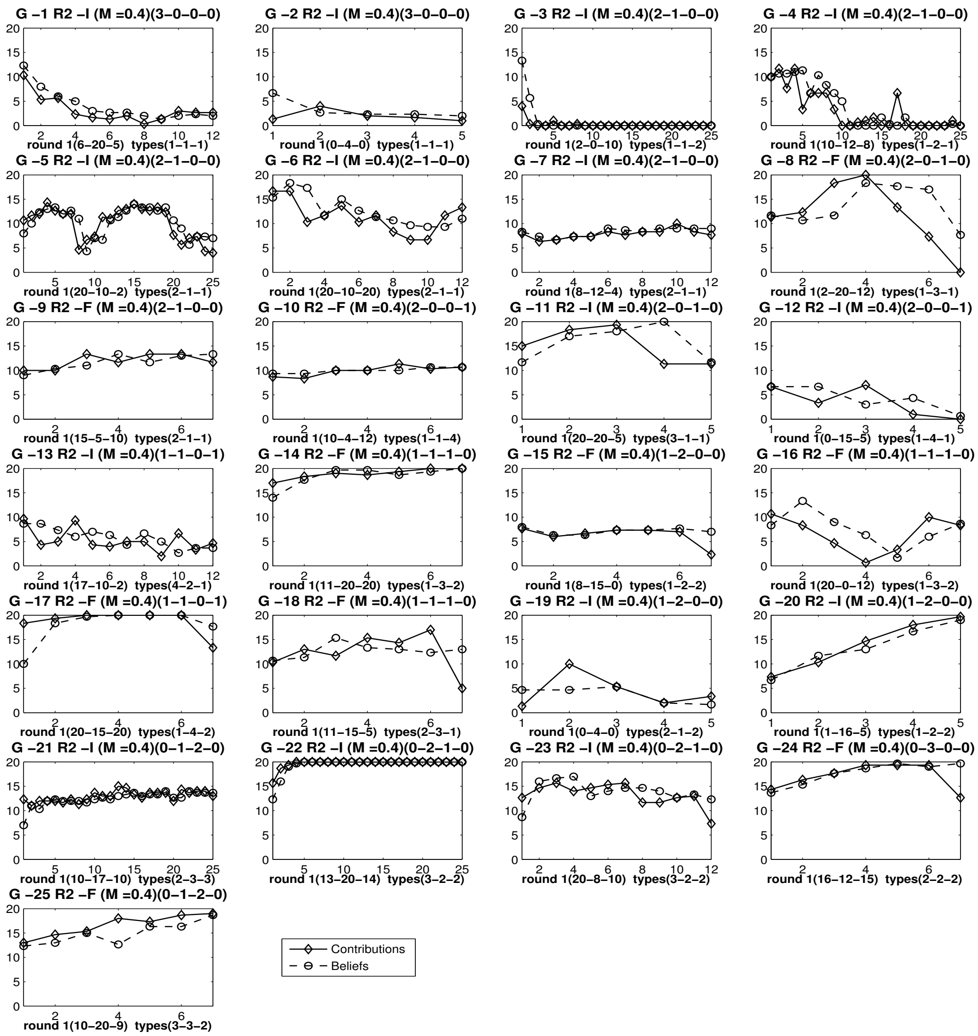

Figure 2.13: Group level contributions and beliefs in R2 - Low MPCR (0.4)

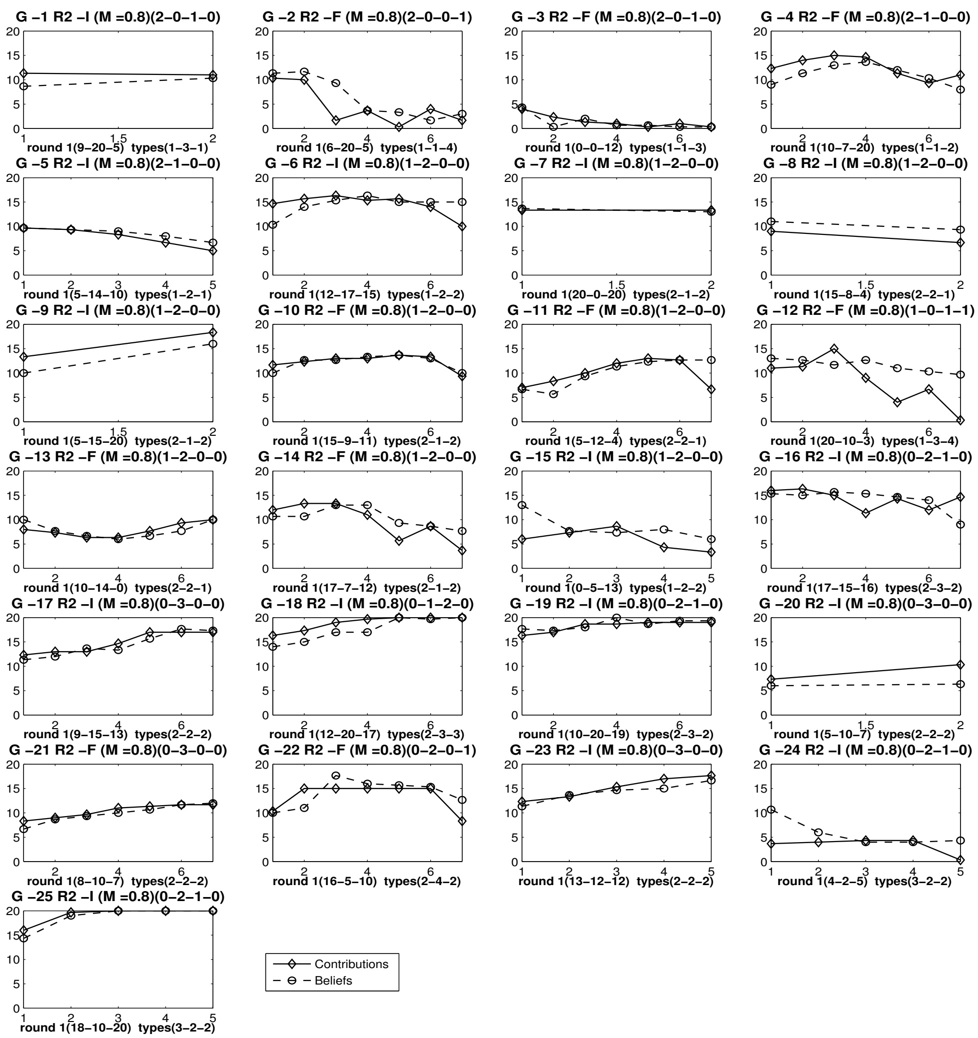

Figure 2.14: Group level contributions and beliefs in R2 - High MPCR (0.8)

amount c when others contribute o on average:

$$u^i(c,o) = \pi^i(c,o) + \beta^i \bar{\pi}(c,o) - \gamma^i max\{0, \bar{\pi}(c,o) - \pi^i(c,o)\}$$

Where $\beta^i \geq 0; \gamma^i \geq 0; \bar{\pi} = \frac{\sum \pi^i}{N}$

From equation (2) in Arifovic & Ledyard (2012), three equilibria are possible for a given combination of $\beta \geq 0$ and $\gamma \geq 0$ of an individual: Free Riding(FR), Conditional Cooperation (CC), and Full Cooperation (FC). We derive conditions for these three possibilities below for each MPCR levels in our experiments:

- When MPCR = 0.4

  FR: $0 \geq 0.07\beta - 0.6$

  CC: $0.67\gamma \geq 0.07\beta - 0.6 \geq 0$

  FC: $0.67\gamma \leq 0.07\beta - 0.6$

- When MPCR = 0.8

  FR: $0 \geq 0.47\beta - 0.2$

  CC: $0.67\gamma \geq 0.47\beta - 0.2 \geq 0$

  FC: $0.67\gamma \leq 0.47\beta - 0.2$

  Our data is unique in the sense, we have observations for a given individual under MPCR = 0.4 and MPCR = 0.8. This makes it possible to see if type switches are in congruence with what theory allows. For example, according to the theory some of the transitions among types are not impossible (for example, a person cannot switch from being a Free Rider to a Conditional Cooperator when MPCR changes from 0.8 to 0.4).

- Possible Switches
  - switching from FR - CC when MPCR changes from 0.4 and 0.8: $\beta \in [0, 8.57]$ and $\beta \in [0.42, 1.40\gamma + 0.42]$
  - switching from FR - CC when MPCR changes from 0.8 and 0.4: $\beta \in [0, 0.42]$ and $\beta \in [8.57, 9.43\gamma + 8.57]$ (Not Possible)
  - switching from FR - FC when MPCR changes from 0.4 and 0.8: $\beta \in [0, 8.57]$ and $\beta \in [1.40\gamma + 0.42, \infty]$
  - switching from FR - FC when MPCR changes from 0.8 and 0.4: $\beta \in [0, 0.42]$ and $\beta \in [9.43\gamma + 8.57, \infty]$ (Not Possible)
  - switching from CC - FC when MPCR changes from 0.4 and 0.8: $\beta \in [8.57, 9.43\gamma + 8.57]$and $\beta \in [1.40\gamma + 0.42, \infty]$
  - switching from CC - FC when MPCR changes from 0.8 and 0.4: $\beta \in [0.42, 0.42 + 1.40\gamma]$ and $\beta \in [9.43\gamma + 8.57, \infty]$ (Not Possible)
- No Switch Conditions
  - FR - FR in both cases $\beta \leq 0.43$ & $\beta \leq 8.57$
  - CC-CC $\beta \in [8.57, 9.43\gamma + 8.57]$ & $\beta \in [0.42, 0.42 + 1.40\gamma]$
  - FC-FC $\beta \in [9.43\gamma + 8.57, \infty]$& $\beta \in [1.40\gamma + 0.42, \infty]$

2.8.4 Separate Estimation of Strategic and Repeated Experiments

Table 2.14 presents the estimates of the social preference parameters and the random choice probability parameter for strategy and repeated games separately.

2.8.5 Econometric Model - Empirical Identification

To investigate empirical identification of the econometric model we used in Results section, we set up a Monte Carlo experiment. Following the procedure outlined in (Cherchi & de

Table 2.14: Separate estimation of β, γ and ω

	Strategy Game Data	Repeated Game Data
Median (β)	10.23***	17.92***
	(0.81)	(0.15)
σ_β	3.09***	2.01***
	(0.029)	(0.0037)
Median (γ)	11.45***	13.28***
	(0.91)	(0.083)
σ_γ	1.79***	1.76***
	(0.016)	(0.0041)
Median (ω)	0.12***	0.36***
	(0.017)	(0.036)
σ_ω	3.34***	1.67***
	(0.24)	(0.19)
N	150	150
LL	-13625	-4964

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Coefficients are transformed back to the original scale. A treatment effect is the (partial) effect of moving from the baseline treatment to the corresponding treatment on the median parameter value. Standard errors are calculated using the *sandwich estimator* and treating all of each subjects choices as a single *super observation*. The entries for σ are the standard deviations of the untransformed normal distributions of the random coefficients.

Dios Ortúzar, 2008), we simulated a collection of data sets following the decision making process outlined in Section 2.5. In this exercise, we generated artificial datasets involving 150 individuals making decisions across P1 and P2 games as described in Section 2.2. We have used the parameter estimates obtained using the minimal model estimation without any co-variates using only strategy game data as reported in Table 2.14 to generate artificial data sets. Monte carlo procedure is as follows:

- **For each data set** i **in** [1..50]

  1. **for each individual** j **in** [1...150]

     - Draw β_j, γ_j and ω_j from distributions specified by original parameters
     - Generate series of choice situation corresponding P1 and P2 experiments with MPCR levels 0.4, 0.8 respectively

Table 2.15: Empirical identification results for log-normal specification of parameters

Parameter	True	Recovered†	Bias	T-stat	Cohen's d	MSE
μ_β	2.22 (9.21)	2.23 (9.80)	0.0073	0.15	0.022	0.11
σ_β	2.68	2.96	0.28	3.31***	0.47	0.43
μ_γ	2.17 (8.81)	2.19 (9.30)	0.018	0.48	0.064	0.072
σ_γ	1.59	2.07	0.48	7.04***	0.99	0.46
μ_ω	-0.73 (0.33)	-0.92 (0.29)	-0.19	-6.94***	-0.98	0.072
σ_ω	1.61	1.64	0.024	1.15	0.016	0.023

Notes:
1. $^{*}p < 0.1, ^{**} p < 0.05, ^{***} p < 0.01$
2. In brackets transformed median value

- In each choice situation
  * compute utilities of each contribution level using β_j, γ_j
  * Add EV1 error to the utility of each contribution and call it the observed utility
  * Draw a random number r from Uniform[0,1]. if $r < (1 - \omega_j)$ choose the contribution level with highest observed utility . Otherwise, choose each contribution level with equal probability.

2. Run estimation as outlined in Section 2.5 and collect the estimated parameters.

Table summarizes results from the Monte Carlo experiment outlined above. It can be seen that original parameters used in generating data are recovered very closely. However in two instances, $\sigma_\beta, \sigma_\gamma$ and μ_ω, the recovered parameters seem to be statistically significantly different from true parameter values. While the t-tests show that we can reject the true parameter being the mean of the distribution represented by recovered parameters, we can observe that MSE ($bias^2 + variance$) is quite small in these cases. For $\sigma_\beta, \sigma_\gamma$, the bias could be arising due to unbounded nature of log-normal distribution.

Table 2.16: Estimation of β, γ and ω using strategy experiments: Logistic-normal specification

Parameter	Estimate
Median (β)	13.34***
	(0.15)
σ_β	5.28***
	(0.09)
Median (γ)	43.83***
	(0.27)
σ_γ	4.27***
	(0.021)
Median (ω)	0.30***
	(0.061)
σ_ω	1.55***
	(0.15)
N	150
LL	-13601

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Coefficients are transformed back to the original scale. A treatment effect is the (partial) effect of moving from the baseline treatment to the corresponding treatment on the median parameter value. Standard errors are calculated using the *sandwich estimator* and treating all of each subjects choices as a single *super observation*. The entries for σ are the standard deviations of the untransformed normal distributions of the random coefficients.

2.8.6 Alternative Specifications of Distributions of Social Preference Parameters and Identification

As alternative to the log-normal specification of β, γ, in this appendix we explored a logistic-normal specification and its empirical identifiability. Both specifications are theoretically identified. Table 2.16 presents estimation results from logistic-normal specification using data from strategy experiments. We use these estimates to do an empirical identification exercise as described in Appendix 2.8.5. The results from the empirical identification exercise are presented in Table 2.17. We conclude from the large bias in identified parameters in Table 2.17 that this specification has significant problems with empirical identification given our design and sample size.

Table 2.17: Empirical identification of logistic-normal specification

Parameter	True	Recovered†	Bias	T-stat	Cohen's d	MSE
μ_β	-1.87 (13.34)	-1.23 (22.99)	0.63	13.35***	1.88	0.51
σ_β	5.29	6.84	1.56	8.032***	1.14	4.34
μ_γ	-0.25 (43.83)	-0.039 (48.92)	0.21	3.36***	0.48	0.23
σ_γ	4.28	4.38	0.10	0.51	0.072	1.99
μ_ω	-0.83 (0.30)	-0.98 (0.27)	-0.15	-7.14***	-1.01	0.043
σ_ω	1.55	1.56	0.018	0.94	0.13	0.018

Notes:
1. $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$
2. In brackets transformed median value

Table 2.18: Types in P1 and P2 across different levels of MPCR computed using the LCP method

Type	P1-Low	P1-High	P2-Low	P2-High
Free Riders	20 (27%)	16 (21%)	34 (45%)	20 (27%)
Conditional Cooperators	45 (60%)	44 (59%)	26 (35%)	43 (57%)
Full Contributors	7 (9%)	10 (13%)	11 (15%)	9 (12%)
Noisy Contributors	3 (4%)	5 (7%)	4 (5%)	3 (4%)
Total	75(100%)	75 (100%)	75(100%)	75(100%)

2.8.7 Distribution of Types Based on Classification using LCP

Table 2.18 presents the distribution of types in P1 and P2 experiments at different levels of MPCR.

2.8.8 Belief Formation in Repeated Games

We used the regression approach of Fischbacher and Gächter (2010) to study belief formation in repeated games. Regressions are conducted with the belief in a given round as the dependent variable and the belief from the previous round and the average contribution of other members of the group in the previous round as explanatory variables. Table 2.19

Table 2.19: Beliefs in repeated games

Belief_t	(1)	(2)	(3)
Belief_{t-1}	0.38***	0.39***	0.38***
	(0.034)	(0.088)	(0.028)
Average Others' $\text{Contribution}_{t-1}$	0.62***	0.62***	0.62***
	(0.034)	(0.084)	(0.032)
Constant	-0.026	0.049	-0.056
	(0.14)	(0.48)	(0.10)
Obs	1920	720	1200
R^2	0.87	0.85	0.88

$^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$
Robust standard errors clustered at session level in parentheses

presents the regression estimates. Model (1) presents results for all sessions, whereas Models (2) and (3) report results for sessions with finitely repeated and indefinitely repeated games respectively. In all cases, the coefficient on the constant is insignificant and the sum of the coefficients on the belief and average contribution of others in the previous round is not different from one (all sessions: $F(1,9) = 0.44, p = 0.522$; finitely repeated sessions: $F(1,3) = 0.01, p = 0.944$; indefinitely repeated sessions: $F(1,5) = 0.89, p = 0.388$). Beliefs in a given round are updated as a weighted average of the beliefs in the previous round and average of the others' contributions in the previous round. The results are consistent with the findings in Fischbacher and Gächter (2010), and they use Strangers' matching in the experiments. The regression coefficients are not different across the finitely and indefinitely repeated games (Chow-test, $F(3,9) = 0.04, p = 0.990$). Thus, we conclude that belief formation is identical across finitely and indefinitely repeated games.

2.8.9 Validation Experiment Instructions

INSTRUCTIONS

Welcome and thank you for participating in todays economic experiment. Please put away all your belongings and turn off your cell phones. You are not allowed to talk to any other participant during the experiment. If you have any questions, please raise your hand. We will come to you and answer your questions in private. The experiment will be run entirely on the computer and all interactions between yourself and others will take place via the computer terminal.

You have earned $5 just for showing up on time. This is yours to keep. In addition, depending upon the decisions you make, the decisions others make and random choice, you can earn more money. These instructions describe in detail the experiment and the tasks you are asked to complete.

During the experiment, your earnings will be described in terms of tokens. At the end of the experiment, the total number of tokens you have earned will be converted to money at the following rate:

20 tokens = $1 (1 token = 5 cents)

The experiment consists of two tasks. You will receive instructions for each task prior to making decisions for that task. Your total earnings from the experiment will be the sum of your earnings in each task. At the end of the session, the total number of tokens earned across the two tasks will be converted to money and paid to you privately in cash, along with the $5 show-up fee.

In each task, all participants will be randomly divided in groups of four members. Participants do not know who is in which group.

The Decision Situation

You will be a member of a group consisting of 4 people. Each group member is provided with 20 tokens and must decide how to allocate the tokens between a private account and a group project. You can choose to contribute none, all, or some of your tokens to the group project. The tokens you choose not to contribute to the group project will remain in your private account. Everyone makes the same decision.

For each token put in the private account you earn exactly one token. For example, if you put 20 tokens into your private account (and therefore do not contribute to the group project) your earnings from this account will be 20 tokens. If you put 6 tokens into your private account, your earnings from this account will be 6 tokens. You are the only one who earns tokens from your private account.

What you earn from the group project depends on the total number of tokens that you and the other members in your group contribute to the group project and a return rate. The return rate, denoted by M, specifies how much a token contributed to the group project returns to every member of the group. The return rate M will always be strictly between 0.25 and 1 in the experiment. For example, if the return rate is 0.5, then each token contributed to the group project returns 0.5 tokens to every member of the group. You will always be told the return rate before you make your contribution decision. The more each member of the group contributes to the group project, the more each member earns from the group project.

Your earnings from the group project are best explained by a number of examples.

Example 1: Suppose that you decided to contribute no tokens to the group project but the three other members in your group contribute a total of 40 tokens. This makes a total contribution of 40 tokens to the group project.

Suppose each token contributed to the group projects returns M = 0.3 tokens to each group member. Then your earnings from the group project would be 40 tokens $\times$ 0.3 = 12 tokens. Everyone else in your group would also earn 12 tokens from the group project.

Suppose instead that each token contributed to the group projects returns M = 0.5

tokens to each group member. Then your earnings from the group project would be 40 tokens $\times$ 0.5 = 20 tokens. Everyone else in your group would also earn 20 tokens from the group project.

Suppose instead that each token contributed to the group projects returns M = 0.8 tokens to each group member. Then your earnings from the group project would be 40 tokens $\times$ 0.8 = 32 tokens. Everyone else in your group would also earn 32 tokens from the group project.

Example 2: Suppose that you decided to contribute 10 tokens to the group project and the three other members in your group contribute a total of 40 tokens. This makes a total contribution of 50 tokens to the group project.

Suppose each token contributed to the group projects returns M = 0.3 tokens to each group member. Then your earnings from the group project would be 50 tokens $\times$ 0.3 = 15 tokens. Everyone else in your group would also earn 15 tokens from the group project.

Suppose instead that each token contributed to the group projects returns M = 0.5 tokens to each group member. Then your earnings from the group project would be 50 tokens $\times$ 0.5 = 25 tokens. Everyone else in your group would also earn 25 tokens from the group project.

Suppose instead that each token contributed to the group projects returns M = 0.8 tokens to each group member. Then your earnings from the group project would be 50 tokens $\times$ 0.8 = 40 tokens. Everyone else in your group would also earn 40 tokens from the group project.

As you can see, it does not matter who contributes tokens to the group project. Everyone will get a return from every token contributed-whether they contributed tokens in the group project or not.

Your total earnings will be the sum of the tokens you earn from your private account and the tokens you earn from the group project. Therefore, your total earnings will be:

Your total earnings = 20 - your tokens contributed to the group project + M $\times$ sum of tokens contributed to the group project by everybody in your group

You will now complete some questions to make sure everyone understands how earnings are calculated.

Questions

Subject number: ________________

Please answer the following questions. These will help you understand how earnings are calculated. Your payoff is not affected by your answers to these questions.

Each token in the private account earns 1 token.

1. Each participant has 20 tokens. Suppose you contribute 10 tokens to the group project and the other three other members in your group contribute 40 tokens in total. Each token in the group project earns M = 0.5 tokens for each participant in the group.

What are your earnings from your private account? ________________

What are your earnings from the group project? ________________

What are the earnings of your group members from the group project? ________________

What are your total earnings? ________________

2. Each participant has 20 tokens. Suppose you contribute 20 tokens to the group project and the other three members in your group contribute 30 tokens in total. Each token in the group project earns M = 0.8 tokens for each participant in the group.

What are your earnings from your private account? ________________

What are your earnings from the group project? ________________

What are the earnings of your group members from the group project? ________________

What are your total earnings? ________________

3. Each participant has 20 tokens. Suppose you contribute 12 tokens to the group project and the other three members in your group contribute 38 tokens in total. Each token in the group project earns M = 0.3 tokens for each participant in the group.

What are your earnings from your private account? ________________

What are your earnings from the group project? ________________

What are the earnings of your group members from the group project? ________________

What are your total earnings? ________________

Instructions for Task 1 - A-Task

In the A-task, you will make a choice for the decision situation described earlier. You will have 20 tokens and must decide how many to put into your private account and the group project. You will be randomly assigned to a 4-member group and will not know who is in your group.

Each token contributed to group project in this task returns M = 0.5 tokens to each group member.

Each participant has two decisions to make in this task: make an unconditional contribution and complete a contribution table.

Unconditional Contribution: In this decision, you must decide how many of the 20 tokens you would like to contribute to the group project. You will make your decision on a screen such as the following.

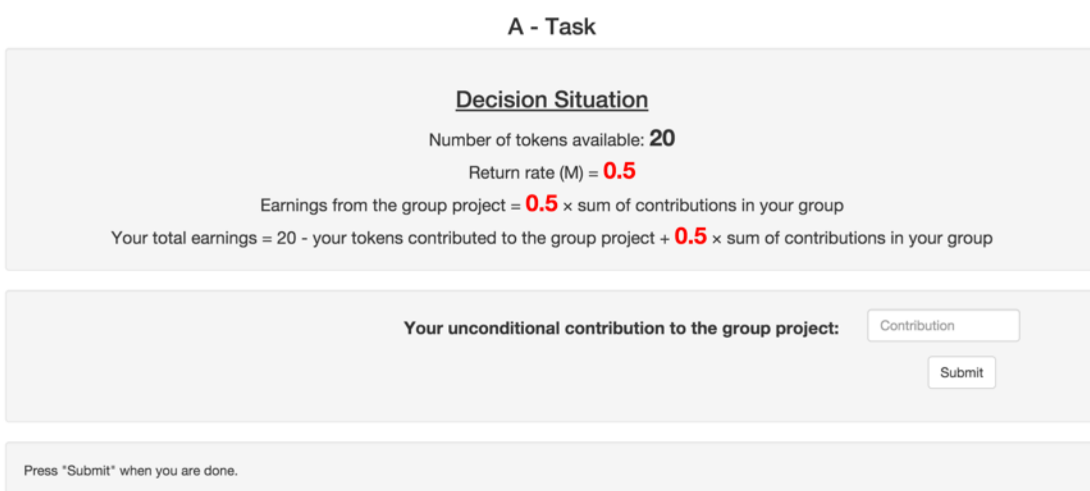

Figure 2.15: Decision screen for unconditional contribution choice in strategy games

Contribution Table: In this decision, you must decide how many tokens you would like contribute to the group project for each possible average contribution of the other group members (e.g. 0, 1, 2,..., 20). That is, if the other group members contributed 0 tokens on average, how much would you contribute? If they contributed 1 token on average, how much would you contribute? If they contributed 2 tokens on average, how much would you contribute? And so on, up to 20 tokens on average.

This means that in total you have to give 21 responses. You will make your decisions on a screen such as the following:

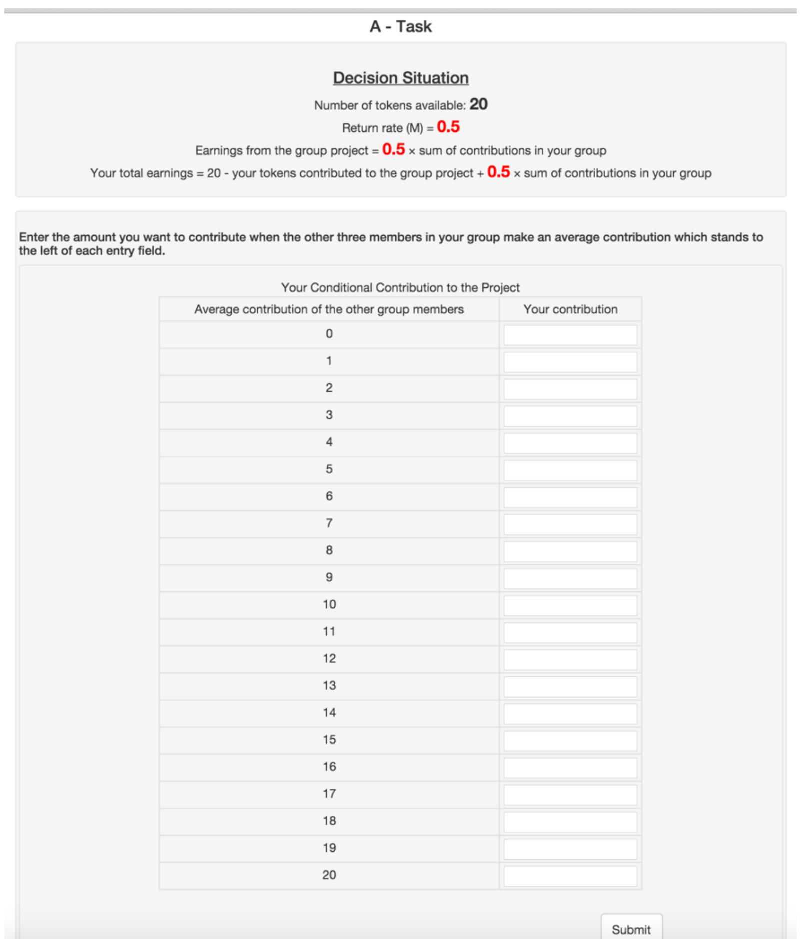

Figure 2.16: Decision screen for conditional contribution table in strategy games

Once each participant has made the unconditional decision and completed the contribution table, the computer will randomly determine if the unconditional contribution or the contribution table will be used to determine your earnings. In each group, one of the four participants is randomly chosen to have the contribution table count to calculate earnings. For the other three participants in the group, the unconditional contribution counts to calculate earnings. How is this done? If the participant is chosen to have his contribution table count for earnings, first the unconditional contributions to the group project of the three other members in his group are averaged and rounded to the nearest integer (e.g. 0,1,2, ,20). Then, the contribution table of the participant is used to determine how many tokens the participant contributes to the group project. The number of tokens contributed is the amount he specified for the average contribution of the other members in his group.

For example, if the average contribution of the other group members in his group is 16 tokens (48 tokens/3), and he specified 10 tokens if the average contribution of the other group members is 16, the total contributed to the group project by everyone in the group would be 58 (10 + 48) tokens.

You will not know in advance which decision, the unconditional contribution or the contribution table, will count to determine your earnings, so you should make each decision carefully as though it will count.

The following examples should help make this procedure clear.

Example 1: Suppose the contribution table was randomly chosen to count for you. This means that the decisions you made in the contribution table determine your earnings. For the other three members of your group their unconditional contributions determine their earnings. Suppose that the total contribution of the other three group members is 39 tokens, and the average contribution 13 tokens (39 tokens/3). In your contribution table, suppose you chose to contribute 11 tokens if the average contribution of the other group members is 13, then your earnings for Task 1 would be: $20 - 11 + 0.5 \times (11 + 39) = 34$. If instead you chose to contribute 1 if the average contribution of the other group members is 13, your payoff would be: $20 - 1 + 0.5 \times (1 + 39) = 39$.

Example 2: Suppose the unconditional contribution was randomly chosen to count for your earnings. Also, suppose that the unconditional contributions of the two group members who were not selected for the contribution table to count are 12, 18. If your unconditional contribution is 20, then the average unconditional contribution is 17 tokens ((20 + 12 + 18)/3). If the group member selected to have his contribution table count chose 10 tokens if the average contribution of his other group members is 17, then your earnings are: 20 - 20 + 0.5 × (20 + 12 + 18 + 10) = 30.

Are there any questions before we begin?

Instructions for Task 2 - B-Task

In the B-Task, you will be randomly assigned to a group of four as described earlier. Your group members in Task 2 may be different from your group members in Task 1, however, you will remain with the same group members for all decisions you make in Task 2.

The B-Task lasts for 10 rounds. In each round, you will be given 20 tokens and must decide how many tokens you would like to contribute to the group project and how many you would like to put in your private account.

Prior to making your contribution decision in each round, you will be informed of the return amount M ($0.25 < M < 1$) from each token contributed to the group project by your group. The return amount M could change across rounds. If each token contributed to the group project returns M tokens to each group member in a given round, your earnings from your contribution decision in that round are:

Your total earnings = 20 - your tokens contributed to the group project + M × sum of tokens contributed to the group project by everybody in your group

You will participate in the decision situation repeatedly with the same group members for all 10 rounds.

You will make decisions on a screen such as the following:

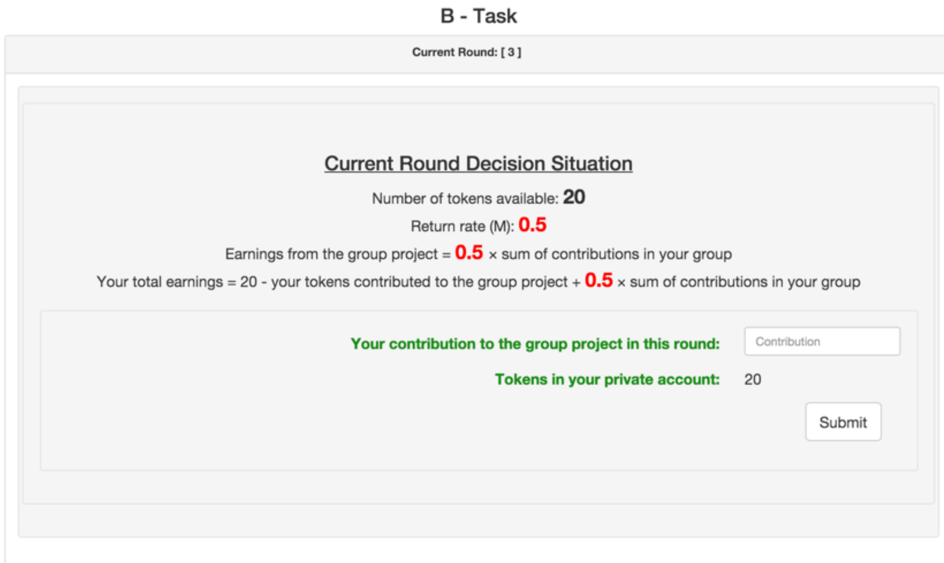

Figure 2.17: Decision Screen for contribution choice in repeated game

Here is an example to explain how earnings are calculated in each round:

Example 1: Suppose the return rate M = 0.5 tokens for a particular round. Suppose you chose to contribute 10 tokens and the other three group members chose to contribute 40 tokens in total in that round. Your earnings in that round would be: 20 - 10 + 0.5 × (10 + 40) = 35 tokens.

In each round, after you decide how much to contribute to the group project, you will be asked to guess the average contribution to the project (rounded to the nearest integer) of the three other members of your group. You will receive tokens for the accuracy of your estimate. If your guess is exactly equal to the average contribution of the other group members, you will receive 3 tokens in addition to your earnings for that round. If your guess was off by 1 token, you will get 2 additional tokens. If your guess was off by 2 tokens, you will get 1 additional token. And, if your guess was off by 3 or more tokens, you will get 0 additional tokens.

When everybody in your group has completed the two decisions in a given round, you will be shown the total contributions to the group project, the average contribution of the other three group members, and your total earnings in that round.

Once all subjects in the experiment have completed the two decisions and are told their earnings and the average contribution of their group members in the current round, the task proceeds to the next round. After the 10th round, the task is done and there will be no more rounds.

Are there any questions before we begin?

Chapter 3: Cooperation in Networked Communities: An Experiment and an Empirical Agent-Based Model

Abstract

Understanding the fundamental patterns and determinants of cooperation within social groups remains a challenge across disciplines. It has become an accepted paradigm that social networks are important determinants of group outcomes. Results from agent-based simulations suggest that presence of a network structure that constrains the interaction among individuals can prevent complete free riding in social dilemmas. Furthermore, social diversity due to degree heterogeneity in networks was found to enhance cooperation significantly. However, the behavioral specification of agents in these studies is often formulated after evolutionary mechanisms inspired from microbiology. In this Chapter, I empirically investigate individual decision making in network public goods games using incentivized experiments. Econometric investigations of the experimental data show that choices in repeated network public goods games are explained by social preferences and reinforcement learning. Using an empirical agent-based model in which behavioral specification of agents is derived from the experimental data, I show that degree heterogeneity, network size, network density, and the average path length of a network have no significant effect on the cooperation levels in public goods games. These results demonstrate that agent-based models based on empirical micro-foundations can lead to very different conclusions than that of agent-based models based on micro-foundations extrapolated from other domains like that of biology. The empirical agent-based model shows that reducing the price of cooperation in the first round can sustain higher levels of cooperation over later rounds of a repeated game. This is due to the way social preferences interact with the price of cooperation and the high inertia associated with reinforcement learning. These insights illustrate that empirical understanding of determinants of behavior in social dilemmas is instrumental to explore effective mechanisms that can promote cooperation using agent-based models.

3.1 Introduction

In this Chapter, I study cooperation in networked communities using empirical agent-based models in which behavioral specification of agents is derived from the econometric analysis of data from behavioral experiments. I use a network public goods game as the prototypical framework to study the problem of cooperation in networks. In a network public goods game, individuals engage in "local" public goods games over a network structure. Numerous studies have used agent-based models to demonstrate that the presence of network structure is sufficient to sustain higher levels of cooperation in cooperative dilemmas including public goods games (Nowak, Bonhoeffer, & May, 1994; Nowak, 2006; Szabó & Fáth, 2007; F. Santos, Santos, & Pacheco, 2008; Poncela, Moreno, et al., 2007; Roca, Cuesta, & Sánchez, 2009). Furthermore, social diversity introduced by heterogeneous graphs like scale-free networks is found to enhance significantly higher levels of cooperation compared to that of homogeneous graphs like regular lattices (F. C. Santos & Pacheco, 2005; F. Santos et al., 2008).

There are two main concerns with the results in the earlier studies about the effect of network structure on cooperation in social dilemmas. First, the behavioral specification of agents in these studies is inspired from insights in microbiology. The agents are assumed to imitate the behavior of more successful individuals in their neighborhoods. The neighborhood of an individual includes itself and all the individuals that are connected to it on a network. In some cases, this imitation is extremely simple in that agents imitate the strategy of the most successful neighbor (Nowak et al., 1994).[1] Success is measured in terms of the payoff obtained in the previous round or cumulative payoff achieved until the current round of play. Sometimes, imitation is stochastic (F. Santos et al., 2008). One of the given agent's neighbors is randomly picked. If the randomly chosen neighbor had achieved a higher payoff then the agent imitates the strategy of the randomly picked neighbor with a probability proportional to the difference between the payoff of the neighbor and that

[1] An agent is a neighbor of itself.

of the agent. While this behavioral description could be best suited to describe decision making of micro-organisms studied in microbiology, it is unclear if it is a good description of how humans make decisions in strategic environments. Second, recent experiments with public goods games over networks have shown no evidence of network structure affecting cooperation levels achieved in a significant manner (Suri & Watts, 2011; Rand, Arbesman, & Christakis, 2011; Wang, Suri, & Watts, 2012). However, the absence network effect in these experiments may not necessarily mean networks are irrelevant in promoting cooperation. A potential problem with these experimental studies is the size of the networks under consideration. In these studies, the networks are too small so they may not be presenting opportunities for significant aggregate differences in outcomes. At the same time, it is also impractical to conduct experiments with thousands of individuals. Large scale experiments can be very expensive and cumbersome. Furthermore, there could be logistic issues in terms of having a lab to accommodate thousands of individuals. In this Chapter, I address the difficulties in empirically investigating the effects networks could have on cooperation in network public goods games by merging experiments and agent-based models. I derive empirical behavioral specification of agents using the data from lab experiments. I then deploy these agents in agent-based models to explore the network effects on cooperation. The empirical agent-based model thus obtained helps one to overcome the limitations of lab experiments in studying large scale networks. Moreover, the empirical behavioral specification derived from experimental data helps one to do away with unverified theoretical behavioral specification of agents. Finally, empirical agent-based models can be used to systematically explore practically countless counterfactual scenarios. For example, the average path length of a network can be varied in the simulations to study the effect of the average path length on cooperation. Likewise, the average clustering coefficient can be varied systematically to study its effects on the aggregate cooperation levels.

While the data from the study of Suri and Watts (2011) on network public goods games can be used to elicit the empirical behavioral specification of agents, there are a few concerns in doing so. First, the experiments in their paper are conducted over internet using

Amazon Mechanical Turk (AMT), and when participants dropped out from the experiments before they ended a computer program made decisions for those participants. If one drops all the sessions where at least one participant dropped out before an experiment ended, the resulting dataset is quite small. Second, some of the experiments involved computer programs making decisions alongside with human participants. The main goal was to embed a full cooperator or a full free rider in the network and study their effect on human participants. This arrangement could potentially have a significant effect on the social preferences and/or learning of human participants. If participants knew that there are non-human agents in the experiment, their behavior could potentially be very different since their concern for the welfare of computer programs could be different to that of their concern for the welfare of human participants. This introduces complications in the estimation of social preferences and understanding their role in the decision making process of individuals. Finally, the data from these experiments consists of choices made only in repeated environments and thus poses a challenge in disentangling the roles of social preferences and leaning in repeated network public goods games. In Chapter 2, I found that both social preferences and learning are important in explaining choices in repeated games. Without having a measurement of social preferences separately, it will be challenging to understand to what extent behavior is due to social preferences and to what extent it is due to learning in a repeated game. Therefore, I do not use data from the study of Suri and Watts (2011) to derive the behavioral specification of agents.

Local public goods games that individuals participate in network public goods games are identical to that of public goods games encountered in Chapter 1 and Chapter 2. The analyses of experimental data in Chapter 1 and Chapter 2 have shown that contribution choices in repeated public goods games are best described by social preferences of participants affecting the choice of first round contributions and then subsequent contributions based on payoff-based reinforcement learning. At first, it is tempting to think that one could directly deploy this empirical behavioral specification for agents to study cooperative behavior in networked environments. However, such an "extrapolation" of behavior from non-networked environments to networked environments implicitly assumes that social preferences and learning are identical in networked and non-networked environments. It remains an open question if these assumptions are true. To this end, I empirically investigate social preferences and learning in network public goods games. To do so, I use data from behavioral experiments. The experimental design is based on that of Fischbacher and Gächter (2010) and is identical to the experimental design in Chapter 2 except that participants are placed on a network. These experiments are designed to observe the behavior of a given participant in four network public goods game settings. These are a one-shot conditional contribution game (labeled the P-task) and a repeated game (labeled the R-task) at two different costs to cooperate (i.e. marginal per capita return (MPCR) levels), one low and one high, that preserve the social dilemma. The one-shot game uses the strategy method to elicit an unconditional contribution and conditional contributions based on several possible average contributions of the neighbors on the network of a given participant. Each one-shot game is then followed by a repeated game at the same MPCR level. In this game, for each participant along with a contribution decision in each round, I also elicited a belief about the average contribution of his neighbors on the network. The elicited beliefs allow for an examination of belief learning over rounds. The one-shot and repeated game sequence is then conducted again at a different MPCR level. In addition, across sessions but not within sessions, I varied whether the repeated games were finite or indefinite. The data from indefinitely repeated games in conjunction with data from finitely repeated games allows me to study if individuals learn differently across these two environments. Finally, the order of which MPCR is seen first is also reversed across sessions to control for experience effects. Thus, the experimental design is identical to that of the experimental design in Chapter 2. The only difference is that groups are separate in the Partners' matching of experiments in Chapter 2 whereas here groups (neighborhoods to be precise) overlap since individuals are connected via a network. This provides a cleaner comparison of data from two experiments to see if the presence of network alone affects contributions in any significant manner.

The results in this Chapter are presented in two stages. The first stage results concern with the decision making of individuals in repeated network public goods games. I find that contributions and beliefs in repeated network public goods games are identical to that of non-network public goods games (Partners' treatment) considered in Chapter 2. Thus, the presence of a network does not seem to be affecting contributions. I also find that the distribution of cooperative types (free riders, conditional cooperators, and full cooperators) in network public goods games is identical to that of the distribution of types in non-network public goods games. I find that the estimated distribution of social preference parameters is reasonably successful in explaining the aggregate changes in the distribution of cooperative types when prices change. The model is able to predict an increasing proportion of full cooperators when the price of cooperation gets smaller. Finally, identical to the non-network case, I find that contribution choices in repeated network public goods games are best described by social preferences of participants affecting the choice of first round contributions and then subsequent contributions based on payoff-based reinforcement learning. Learning is identical across finitely and indefinitely repeated network public goods games. The second stage results are derived from an empirical agent-based model. Agents in the empirical agent-based model are modeled after the results in the first stage. The agent-based model is then deployed to investigate the effects of network size, degree heterogeneity, neighborhood size, network density, and the average path length on contributions. All of these structural characteristics of networks do not affect contributions. These results stand in contrast with the results from agent-based models that modeled agents after insights from evolutionary dynamics. Finally, I show that subsidizing cooperation in the first round can lead to higher levels of contributions in network public goods games. This is due to the fact that frequency of full cooperators increases in the first round due to the low price of cooperation and then the high inertia of reinforcement learning sustains these high levels of contributions over later rounds.

This Chapter also contributes to the methodology of agent-based modeling. Realistic representation of behavior and its heterogeneity across individuals is one of the well documented strengths of agent-based modeling framework (Axtell, 2000b, 2007; Gilbert, 2008).

However, there are two long standing methodological challenges in this framework. First challenge concerns with the way behavioral representation is actually reified in agent-based models and the validity of such reification. Second challenge concerns with how heterogeneity is incorporated into the behavioral representation of agents.

Many agent-based modelers subscribe to the notion of near-decomposability of social systems (Simon, 1996). According to this paradigm, a given system is organized into hierarchical layers of components, components of components, components of components of components and so on. The hierarchical structure is so that the behavior at a given level emerges primarily due to the components at the immediate lower level and their organization rather than the components at much lower levels. Many agent-based models subscribe to this paradigm by modeling the aggregate outcome of interest as emerging due to the interaction of components at the level immediately below the level of outcome variable of interest. In other words, behavioral specification of agents is often constrained to one level, the level that is supposed to be immediately below the level of the outcome variable of interest. For example, agent specification often is derived completely from cognitive foundations or social foundations but not mixed. This is because, traditionally social determinants of behavior are put strictly above cognitive determinants of behavior. But, the results from my investigations using the econometric analysis of experimental data show that both cognitive determinants of behavior, i.e. the learning mechanism as characterized by reinforcement learning, and social determinants of behavior, i.e. the social preferences, are important in explaining observed behavior in public goods games. However, this does not necessarily mean both cognitive and social determinants of behaviors are simultaneously important in other domains. For example, when one studies an environment like that of double auction markets the relevance of social preferences in explaining emergent aggregate patterns might be insignificant as attested by experiments (Gode & Sunder, 1993). Thus, when thinking about the relevant ingredients of behavioral specification of agents in studying social systems, it is not straightforward whether it is sufficient to stick to one level of behavioral determinants or to consider behavioral determinants across levels. There is no silver bullet.

Theoretical understanding of hierarchies of behavioral determinants may not help one to come up with a behavioral specification of agents' that sufficiently captures the decision making of individuals across social domains. Significant determinants of behavior in one domain (e.g. a domain characterized by a cooperative environment) may or may not be effective in explaining behavior in a different domain (e.g. a domain characterized by a competitive environment). While observational data can be used to determine the behavioral determinants in a given domain, identification could be problematic due to the lack of data from counterfactuals. Furthermore, determinants and the resulting behaviors could be coevolving. Experimental economics solves this problem using randomized experiments to tease out the causal relations in microeconomic systems (Smith, 1976, 1982). In this Chapter, I illustrate how incentivized experiments can be used to identify relevant behavioral determinants that can then be used to specify the behavior of agents in agent-based models. My study involves network public goods games, but it is straightforward to extend this type of approach to other domains.

Even after one identifies behavioral determinants with respect to the problem at hand, there is ambiguity in the agent-based modeling literature about how one goes about characterizing and implementing the heterogeneity in behavioral determinants across agents. It is well argued that it is realistic to assume agents are heterogeneous in their behaviors and the incorporation of heterogeneity is one of the foundational principles of agent-based modeling. However, the literature on how this is to be done in a systematic way is scant. Sometimes, the approaches are ad-hoc. For example, the parameters that characterize behavioral heterogeneity are assumed to be in some range or distributed according to some probability distribution. These decisions are more often made based on intuition or established rules of thumb. Even in cases where the distributional assumptions about the parameters are based on well-documented domain specific stylized facts, more often than not, the distributions are not estimated from the data. In this Chapter, I illustrate that econometric methods that allow for observed and unobserved parameter heterogeneity across individuals like that of Mixed Logit methods (Train, 2009) provide a systematic foundation for characterizing and incorporating behavioral heterogeneity across agents in agent-based models. These methods are readily available in the main stream econometric packages of statistical platforms like that of MATLAB, R, or STATA for conducting estimations with empirical or experimental data. It is straightforward to use the estimated distributions of behavioral parameters in specifying the heterogeneity across agents. This approach has potential to standardize how agent-based modelers implement agent heterogeneity.

The Chapter is organized as follows. In Section 3.2, I describe the network public goods environment. Section 3.3 presents the experimental design, procedures, data, and econometric investigations. Section 3.4 describes the empirical agent-based model and the results. I close with conclusions in Section 3.5.

3.2 Network Public Goods Game

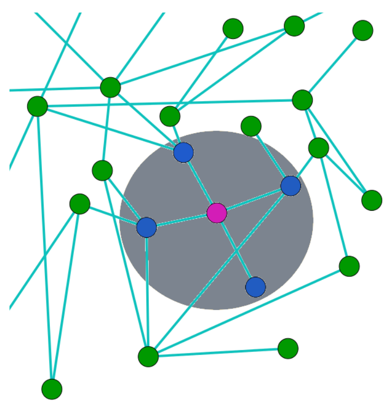

Figure 3.1: A network public goods game

Network public goods games have been the workhorse of studying cooperation in networked communities (F. Santos et al., 2008; Suri & Watts, 2011). In a networked public goods game, players are located on the nodes of a given network and play "local" public goods games with their neighbors. For each player, its neighborhood includes itself and all other players that are connected to it in the network. The number of other players connected to a given player is known as the degree of the player. Therefore, the size of the neighborhood of a player is its degree plus one. In a network public goods game, every player makes a contribution decision and receives a payoff from the local public goods game that is centered on it. Importantly, each member of the network receives information on the contributions of others only within its neighborhood. Figure 3.1 shows an example of a networked public goods game. Here the player situated on the pink node receives a payoff from the public goods game centered on it that includes all the players situated on the blue nodes in the highlighted neighborhood.

Consider a network public goods game consisting of N individuals. Index these individuals as $i = 1, 2, .., N$. Denote the neighborhood of i using Γ_i and the size of i's neighborhood as $N_i = |\Gamma_i|$. Note that i's neighborhood includes i itself. Each individual is endowed with an amount w_i and must decide on a contributions decision $c_i \in [0, w_i]$. In this paper, I study public goods games where each individual is endowed with the same amount $w_i = E, \forall i$. The payoff from the local public goods game of i is given as:

$$\pi^i = E - c_i + \frac{r}{N_i} \sum_{j \in \Gamma_i} c_j \tag{3.1}$$

Where r is called the enhancement factor. As long as $1 < r < N_i, \forall i$, it is easy to see that individually rational decision is not to contribute anything but the social optimum is achieved when everybody contributes their entire endowment. In this way, a network public goods game extends the tension between private and public interests that is the basis of any standard social dilemma to a networked community. Note that, in this formulation $M = \frac{r}{N_i}$ is the marginal per capita return (MPCR or simply M) from the local public goods game of individual i.

The network public goods game is often framed in two ways. In the first way, a fixed r is chosen to model situations where an individual's contribution is split evenly across all the local public goods games of its neighbors and itself. Note that when one uses this formulation, if the degrees of players are heterogenous, it is possible that for some players the social dilemma may not hold for a fixed r. For example, consider $r = 3$. For a player whose degree is 4, the MPCR will be $M = \frac{3}{5}$ and the local public goods game is a social dilemma. Now consider a player whose degree is 2, i. e. he is connected to two other players. For this player, MPCR is $M = \frac{3}{3} = 1$. Thus, the local public goods game is not a social dilemma because the full contribution is the Nash equilibrium. This problem can be overcome by choosing r that is greater than 1 but strictly less than that of the neighborhood size of a node whose degree is the minimum in the network. For example, in a network with a minimum degree of m, the dilemma can be enforced for every player by choosing $1 < r < m + 1$. The second way of framing network public goods game is by choosing a fixed MPCR level. In this framing, an individual's contribution is replicated evenly across all the local public goods games of its neighbors and itself. This framing is useful to capture situations involving, for example, informational goods that can be replicated without any extra cost. This formulation also faces the problem of preserving the dilemma for all players in the network. The problem can be overcome by fixing MPCR at a level so that the social dilemma is preserved for all players. This is achieved by choosing MPCR such that it satisfies $\frac{1}{m+1} < M < 1$ where m is the minimum degree in the network. The payoff from the local public goods game of i is given in this formulation as:

$$\pi^i = E - c_i + M \sum_{j \in \Gamma_i} c_j \tag{3.2}$$

In a repeated game, the stage described above is repeated over T number of rounds. In a finitely repeated game, the number of rounds, T , is determined beforehand. It is easy to see that zero contribution in each round is individually rational in a finitely repeated network public goods game when the social dilemma is preserved for each player in the network.

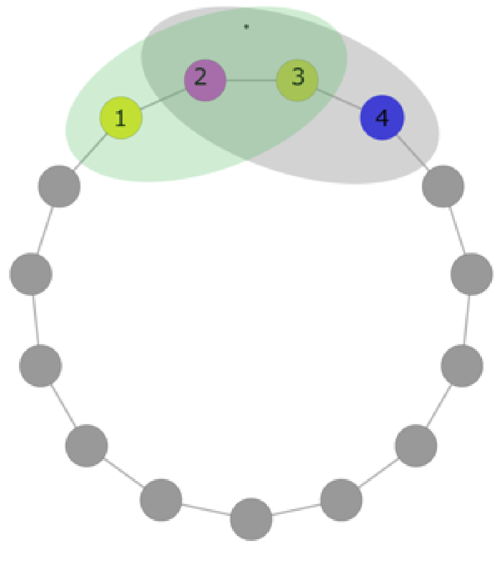

Figure 3.2: Overlapping neighborhoods in a network public goods game

In a network public goods games, neighborhoods of different players overlap. The neighbors of player i are also members of other neighborhoods, and i is a member of other neighborhoods. This is illustrated in Figure 3.2. Consider Player 2 on the node that is colored pink. His neighborhood is shown by the shaded area around the yellow, pink and yellow positions (numbered 1, 2, 3) in the network. The neighborhood of player 2's neighbor to the right is shown by the shaded area around the pink, yellow and blue positions (numbered 2, 3, 4). In other words, player 2 is in its neighborhood and is in the neighborhood of its neighbor to the right. Likewise, player 2 is in the neighborhood of its neighbor to the left. This stands in contrast with the Partners' matching that we considered in Chapter 2. The overlapping neighborhoods can provide different opportunities for strategic behavior and learning. In the next section, I conduct an experiment with a network public goods game and compare its data with that of the experiment using Partners' matching gathered in Chapter 2 to assess if the presence of a network leads to different outcomes and behaviors.

3.3 Social Preferences and Learning in Network Public Goods Games: An Experiment

In this section, I describe an experiment involving a network public goods game. The data from the experiment is used to investigate the roles of social preferences and learning in explaining decision-making in repeated network public goods games. This section is organized into eight subsections. First, I use the model of social preferences by Arifovic and Ledyard to model social preferences in network public goods games. Second, I describe the experimental design. Third, I present a descriptive analysis of the experimental data and its comparison with the data from Partners' treatment in Chapter 2. Fourth, I present estimation results for social preferences using the data on strategy games of network public goods games and compare them with the estimation results from the data on Partners' treatment in Chapter 2. Fifth, I estimate learning models and determine which learning model describes the repeated game data best in network public goods games. Sixth, I disentangle the roles of social preferences and learning in explaining choices in repeated games. Seventh, I investigate if learning is different across finitely and indefinitely repeated network public goods games. I conclude with a short summary in the eighth subsection.

3.3.1 Social Preferences in Network Public Goods Game

It is straightforward to extend the linear social preference model of Arifovic and Ledyard discussed in Chapter 2 to network public goods games. Say Γ_i is the neighborhood of individual i and N_i is the neighborhood size of i.

The payoff that an individual i receives by contributing c^i when others in his/her neighborhood contribute on average o can be written as $\pi^i(c^i, o)) = w - c^i + M(c^i + (N_i - 1)o)$. Where M is the MPCR. Similarly, the average payoff of the neighborhood can be written as $\bar{\pi}(c^i, o) = w - \bar{c} + MN_i\bar{c}$, where $\bar{c} = \frac{c^i+(N_i-1)o}{N_i}$. The utility derived by the individual i:

$$u^i(c^i, o) = \pi^i(c^i, o) + \beta^i \bar{\pi}(c^i, o) - \gamma^i max\{0, \bar{\pi}(c^i, o) - \pi^i(c^i, o)\} \tag{3.3}$$

Where $\beta^i \geq 0; \gamma^i \geq 0$ are social preference parameters. $\beta^i > 0$ implies that individual i has a preference for a higher average payoff to all individuals in his neighborhood. In other words, β^i characterizes an individual's altruistic preference. $\gamma^i > 0$ implies that individual i obtains a disutility when his/her payoff is smaller than the average payoff of the neighborhood, i.e. when $\bar{\pi}(c, o) > \pi^i(c, o)$. γ^i captures the discomfort the individual i face when being taken advantage of by his neighbors.

In the equilibrium individual i would choose a contribution c^i as follows:

$$c^i = \begin{cases} 0 & \text{if } 0 \geq \left(M - \frac{1}{N_i}\right)\beta^i + M - 1 \\ \bar{c} & \text{if } \gamma^i \left(\frac{N_i - 1}{N_i}\right) \geq \left(M - \frac{1}{N_i}\right)\beta^i + M - 1 \geq 0 \\ w & \text{if } \gamma^i \left(\frac{N_i - 1}{N_i}\right) \leq \left(M - \frac{1}{N_i}\right)\beta^i + M - 1 \end{cases} \tag{3.4}$$

3.3.2 Experimental Design & Procedures

The experiment uses a network public goods involving a network as depicted in Figure 3.3. Every individual in the network has exactly two neighbors. The motivation behind choosing this design is that it makes the experiment directly comparable to the experiment in Chapter 2 that involves fixed groups of 3 individuals. The neighborhood size in the network under consideration in the experiments is also exactly 3. By comparing data from the experiments here with the data on public goods games in Chapter 2, I can study the effect of overlapping neighborhoods on the contributions in public goods setting. This is possible since everything other than overlapping neighborhoods is identical across experiments in Chapter 2 and here.

Each individual i is connected to exactly two other individuals on the network. Thus, the neighborhood sizes are identical across all individuals. Each individual is endowed with an amount w_i and must decide on a contributions decision $c_i \in [0, w_i]$. In this paper, I study

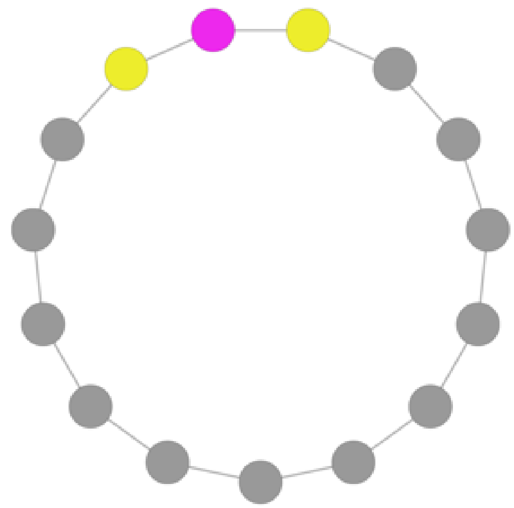

Figure 3.3: Network in the experimental network public goods game

public goods games where each individual is endowed with the same amount $w_i = E, \forall i$. The payoff from the local public goods games of i is:

$$\pi^i = E - c_i + \frac{r}{N_i} \sum_{j \in \Gamma_i} c_j \tag{3.5}$$

Where r is called the enhancement factor. As long as $1 < r < N_i, \forall i$, it is easy to see that individually rational decision is not to contribute anything but the social optimum is achieved when everybody contributes their entire endowment. Since each individual is connected to exactly two other individuals $N_i = 3$ here. Note that in this formulation $M = \frac{r}{N_i} = \frac{r}{3}$ is the marginal per capita return (MPCR or simply M) from the local public goods game of individual i. In terms of marginal per capita return, the payoff function can be written as:

$$\pi^i = E - c_i + M \sum_{j \in \Gamma_i} c_j \tag{3.6}$$

As long as $1 < r < 3$, M will be greater than 0.33 and less than 1. The social dilemma is evident. For any given level of contribution by i's neighbors other than itself, a payoff-maximizing individual's best response is to contribute nothing. Thus Nash equilibrium predicts a zero contribution from everybody in the network. However, a Pareto optimum is achieved when everybody contributes the entire endowment to the public good. In all decisions, there are three individuals in each neighborhood, and the endowment is 20 tokens. All decisions were made with tokens, and converted into a monetary payoff at a rate of 20 tokens = \$1.

Similar to the experiment in Chapter 2, the experiment is based on the Fischbacher and Gächter (2010) design. Each participant is asked to complete four tasks in the following order: P1, R1, P2, R2. In the $P-task$, the participant completes a conditional contribution table in which he decides how much to contribute to the public good account for 21 possible average contribution amounts of his neighbors (e.g. 0, 1, 2, ..., 20). He also makes an unconditional contribution decision by choosing how many of the 20 tokens to contribute to the public good. Payoffs from this task are determined as follows. Once all decisions are completed by the participants, five members are chosen at random. These five members are chosen in such a way that in any neighborhood there is exactly one of the chosen five members present. For example, the five members can be 1, 4, 7, 10, 13. It is easy to see that in any neighborhood exactly one of these members will be present. For these randomly chosen five members, contribution table is payoff relevant. For each of these five members, the unconditional contributions of their two neighbors is averaged and then conditional contribution table is used to determine his contribution decision. For all the remaining participants, their unconditional contributions are pay-off relevant. Everyone knows these procedures before making their decisions.

In the $R-task$, participants make an unconditional decision of how many tokens to put in the public good account. This decision is repeated over several rounds. After deciding how much to contribute to the public good account, a participant is asked to state his belief of the average contribution of the other two members in his neighborhood in the current round. Participants were paid for the accuracy of this stated belief.[2] At the end of a round, a participant is informed of the exact contribution of each of his other two neighbors, the average contribution, and his payoff for that round.

In sessions with finitely repeated games, the number of rounds was fixed at seven. In the sessions with indefinitely repeated games, there was at least one round and then after that the probability of a subsequent round was 0.85. The continuation probability of 0.85 yields, on average, seven rounds of play so the finitely repeated games was fixed at seven rounds to make the games comparable in the expected number of rounds.

Each session had exactly 15 participants. The placement of participants in the network is random, and participants do not know who is connected to whom. Participants were randomly reshuffled on the network before each of the tasks and participants were aware of this before making decisions. Tasks P1 and R1 use the same MPCR level, and tasks P2 and R2 also use the same MPCR level. The MPCR used for P1 and R1 is different than that used for P2 and R2. The two MPCR levels are 0.4 (Low) and 0.8 (High). In five sessions, the low MPCR was used for P1 and R1 and the high MPCR was used for P2 and R2. To control for order effects, in five sessions, this was reversed. Each participant completed all four tasks. Participants knew that there would be four tasks, and the instructions for each task were distributed and read out loud prior to the start of each task.[3] To make sure participants understood the decisions they are asked to make and how to calculate payoffs, a short quiz was administered prior to the start of Task P1. Answers to the quiz were explained before proceeding to the experiment. Participants were paid their earnings for all tasks and all rounds within a task. Earnings were paid privately and in cash at the end of the session.

[2]Specifically, the individual was asked to guess the average contribution of the two other neighbors rounded to the nearest integer. If the guess was exactly equal to the rounded average contribution of the two other neighbors, an individual was paid three tokens ($0.15) in addition to the other experimental earnings. If the guess deviated by only one point, payment was 2 tokens, and if the guess deviated by two points, payment was 1 token. If the guess was off by three points, no tokens were paid. The financial incentives to elicit beliefs were small to avoid hedging.

[3]Participant instructions for sessions with indefinitely repeated games are in Appendix 3.6.1. These include screen shots of the decision screens participants used to make decisions during the experiment. The instructions for the sessions with finitely repeated games are identical except that they mention the repeated games last exactly seven rounds. The instructions for finitely repeated games are available online at *http://www.chennacotla.org/research*

Participants made their decision on a computer, using a web-based software. The experiments were run in the Interdisciplinary Center for Economic Science (ICES) at George Mason University during September and October of 2014. Ten sessions were run, and there were a total of 150 participants. Participants were randomly assigned to cubicles and made their decisions privately and anonymously. Six of the ten sessions involved indefinitely repeated games and four sessions involved finitely repeated games. No one participated in more than one session. Participants were recruited via email from a pool of George Mason University students who had all previously registered to receive invitations for experiments. Each experimental session lasted for approximately 1.5 hours. Average participant earnings were $23.99.

3.3.3 Data

In this section, I describe the conditional contribution decisions from the P-tasks and the contribution and belief decisions from the R-tasks.

To examine the conditional contributions in the P-tasks, I classified participants into cooperative types using the statistical classification algorithm of Kurzban and Houser (2005). This algorithm uses a linear conditional-contribution profile (LCP) to determine a given participant's type. The LCP is the result of an ordinary least squares regression of a participant's conditional contribution in the P-task on each of the 21 possible average contributions of his other two neighbors.. If the estimated LCP is strictly below half the endowment everywhere, then the participant is classified as a Free Rider (FR). A participant is classified as a Full Cooperator (FC) if the LCP lies at or above half the endowment everywhere. If the LCP of a participant has a positive slope and lies both above and below half the endowment then he is a Conditional Cooperator (CC). Any participant who did not fall into one of these three categories is classified as a Noisy Contributor (NC).[4]

The distribution of types identified with the classification algorithm is presented in

[4]In Kurzban and Houser (2005), the authors classify individuals into only the first three groups and exclude analysis on three participants who did not fall into these categories.

Table 3.1: Distribution of cooperative types in strategy games of network public goods games computed using the LCP method.

Type	P1	P2	Low MPCR (0.4)	High MPCR (0.8)
Free Riders	53 (35%)	73 (49%)	71 (47%)	55 (37%)
Conditional Cooperators	77 (51%)	61 (41%)	64 (43%)	74 (50%)
Full Contributors	10 (7%)	11 (7%)	8 (5%)	13 (9%)
Noisy Contributors	10 (7%)	5 (3%)	7 (5%)	8 (5%)
Total	150 (100%)	150 (100%)	150(100%)	150 (100%)

Table 3.1. The first two columns present the distributions observed in the P1 and P2 tasks. The third and fourth columns present the distributions observed in the low MPCR and high MPCR P-tasks respectively. Consistent with previous studies and results in Chapter 2, most participants are classified as conditional cooperators, and roughly one quarter are classified as free riders. These two types account for more than 80% of participants in the combined P1 and P2 tasks and also when separated by the high and low MPCR treatments. Full cooperators and noisy contributors are less frequent. From the P1 to P2 task, as experience increases, there is a significant change in the distribution of types, with an increase in the proportion of free riders and a decrease in the proportion of conditional cooperators (Chi-Squared Test, $\chi^2(3) = 13.89; p = 0.003$). From the low to high MPCR, as the cost to cooperate declines, there is also a significant shift in the distribution, with the proportion of full cooperators rising (Chi-Squared Test, $\chi^2(3) = 8.04; p = 0.045$). A social preference model should be able to explain the transition of the distribution of types when prices change. And, I test this in the later parts of this chapter.

Comparing the distribution of types in network public goods games with that of the distribution of types in Partners' data in Chapter 2, I find that there are no significant differences. For P1, comparing the distribution of types in Table 3.1 with that of distribution of types in Partners' case presented in Table 2.1, the p-value obtained from a Chi-Squared test is 0.104 indicating no significant difference. Similarly, for P2, the p-value is 0.10. At the low MPCR level of 0.4, I find significant difference at the 10% level but not at the 5% level.

Table 3.2: Contributions over rounds in repeated network public goods games

	Contribution			
	(All-1)	(All-2)	(MPCR = 0.4)	(MPCR = 0.8)
Constant	9.18**	9.41***	8.16***	11.59***
	(0.98)	(1.04)	(0.63)	(1.52)
Round	-0.10	-0.11	-0.31**	-0.18
	(0.14)	(0.13)	(0.12)	(0.13)
Finitely Repeated		-0.43	0.41	-1.12
		(2.13)	(1.99)	(2.50)
N	1710	1710	840	870
R^2	0.0013	0.0023	0.0094	0.0092
F	0.51	0.38	3.24	1.48
$Prob > F$	0.4936	0.6975	0.0871	0.2786

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Robust standard errors clustered at the session level are reported. There are in total 10 sessions.

The corresponding p-value is 0.090. At the high MPCR level of 0.8, I find no significant difference in the distribution of types computed using LCP method in both Partners' and Network treatments. The p-value obtained is 0.11. Therefore, it is reasonable to conclude that there is no significant effect of network on the distribution of types observed at different experience levels or at different MPCR levels.

Looking at the contribution behavior of each type, in P1, the average contribution across all conditional contributions is 3.04 tokens for free riders, 9.65 for conditional cooperators, 18.07 for full contributors and 10.17 for noisy contributors. In P2, the average contribution is 2.18 for free riders, 9.33 for conditional cooperators, 18.27 for full contributors, and 11.14 for noisy contributors. In both P1 and P2, average contributions across types are significantly different (Kruskal-Wallis tests, $p < 0.0001$).

Turning now to the R-task, four of my ten experimental sessions involved finitely repeated games of seven rounds each. The remaining six sessions used indefinitely repeated games. The probabilistic continuation rule produced rounds of the following length for R1

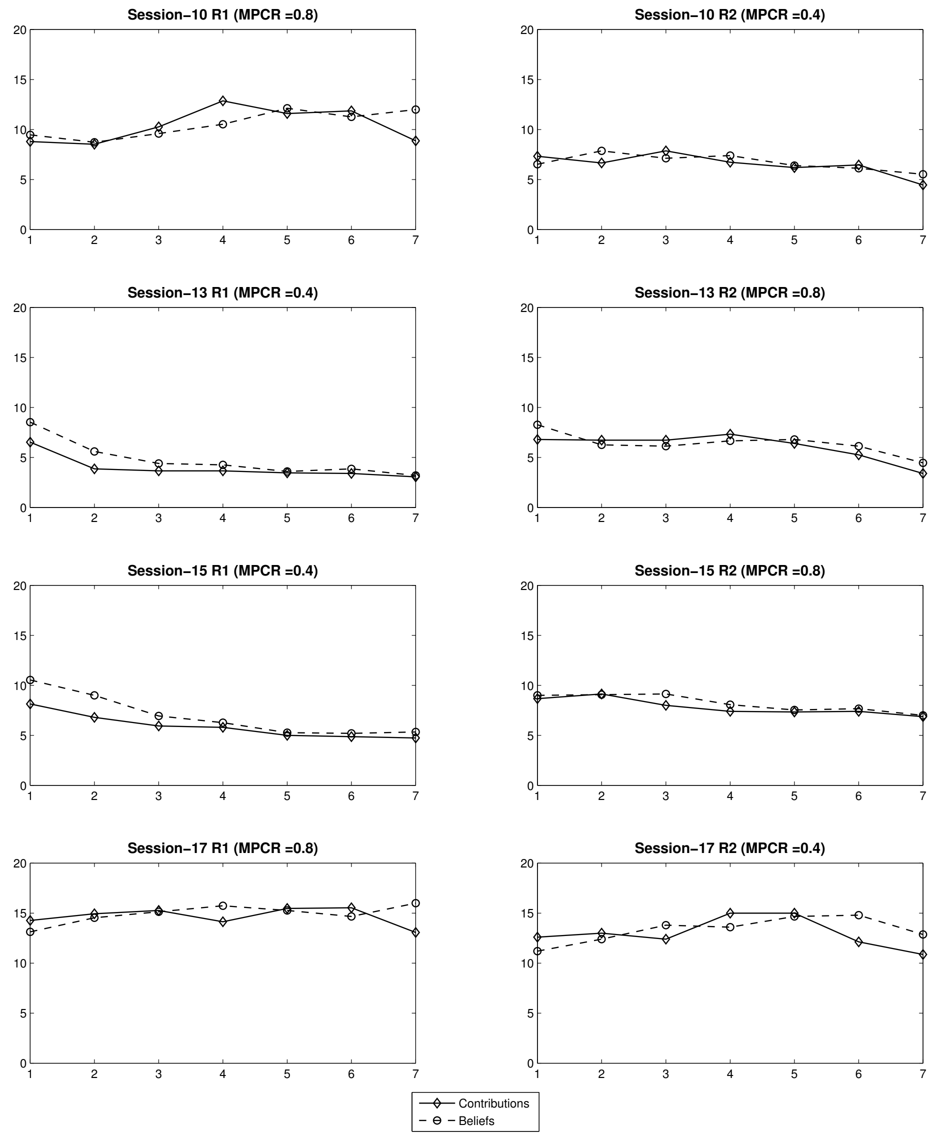

Figure 3.4: Contributions and beliefs in finitely repeated sessions of network public goods games. Average contributions and beliefs are reported for each round in a session.

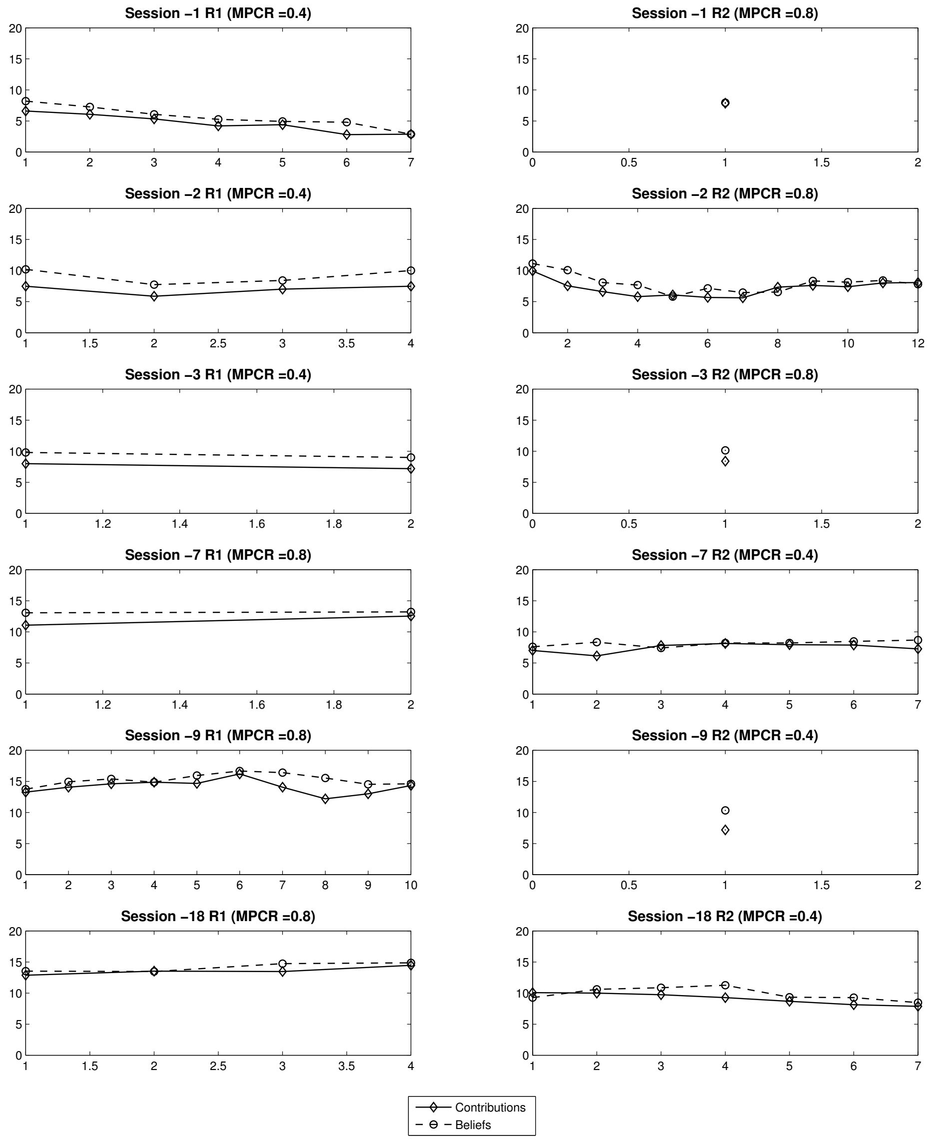

Figure 3.5: Contributions and beliefs in indefinitely repeated sessions of network public goods games. Average contributions and beliefs are reported for each round in a session.

and R2 across the six sessions: {{7,1}, {4,12}, {2,1}, {2,7}, {10,1}, {4,7}}. While the finite and indefinitely repeated games had different lengths, I find no significant difference in average contributions across these games.[5] The high MPCR results in higher contributions ($F(1,9) = 8.30, p = 0.0182$). Surprisingly, when I regress contributions over rounds, I do not find a significant effect on round number. See columns All-1 and All-2 in Table 3.2. However, the combined data has data for both low and high MPCR levels. When I do separate regressions for two MPCR levels, I do find contributions decline significantly in MPCR = 0.4 case. Contributions do not decline when MPCR is 0.8 as shown in the Table 3.2. Sustenance of high levels of contributions at a higher MPCR level is also reported in other studies (Botelho et al., 2009). The table also shows that contributions are not different across finitely and indefinitely repeated games at both levels of MPCR. I also find that elicited beliefs are identical across finitely and indefinitely repeated games.[6] The round level contributions and beliefs are positively and significantly correlated in all ten sessions, as also found in the Partners treatment studied in Chapter 2.[7] Figures 3.4 and 3.5 show session level average contributions and beliefs in finitely and indefinitely repeated games.

Finally, in Table 3.3, I compare contributions and beliefs across Partners' and Network treatments. I combined data from the Partners' treatment in Chapter 2 with the data from the network public goods game in this Chapter. Table 3 shows that contributions and beliefs are identical across Partners' and Network treatments. Thus, presence of a network involving overlapping neighborhoods does not seem to affect contribution choices or belief formation of individuals. The table also shows that contributions and beliefs are also identical across finitely and indefinitely repeated games in both Partners' and Networked treatments.

[5] I test for the effect of indefinite repetition by clustering at the session level and find no effect on contributions ($F(1,9) = 0.04, p = 0.853$).

[6] I test for the effect of indefinite repetition by clustering at the session level and find no effect on beliefs ($F(1,9) = 0.24, p = 0.639$).

[7] The correlations in each of the ten sessions are: 0.74, 0.59, 0.27, 0.65, 0.49, 0.64, 0.65, 0.75, 0.43, and 0.72. Spearman rank correlation tests in all cases, $p < 0.0001$ except in the third case where $p = 0.0723$.

Table 3.3: Contributions and beliefs across Partners and Network treatments

	Contributions		Beliefs	
Constant	10.86***	10.54***	11.23***	10.86***
	(0.72)	(0.60)	(0.70)	(0.63)
Network	-2.10	-1.58	-1.73	-0.86
	(1.26)	(1.36)	(1.24)	(1.40)
Finitely Repeated		0.85		0.98
		(1.68)		(1.57)
Network $\times$Finitely Repeated		-1.26		-2.00
		(2.68)		(2.57)
N	3930	3930	3930	3930
R^2	0.021	0.024	0.022	0.029
F	2.77	0.96	1.93	0.71
$Prob > F$	0.1126	0.4317	0.1812	0.5586

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Robust standard errors clustered at the session level are reported. There are in total 20 sessions. 10 sessions for Partners' treatment and 10 sessions for Network treatment.

3.3.4 Social Preferences

Using the estimation strategy outlined Section 2.5 of Chapter 2, I estimate the distribution of social preference parameters from the data on P-tasks. The aim of the analysis here is two-fold: parallel to what I have done in Chapter 2, I would like to study if experience and MPCR have significant effects on social preferences and if estimated social preference can explain the changes in the distribution of types across MPCR levels.

Table 3.4 reports the estimates for the three parameter vector $\delta^{\eta}, \eta \in \{\beta, \gamma, \omega\}$ for both minimal and full estimation cases. Notation adopted here is from Section 2.5 of Chapter 2. β, γ are modeled as log-normal distributions and ω is modeled as a logistic-normal distribution. The minimal model contains only intercept terms for the parameters. The full model contains an intercept and also treatment dummies for experience and high MPCR for each parameter. The values in the table are in the original parameters' scale. In the full estimation, the constant terms represent the median parameter values for the baseline treatment, which is the P1 task with a low MPCR. Treatment effects reported in the table

Table 3.4: Distribution of social preferences (β, γ) and random choice propensity (ω) parameter estimated using data from the strategy games of network public goods games

	β		γ		ω	
	Min	Full	Min	Full	Min	Full
Constant	16.24***	26.35***	8.88***	2.69***	0.13***	0.028
	(0.54)	(3.44)	(0.30)	(0.47)	(0.013)	(0.052)
Experience		-16.69***		-0.25		-0.028
		(2.75)		(0.35)		(0.051)
High MPCR		-24.31***		1.10		-0.018
		(3.21)		(0.90)		(0.036)
σ	5.05***	2.20***	1.86***	1.67***	2.70***	4.99***
	(0.013)	(0.16)	(0.011)	(0.13)	(0.15)	(1.06)
N	150	150	150	150	150	150
LL	-13856	-13121	-13856	-13121	-13856	-13121

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Coefficients are transformed back to the original scale. A treatment effect is the (partial) effect of moving from the baseline treatment to the corresponding treatment on the median parameter value. Standard errors are calculated using the *sandwich estimator* and treating all of each subjects choices as a single *super observation*. The entries for σ are the standard deviations of the untransformed normal distributions of the random coefficients.

are the calculated partial effects of moving from the baseline treatment to the treatment under consideration using the median parameter value.

According to the minimal model the medians of the estimated β and γ are 16.24 and 8.88. Using the categorization outlined in Appendix 2.8.3 [8], a median participant is a conditional cooperator at both MPCR levels M = 0.4 and M = 0.8. This is in line with the classification results using the LCP method presented in Table 3.1. The median random choice propensity parameter is 13%, which means 13% of the choices in P-tasks can be categorized as being made randomly. According to the estimates for the full model, there are significant effects of MPCR and experience on the median values of social preference parameters. This indicates that the distribution of social preferences need to be shifted significantly to explain the changes in the proportion of cooperative types (free riders, conditional cooperators, and full

[8] note that the categorization can be used as is for the network public goods games since the neighborhood size of each participant is fixed at 3 which is the group size considered in Chapter 2

Table 3.5: Distribution of cooperative types computed from the estimated distribution of social preferences from the strategy games of network public goods games. Full refers to the case where the proportions of types are computed from the distribution of social preferences derived from the full estimation. Min refers to the case where the proportions of types are computed from the distribution of social preferences derived from the minimal estimation.

	P1-Low		P1-High		P2-Low		P2-High	
	Full	Min	Full	Min	Full	Min	Full	Min
Free Riders	22%	39%	3 %	5%	47%	39%	2%	5%
Conditional Cooperators	36%	38%	17%	42%	31%	38%	32%	42%
Full Contributors	41%	23%	80%	53%	22%	23%	66%	53%

cooperators) due to experience and due to changes in prices. Or, in other words, a single distribution of social preferences estimated in the minimal model is not able to account for the changes in the distribution of cooperative types due to the changes in experience and prices. There is also a strong effect of MPCR on random choice propensity parameter. These estimates and effects are very close to what I found in the case of Partners' treatment in Chapter 2 presented in Table 2.1.

To understand the estimates of treatment effects on the social preference parameters, it is important to keep in mind that the distribution of cooperative types is determined by the joint distribution of β, γ conditional on the group size (N) and the MPCR (M).[9] Table 3.5 shows the proportions of types by MPCR and experience level using the estimates in Table 3.4.[10]

Table 3.5 shows that the estimation results from the full and minimal models predict that the proportion of full contributors rises with an increase in the MPCR level. This is consistent with changes in the distribution of types based on the LCP classification

[9] The social preference parameter ranges that lead to different type behaviors are reported in Appendix 2.8.3.

[10] These proportions are determined in a given treatment using the estimated joint distribution of β and γ in that treatment. The joint distribution of β and γ is identical across treatments when results from the minimal model are used. When the estimates from the full model are considered, the distributions of β and γ in a treatment are adjusted based on the effects of experience and MPCR. The proportions of different types are computed using Appendix 2.8.3. For example, the proportion of free riders in P1-Low is computed as the probability that $\beta \leq 8.57$, the proportion of conditional cooperators in P1-Low is computed as the probability that $\beta > 8.57$ and $\beta \leq 8.57 + 9.43\gamma$, and finally, the proportion of full contributors in P1-Low is computed as the probability that $\beta \geq 8.57 + 9.43\gamma$ given the joint distribution of β and γ for P1-Low.

Table 3.6: Changes in the percentages of cooperative types in network public goods games predicted by the full model and minimal model in P1 and P2 tasks across low and high MPCR levels. Reported are the differences in percentages obtained by subtracting the percentage of a type at low MPCR from the percentage of a type at high MPCR.

	P1		P2	
	Full	Min	Full	Min
Free Riders	-19%	-34%	-45 %	-34%
Conditional Cooperators	-16%	4%	1%	4%
Full Contributors	39%	30%	44%	30%

Table 3.7: Cross tabulation of types in the strategy games of network public goods games at two MPCR levels: $MPCR = 0.4$ and $MPCR = 0.8$. Types are computed using the Linear Contribution Profile (LCP) method.

		$MPCR = 0.8$		
		FR	CC	FC
	Free Riders (FR)	43	22	5
$MPCR = 0.4$	Conditional Cooperators (CC)	11	47	4
	Full Cooperators (FC)	0	4	4

algorithm reported in Table 3.1. However, the exact proportions of types predicted by the full and minimal model are different across treatments. The full model estimates show that the proportion of free riders substantially rises with experience, moving from 22% in P1-Low to 47% in P2-Low. At high MPCR levels, the experience effect seems to be very small. Percentage of free riders at high MPCR level almost does not change (compare percentage of free riders in P1-High and P2-High). The minimal model cannot explain this effect of experience at low MPCR levels.

One of my primary objectives is to test if the social preference model is able to explain the changes in the distribution of cooperative types when prices change in network public goods games. The estimates from the full model show that this is not the case since they show a significant effect of MPCR on the social preference parameters even after controlling for experience effects. In other words, the distribution of social preferences has to be shifted significantly to explain the changes in the distribution of cooperative types when the prices change. A potential candidate of explanation for the significant effect of MPCR on the social preference parameters can be the misspecification of the distribution of social preferences as mentioned in Chapter 2. As done in Chapter 2, given the log-normal specification assumption for social preferences, I evaluate how worse off the minimal model is in predicting the changes in the distribution of cooperative types across different levels of MPCR compared to that of the full model that incorporates the effects of MPCR. Despite the significant effects of MPCR on the median values of the parameters, if the minimal model is able to capture the general direction of changes in the distribution of cooperative types one could use it as a practical unifying model of organizing different distributions of types under varying prices.

Table 3.6 presents the changes in the percentages of each cooperative type when one moves from low MPCR (0.4) to high MPCR (0.8). The table reports the changes in the percentages of types using the minimal model and full model. The changes in the percentages of types is reported for P1 and P2 tasks separately. This helps to understand separately how successful is the minimal model in explaining the changes in the distribution of types due to the changes in MPCR when participants do and do not have experience. In both P1 and P2, the full and minimal models predict that percentage of free riders decreases and the percentage of full cooperators increases when one moves from low MPCR to high MPCR. In P1, there is disagreement between the full model and minimal model in terms of change in the percentage of conditional cooperators. There is a substantial difference in terms of the magnitude of change in the percentages of free riders and conditional cooperators between the full and minimal models in P1. However, both models provide a very close prediction of changes in percentages of types in P2. Thus, if one considers that with experience participants' responses are less noisy, the minimal model is quite good at describing how the proportions of cooperative types change when the MPCR changes.

Another way to evaluate the success of social preference model is to see how many individual level type transitions are consistent with the model. A type transition at the individual level is consistent if one can find at least one pair of social preference parameters according to which the transition is possible. Appendix 2.8.3 describes which individual level transitions are possible and which are not possible according to the social preference model. The only individual level transitions that are not possible are the transitions from CC to FR, from FC to FR, and from FC to CC when MPCR changes from 0.4 to 0.8 (or the transitions in the opposite direction when MPCR changes from 0.8 to 0.4). A small number of these transitions in the data would support that social preference model is relatively successful. Table 3.7 shows the individual level types at the two MPCR levels in the experiment. Individual level type information is computed using the LCP method as described in Section 2.3. I have dropped the cases where the transitions involved a noisy type. In total, the table reports the type transition for 140 individuals in the data. There are only 15 cases where type transitions are not consistent with the social preference model. Eleven of them involve a transition from CC to FR when MPCR changes from 0.4 to 0.8, zero of them involve a transition from FC to FR when MPCR changes from 0.4 to 0.8 and four of them involve transition from FC to CC when MPCR changes from 0.4 to 0.8 . Thus, only 11% of individual level transitions cannot be accounted by the social preference model indicating that the model is reasonably successful in organizing changes in behavior when prices change.

3.3.5 Learning

In this subsection, I investigate how do people learn in network public goods game. The analysis is parallel to that of the analysis done in Chapter 2 with data on Partners' treatment. The estimation strategy is identical to that of what is presented in Section 2.5 of Chapter 2. I only consider the learning models that allow for parameter heterogeneity since I found parameter heterogeneity improves the fit significantly in Chapter 2. Furthermore, when I parameterize empirical agent-based models, parameter heterogeneity provides a natural way to introduce heterogeneity in the behavior of agents.

I consider REL as the representative model of reinforcement learning and NFP as the representative example for belief learning. These models can help me to examine whether belief learning or reinforcement learning best explains the choices when parameter heterogeneity is allowed and when merged with social preferences. For each of these learning models, I estimate two variations: payoff based learning and utility based learning. In the payoff based version of a model, attractions of the strategies in the learning model are updated using the payoffs in the game. In the utility based model, attractions of the strategies are updated using the utilities computed with social preferences. I use the distribution of social preferences estimated from strategy experiments (P-tasks) in computing the utilities. I consider estimates from the minimal estimation in Table 3.4. The payoff based variation of a learning models is suffixed with PAYOFF and the social preferences based variation of a learning model is suffixed with UTIL. Finally, AL models individuals as best responding to their elicited beliefs and evaluates how well the elicited beliefs capture learning of individuals. Here, the best responses are computed by integrating over the distribution of social preferences estimated from P-tasks. AL has two free parameters which characterize the distribution of its random choice propensity parameter.

All the learning models are estimated using the aggregated repeated game data. To keep all models on the same footing initial attractions of all strategies in the first round are set to zero. Thus in the first round, all choices have equal probabilities for being chosen according to all these models. I revisit this assumption in the next subsection. Table 3.8 reports the parameter estimates of each of the models and the corresponding fit metrics. I report the maximum log-likelihood (LL) achieved, AIC, BIC, and Pseudo-R^2 for each learning model.[11] The performance of a random choice model, RAND, is also reported for

[11]
$$AIC = LL - k$$
$$BIC = LL - \frac{k}{2} ln(Ntotal)$$
$$Pseudo\text{-}R^2 = \frac{AIC - LLR}{LLR}$$

Where k is the number of the parameters of the learning model, $Ntotal$ is the total number of observations, and LLR is the log-likelihood obtained by a random choice model. Since choices for an individual are not independent I consider the number of individuals as the effective size of the sample in the computation of BIC. Therefore, $Ntotal$ is equal to 150 here. AIC and BIC penalize models for their complexity and therefore

Table 3.8: Comparison of learning models using data from repeated network public goods games. UTIL means utilities computed using social preferences are used for attraction updating and PAYOFF means pure payoffs in the repeated game are used for attraction updating in the learning model.

Model	Parameter Estimates				LL	AIC	BIC	Pseudo-R^2
AL	ν_ω	0.3172	σ_ω	2.0347	-3938	-3940	-3944	0.2430
REL-UTIL	ν_λ	1.6456	σ_λ	0.2568	-3962	-3966	-3972	0.2381
	ν_{N1}	0.2647	σ_{N1}	6.6453				
REL-PAYOFF	ν_λ	1.4277	σ_λ	0.3255	-3956	-3960	-3966	0.2394
	ν_{N1}	0.0850	σ_{N1}	4.0878				
NFP-UTIL	ν_λ	0.9775	σ_λ	3.4301	-4308	-4312	-4318	0.1718
	ν_w	0.1836	σ_w	3.5615				
NFP-PAYOFF	ν_λ	71.962	σ_λ	1.6453	-4717	-4721	-4727	0.0932
	ν_w	0.0152	σ_w	3.5893				
Random Choice Benchmark								
RAND‡					-5206	-5206	-5206	0.0000

† ν stands for the median of the distribution of a parameter. Therefore, ν_λ refers to the median value of λ. σ stands for the standard deviation of the untransformed normal distribution of a random coefficient. Therefore, σ_λ is the standard deviation of the normal distribution that was exponentiated to obtain the log-normal distribution of λ.

‡ RAND is the random choice model. According to this model any choice has a probability of $\frac{1}{21}$ of being chosen.

Table 3.9: Comparison of learning models using data from repeated network public goods games. UTIL means utilities computed using social preferences are used for attraction updating and PAYOFF means pure payoffs in the game are used for attraction updating in the learning model. For each pair of learning models, a two sided p-value favoring the learning model in the row is reported. The two-sided p-values are computed using the Vuong Test for Non-Nested Models.

	REL-UTIL	REL-PAYOFF	NFP-UTIL	NFP-PAYOFF	RAND
AL	0.7092	0.7723	0.0000	0.0000	0.0000
REL-UTIL		0.4283	0.0000	0.0000	0.0000
REL-PAYOFF			0.0000	0.0000	0.0000
NFP-UTIL				0.0000	0.0000
NFP-PAYOFF					0.0000

comparison purposes.

Comparing the variations of the models that consider payoffs in the learning and utilities based on social preferences in the learning, it can be seen that the results are mixed. For REL gains in the fit are negative when merged with social preferences. The gains due to the consideration of social preferences are very prominent in the case of NFP. AL achieves the best fit in terms of all four fit metrics. However, one needs to keep in mind that AL uses twice as much of the data as all other learning models use since AL uses elicited beliefs in each round along with choices. Since, all the models considered here, AL, REL-PAYOFF, REL-UTIL, NFP-PAYOFF, and NFP-UTIL are non-nested, I assess their performance in pairs using the Vuong's test for non-nested models (Vuong, 1989) as a model selection criterion. Table 3.9 presents p-values computed using the Vuong test for each pair of learning models. A two sided p-value favoring the learning model in the row is reported. AL, REL-UTIL, REL-PAYOFF perform equally well. This is significant since it indicates that the REL variations are able to perform as well as AL though AL uses twice the amount of the data compared to that of REL variations. The p-values computed for AL and REL provide an unfair advantage to AL since the Vuong test does not take into account the fact that AL uses twice the amount of data compared to that of REL. Interestingly, for REL it does not seem to matter whether attractions are updated using the payoffs or utilities. This is in contrast with the results from NFP. NFP-UTIL which uses the utilities computed with social preferences to update attractions perform significantly better.

The results on this section reinforce my findings in Chapter 1 and 2. REL provides the best account of observed choices in repeated network public goods games. NFP performs significantly better when it is merged with social preferences. However, it does not beat the both variations of REL. For REL, the payoff based attraction updating works as well as the utility based attraction updating. Thus, I conclude that payoff-based REL provides the best and parsimonious account of learning in repeated network public goods games.

---

are superior metrics than log-likelihood for model comparison.

3.3.6 Roles of Learning and Social Preferences in Explaining Repeated Network Public Goods Game Data

In Chapter 2, I argued that choices in the repeated games are best described by social preferences affecting the choice of first round contributions and then subsequent contributions based on payoff based learning in public goods games with Partners' matching. I follow up that observation here in the context of network public goods games.

First I note that there are substantial differences in contributions across cooperative types. For example, using the LCP classification of types in P1-task as in Section 3.3.3, the average contributions of free riders, conditional cooperators, and full cooperators are 5.73, 10.85, and 16.09 respectively in R1. The contributions are significantly different across types (Kruskall-Wallis Test: $p < 0.0001$), and there is a significant increasing trend in contributions across the groups (Jonckheere-Terpstra Test, $p < 0.0001$). Similarly, using the type information from the P2-task, the average contributionss of free riders, conditional cooperators, and full cooperators are 5.53, 9.21, and 16.60 respectively in R2. The contributions are significantly different across types (Kruskall-Wallis Test: $p < 0.0001$), and there is a significant increasing trend in contributions (Jonckheere-Terpstra Test, $p < 0.0001$). This indicates that social preferences elicited in P-tasks have a nontrivial relevance in explaining the average contributions in the repeated games.

It is possible that aggregate differences across types are arising simply because they are starting at different levels of initial contributions and their learning in later rounds is characterized by high inertia like that of REL's. This stands in contrast with the assumption in earlier literature that the aggregate differences across types arise because types start with random contributions and learn their equilibrium contributions over time. It is also possible that aggregate differences are emerging because individuals start with different initial contributions and then they also proceed to learn equilibrium contributions. To be able to systematically disentangle the roles of social preferences and learning in repeated games, I consider four models that involve different assumptions about learning and social preferences.

- **RAND-PAYOFF**: REL learning based on payoffs in the game with random first round contributions. This model does not use any information about the social preferences of the participants. In the learning phase, individuals use the payoffs in each round to update the attractions of strategies in the REL learning model.

- **RAND-UTIL**: REL learning based on utilities in the game with random first round contributions. In the learning phase, individuals use the utilities computed with estimated distribution of social preferences from the P-tasks to update the attractions of strategies in each round in the REL learning model.

- **SP-PAYOFF**: REL learning based on payoffs in the game with the first round contributions determined from the social preference parameters of the participants. The first round contribution decisions of individuals are modeled as best responses to their elicited beliefs in the first round given the distribution of social preferences estimated from P-tasks. In the learning phase, individuals use payoffs in each round to update attractions of strategies in reinforcement learning model.

- **SP-UTIL**: REL learning based on utilities in the game with the first round contributions determined from the social preference parameters of the participants. The first round contribution decisions of individuals are modeled as best responses to their elicited beliefs given the distribution of social preferences estimated from the P-tasks. In the learning phase, individuals use the utilities computed with estimated distribution of social preferences from the P-tasks to update the attractions of strategies in each round in the REL learning model.

The parameter estimates and fit of the four models are reported in Table 3.10. Models whose first round contributions are determined by social preferences achieve higher likelihood compared to the models that use random contributions in the first round. It does not seem to matter much whether payoffs or utilities are used in updating the attractions in the

Table 3.10: Disentangling the role of social preferences and learning in repeated network public goods games. REL is used to model learning. SP means social preferences determine the first round choices and RAND means the first round choices are drawn from a uniform distribution. UTIL means utilities computed using social preferences are used for attraction updating and PAYOFF means pure payoffs in the game are used for attraction updating in the learning model.

Model	Parameter Estimates				LL	AIC	BIC	Pseudo-R^2
SP-UTIL	μ_λ	1.7891	σ_λ	0.3586	-3885	-3889	-3896	0.2529
	μ_{N1}	0.0335	σ_{N1}	5.6385				
SP-PAYOFF	μ_λ	1.4250	σ_λ	0.3169	-3888	-3892	-3898	0.2523
	μ_{N1}	0.2601	σ_{N1}	4.2001				
REL-UTIL	ν_λ	1.6456	σ_λ	0.2568	-3962	-3966	-3972	0.2381
	ν_{N1}	0.2647	σ_{N1}	6.6453				
REL-PAYOFF	ν_λ	1.4277	σ_λ	0.3255	-3956	-3960	-3966	0.2394
	ν_{N1}	0.0850	σ_{N1}	4.0878				
RAND‡					-6759	-6759	-6759	0.0000

‡ RAND is the random choice model. According to this model any choice has a probability of $\frac{1}{21}$ of being chosen.

REL learning model. Table 3.11 reports two-sided p-values computed using the Vuong's test for each pair of the four models. SP-UTIL and SP-PAYOFF perform equally well in explaining contributions. They outperform RAND-UTIL and RAND-PAYOFF models. RAND-UTIL and RAND-PAYOFF are indistinguishable from each other in terms of their performance. The results highlight the fact that initial contributions determined by social preferences combined with high inertia associated with REL learning model explain the aggregate differences in contributions across cooperative types. This result stands in contrast to the assumption in the existing literature that individuals start with random contributions and then move towards their equilibrium contributions based on learning with utilities derived from social preferences (Cooper & Stockman, 2002; Janssen & Ahn, 2006; Wendel & Oppenheimer, 2010; Arifovic & Ledyard, 2012). No significant difference between the fit of SP-PAYOFF and SP-UTIL indicates that once initial round contributions are determined by the social preferences, there is no additional information to be gained about the

Table 3.11: Disentangling the role of social preferences and learning in repeated network public goods games. REL is used to model learning. SP means social preferences determine first round choices and RAND means first round choices are drawn from a uniform distribution. UTIL means utilities computed using social preferences are used for attraction updating and PAYOFF means pure payoffs in the game are used for attraction updating in the learning model. For each pair of learning models, a two sided p-value favoring the learning model in the row is reported. The two-sided p-values are computed using the Vuong Test for Non-Nested Models.

	SP-PAYOFF	RAND-UTIL	RAND-PAYOFF	RAND
SP-UTIL	0.7493	0.0000	0.0006	0.0000
SP-PAYOFF		0.0001	0.0004	0.0000
RAND-UTIL			0.4283	0.0000
RAND-PAYOFF				0.0000

path of the contributions by using social preferences based utilities to update the attractions of strategies rather than using simple payoffs within the a repeated game. In other words, social preferences matter insofar they determine the first round contributions and then individuals behave solely based on payoff-based learning.

Identical to the findings in the context of public goods games in fixed groups in Chapter 2, my analysis here shows that choices in the repeated network public goods games are best described by social preferences affecting the choice of first round contributions and then subsequent contributions based on payoff based average reinforcement learning.

3.3.7 Effect of Finite and Indefinite Repetition on Learning

My analysis thus far has shown that choices in the repeated network public goods games are best described by social preferences affecting the choice of first round contributions and then subsequent contributions based on payoff-based reinforcement learning. In this subsection, I investigate if the type of repetition has any effect on the learning parameters of REL model.

Table 3.12 reports estimation results using only repeated game data from the R1 and

Table 3.12: Effects of the type of repetition on the parameters of REL learning model in network public goods games

	λ	$N1$
Constant	1.67***	0.028**
	(0.18)	(0.013)
Infinitely Repeated	0.18	0.028
	(0.24)	(0.049)
σ	0.37***	5.68***
	(0.095)	(0.28)
N	150	
LL	-3883	

Notes: $*p < 0.1, **p < 0.05, ***p < 0.01$. Coefficients are transformed back to the original scale. A treatment effect is the (partial) effect of moving from the baseline treatment to the corresponding treatment on the median parameter value. Standard errors are calculated using the *sandwich estimator* and treating all of each subjects choices as a single *super observation*. The entries for σ are the standard deviations of the untransformed normal distributions of the random coefficients.

R2-tasks. It contains the estimates for the two parameters of REL model, λ and $N1$ for the model that contains an intercept and a dummy for the finitely repeated game. All values are on the original scale. There are no significant effects of finite repetition on the estimated parameters of the REL learning model. Thus learning is identical in both finitely and indefinitely repeated games. This result explains the similar observed contribution patterns across finitely and indefinitely repeated games in network public goods games.

3.3.8 Section Summary

In this section, I report results from an experimental network public goods games. The goal of conducting this experiment is to identify if contributions and behavior are affected in any significant way due to the presence of a network structure. When I compare contributions in the network public goods game with that of the data from Partners' treatment which does not involve overlapping groups from Chapter 2, I find no significant difference indicating presence of a network does not affect contribution choices. I also find that distribution of cooperative types in network public goods games are identical to that of the distribution of types in non-network public goods games. I find that the estimated distribution of social preference parameters is reasonably successful in explaining the aggregate changes in the distribution of cooperative types (free riders, conditional cooperators, and full cooperators) when prices change. The model is able to predict an increasing proportion of full cooperators when the price of cooperation gets smaller. Finally, identical to the non-network case, I find that contribution choices in repeated network public goods games are best described by social preferences of participants affecting the choice of first round contributions and then subsequent contributions based on payoff-based reinforcement learning. I find that learning is identical across finitely and indefinitely repeated network public goods games.

3.4 Cooperation in Network Public Goods Games: An Empirical Agent-Based Model

Having identified the roles of social preferences and learning in network public goods games in the previous section, I now turn to deploying these empirical insights in an agent-based model. The "empirical" agent-based model thus created will be deployed in thought-experiments to generate new insights about contribution dynamics in networked communities. This section has four subsections. In the first subsection, I describe the details of the empirical agent-based model. Second subsection explores the effect of network size on the contribution levels in a network public goods game using the empirical agent-based model. In the third subsection, I examine the role of degree heterogeneity on the contribution levels. Fourth subsection investigates if different activation regimes, network density, neighborhood sizes, and the average path length make a difference. Finally, in the fifth subsection I show that subsidizing cooperation in the first round can sustain higher levels of cooperation in the later rounds of a repeated game.

Table 3.13: Means and standard deviations of the underlying normals of the distributions of social preference and learning parameters

	β	γ	λ	$N1$
μ	2.7875	2.1843	0.3542	-1.3468
σ	5.0538	1.8585	0.3169	4.2001

3.4.1 Empirical Agent-Based Model

The empirical agent-based is formulated after the estimation results in Section 3.3. Given a network size, the required number of agents are randomly located on the the nodes of the network. For each agent, the social preference parameters and the learning parameters are drawn from the corresponding estimated distributions. The minimal model estimated in Table 3.4 is used to specify social preferences of individuals. The learning parameters are specified using the estimates for SP-PAYOFF model in Table 3.10. Note that social preference parameters and learning parameters were specified and estimated as log-normal distributions. Table 3.13 presents the means and standard deviations of the underlying normals of the log-normal distributions of the parameters. For each agent, each parameter value is assigned using a random draw taken from the corresponding log-normal distribution. In the simulations of repeated games, in the first round each agent makes a contribution decision based on its social preferences and random belief. In the later rounds, the agents use REL learning model with its parameters to make contribution decisions. The activation scheme is synchronous meaning that in each round each agents make their contribution decisions simultaneously. I will relax this assumption in the third subsection. In all simulations MPCR is fixed at 0.5 unless otherwise specified. And, the initial endowment is 20 tokens.

3.4.2 Effect of Network Size

Results in Table 3.3 show that the presence of a network structure alone does not affect contributions in public goods games. However, this can be due to the size of the network. The network considered in my experiments is rather very small and consists of only 15 nodes. The empirical agent-based model allows me to experiment with arbitrarily large network sizes to see if the size of the network matters for the network effects on contributions to pan out. To this end, I vary the size of the network systematically from 10 to 200. The structure of the network is identical to that of the network considered in my experiments. The players are placed on a circle and are connected to the players on the left and right. Thus, the degree of each player is 2. By keeping the structure of the network and neighborhood sizes identical to that of the networks in my experiment, I can study the effect of the size of the network on contributions. Figure 3.6 depicts average contribution levels achieved in the 20th round for each network size that was considered. I ran the model for 20 rounds since contribution levels were observed to be stabilized by then in my investigations. The error bars stand for 95% confidence intervals. The confidence intervals will shrink as network size increase. It is clear that the network structure does not influence contributions significantly.

3.4.3 Effect of Degree Heterogeneity

A key insight from the literature on the effect of a network that constrains the interaction among individuals on cooperation is that a network with degree heterogeneity can promote higher levels of cooperation than that of a network without degree heterogeneity. In this section, I evaluate this insight using the empirical agent-based model that is outlined earlier. I study the effect of degree heterogeneity by considering different types of networks. In my computational experiments, four types of network topologies are examined: regular graph, small world graph, scale-free graph, and complete graph. In a regular graph each node is connected to exactly the same number of other nodes. The graph that is considered in the experiment in Section 3.3 is an example of a regular graph. Regular graphs have no degree

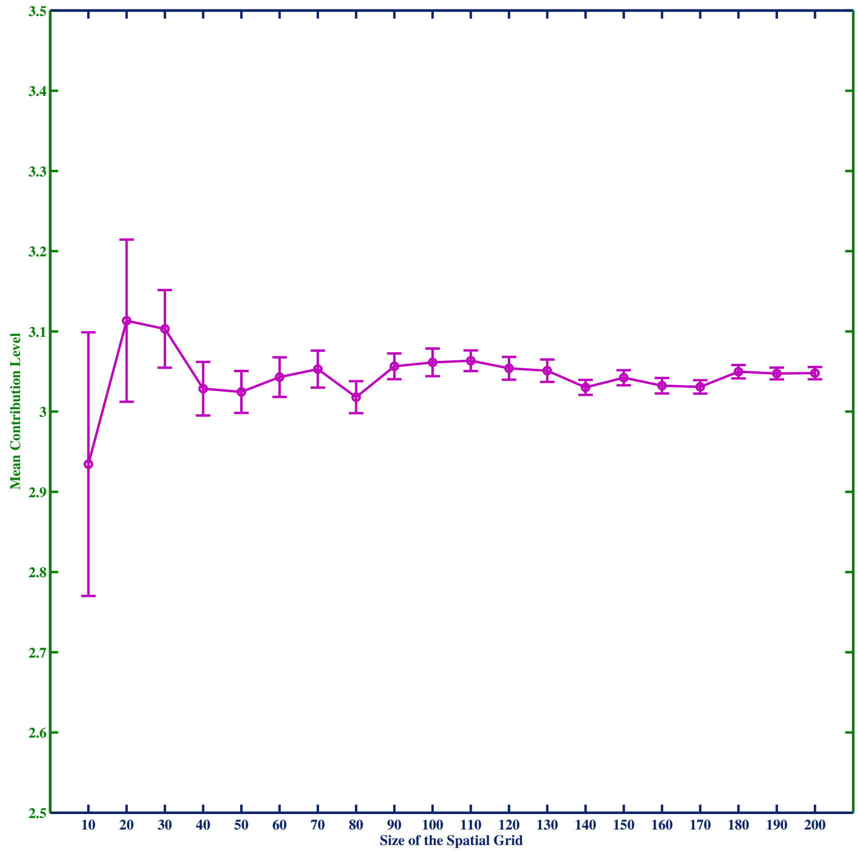

Figure 3.6: Simulated mean contribution levels in the 20^{th} round of a repeated network public goods game for different sizes of the network. Error bars represent 95% confidence intervals.

heterogeneity since all nodes have exactly the same degree. In the simulations, I considered a regular graph of size 30 and in which the degree of each node is 8 or neighborhood radius is $d = 4$. Small world graphs are obtained by randomly rewiring connections in a regular graph (Watts & Strogatz, 1998; Watts, 2003). This makes the average path length in these networks small as observed in real world networks. A small world graph has two parameters that characterize how it is created: neighborhood radius (k) and the rewiring probability (p). Given k and p a small world graph is created as follows: first a regular graph in which all nodes have a degree $2k$ (or equivalently a neighborhood radius of k) is created and then

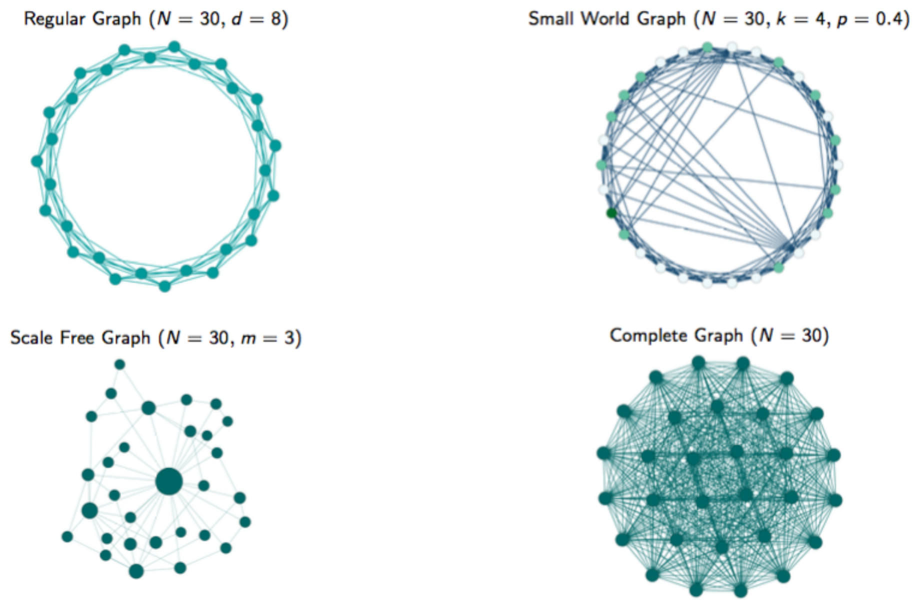

Figure 3.7: Networks with varying degree heterogeneity that are considered in simulations.

each edge in the network is rewired with a probability of p. In my simulations, I considered two small world networks of size 30: first one with parameters $k = 4, p = 0.3$, and second one with parameters $k = 4, p = 0.7$. A small world network also does not have degree heterogeneity like a regular graph. However, it has a relatively small average path length. A scale free network (Barabási & Albert, 1999; Barabási, 2009) is a prototypical example of a network with degree heterogeneity. To be more precise, the degree distribution in a scale-free network follows a power law. A scale free network is obtained via the so-called "preferential attachment" algorithm (Barabási & Albert, 1999). The idea is that a scale free network grows in steps where in each step a new node is added to the network. The new node is attached to an existing node in the network with a probability that is proportional to the degree of the existing node. This is what is meant by preferential attachment. Nodes with a lot of neighbors will keep adding more and more neighbors. There is one parameter in this algorithm called the number of new edges (d) introduced with a new node. In my simulations, I considered two scale networks of size 30 with d= 3 in one case and d = 7 in another case. Finally, in a complete graph every node is connected to every other node in the network. This is a limiting case of a network and produces a scenario equivalent to that of typical Partners' treatment in public goods games. In the simulations, I considered a complete graph of size 30. Note that the degree of each node in a complete graph is the same, and in the network of size 30 it would be 29. Figure 3.7 depicts the networks that were used in the simulations.

Figure 3.8 plots the average contribution in each period in a 20 period networked public goods game for the network topologies under consideration. I find that the differences in stable contributions among different topologies are almost negligible. These results are in contrast with the findings from previous studies involving agent-based models (F. C. Santos & Pacheco, 2005; F. Santos et al., 2008).

3.4.4 Effects of Activation Schemes, Neighborhood Size, Average Path Length, and Network Density

An implicit but a very crucial assumption in the previous simulations is that activation scheme is synchronous. In other words, agents make their contribution decisions in unison. In a social system, often such an assumption is unrealistic. An asynchronous decision making scheme is more realistic. Furthermore, previous research shows that strikingly different results are obtained from agent-based models when different activation schemes are considered (Huberman & Glance, 1993; Axtell, 2000a). Therefore, I examine the results obtained in the previous section in light of different activations schemes.To understand the effects of asynchronous activation regimes where decision making of agents is temporally minimally correlated, I consider two activation regimes: random activation and uniform activation. In these schemes, a period refers to the number of activations equal to the size of the population. In random activation, every player is activated on average once every period, whereas in uniform activation each player is activated exactly once but the order of activation of agents is random. For each activation scheme, I also consider if the

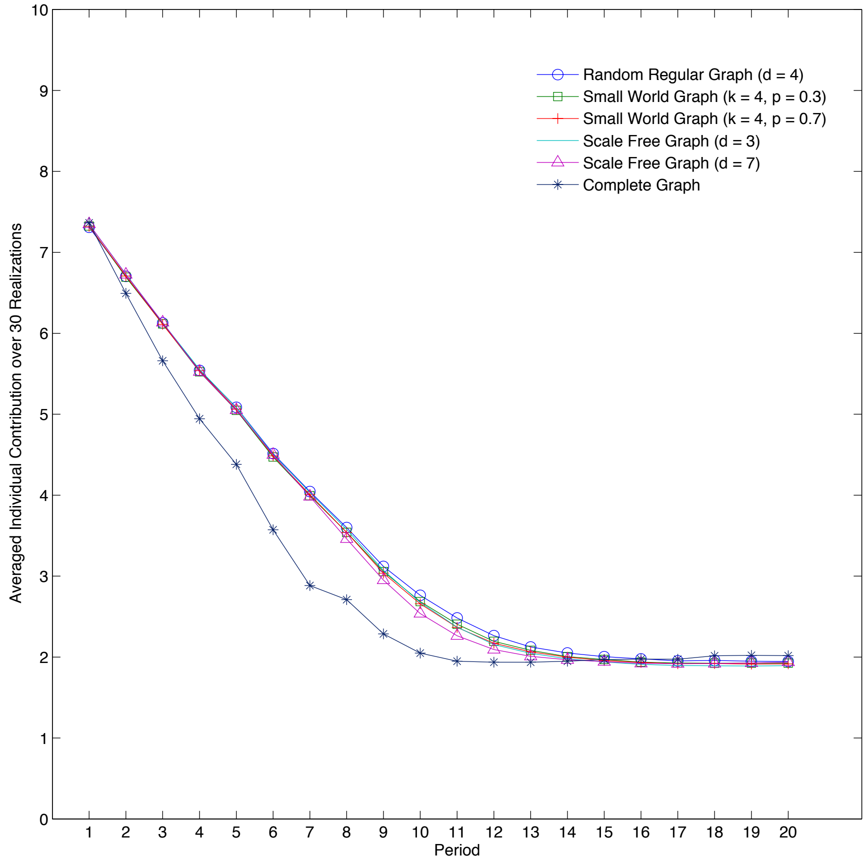

Figure 3.8: Simulated mean contributions over the rounds of network public goods games. All networks are of size 30.

variation in neighborhood sizes and average path lengths leads to different contribution dynamics. I study neighborhood size effects, by varying the size of the radius (d) of regular networks (note that a radius size of d means neighborhood size of $2d$). I study the effect of the average path length by varying the probability of rewiring (p) in small world networks with $k = 2$. Finally, I vary the number of edges introduced with a new node in the preferential attachment algorithm that generates a scale free network. A higher number of edges introduced in each step lead to higher density of the scale free network.

Figure 3.9 presents average contribution levels achieved in the 20th round at different sizes of neighborhoods in regular networks for three different activation regimes. Again, I

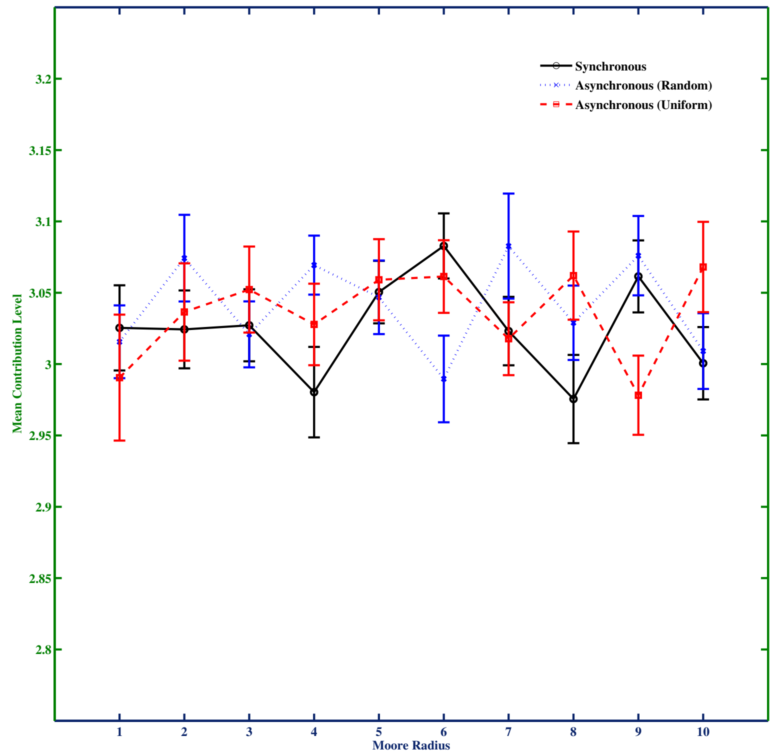

Figure 3.9: Simulated mean contribution levels in the 20^{th} round for different radii of a regular network of 30 nodes for three different activation regimes, viz. synchronous activation and the two asynchronous activation regimes: random activation and uniform activation. Error bars represent 95% confidence intervals.

chose 20th round contributions for reporting since in my simulations I found that contributions stabilized by 20th round of the repeated game. There are no discernible differences in contributions levels either across activation regimes or across different neighborhood sizes (Mann-Whitney Tests, $p > 0.10$ in all cases). Figure 3.10 presents average contribution levels achieved in the 20th round at different probabilities of rewiring in small world networks for three different activation regimes. There are no discernible differences in contributions levels either across activation regimes or across different rewiring probabilities that affect the average path length of the small world network (Mann-Whitney Tests, $p > 0.10$ in all

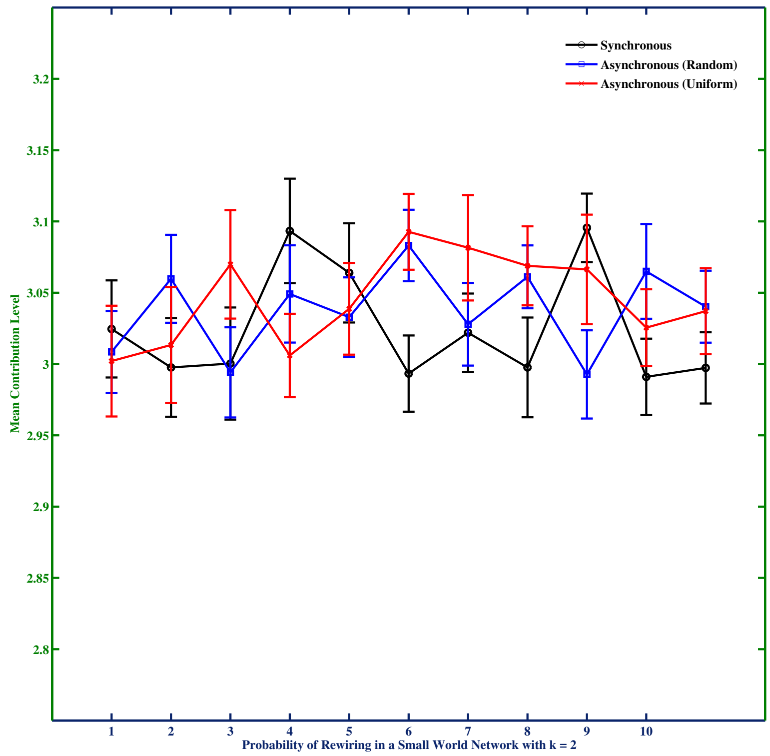

Figure 3.10: Simulated mean contribution levels in the 20^{th} round for different rewiring probabilities for a small world network with k = 2 for three different activation regimes. Probability of rewiring is varied from 0 to 1 in increments of 0.1. Error bars represent 95% confidence intervals.

cases). Finally, Figure 3.11 presents average contribution levels achieved in the 20th round for different number of new edges introduced in the generation of scale-free networks for three different activation regimes. There are no discernible differences in contributions levels either across activation regimes or across different densities of scale-free networks (Mann-Whitney Tests, $p > 0.10$ in all cases). These results show neither activation regimes nor structural details of networks make any significant difference in contributions of a network public goods game.

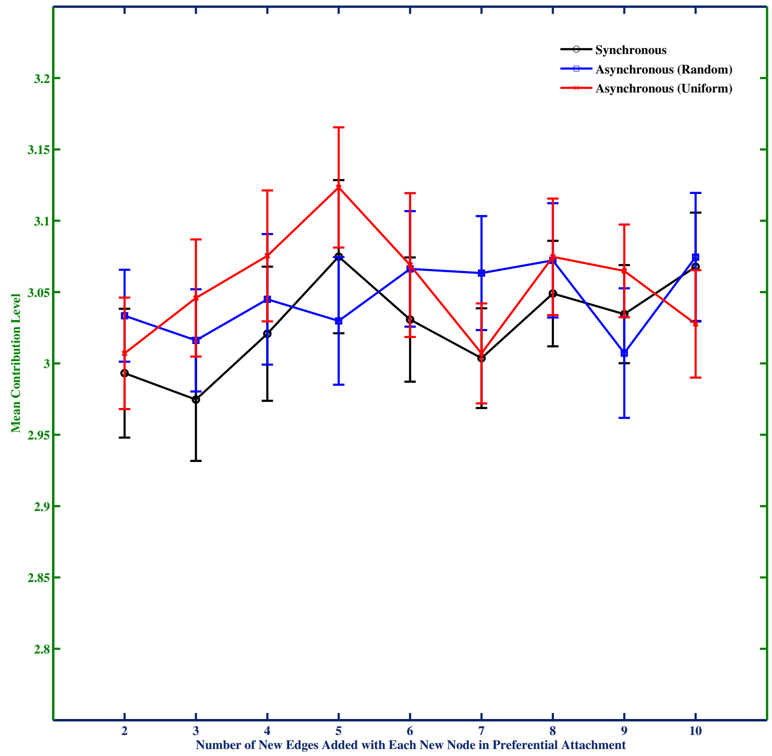

Figure 3.11: Simulated mean contribution levels in the 20^{th} round for different number of new edges introduced with a new node in the preferential attachment algorithm that produces a scale free network for three different activation regimes. Error bars represent 95% confidence intervals.

3.4.5 Subsidizing Cooperation in the First Round Can Sustain Higher Levels of Cooperation

In Section 3.3, I showed that first round contributions are in line with that of social preferences and in later rounds decision making is in line with reinforcement learning. This means that by reducing the cost of cooperation in the first round, one can make individual contributions higher in the first round. This is because, at low prices the proportion of free riders decreases and full cooperators increases. Given my finding that individuals learn using a

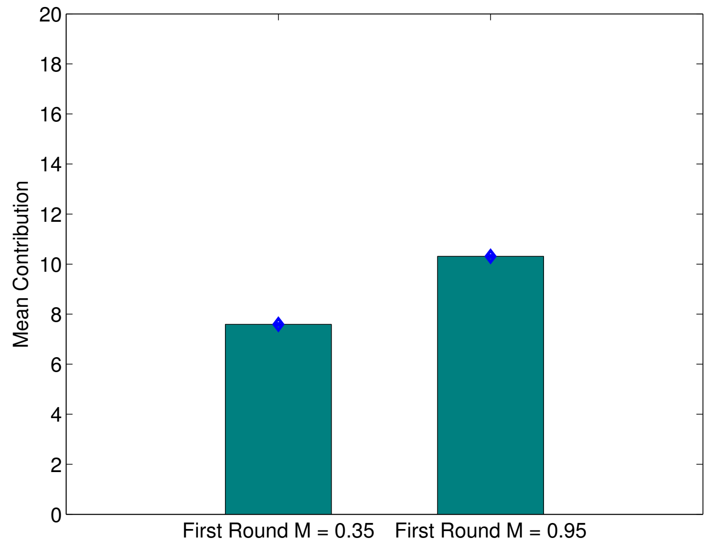

Figure 3.12: Mean contribution levels in the $2-10$ rounds of a repeated network public goods games for two (simulated) cases where the first round MPCR is varied. MPCR in later rounds is fixed at 0.5. Error bars represent 95% confidence intervals.

high inertia learning rule like that of reinforcement, the higher level of contributions in the first round would mean that higher levels of contributions will be sustained in later rounds too. I examine this by using the empirical agent-based model. I consider two treatments. I consider a regular network as that of the network considered in the experiment in Section 3.3. Each agent is connected to exactly two other agents. In the first treatment, I fix the first round MPCR at 0.35 and the MPCR in later rounds is 0.5. In the second treatment, I fix the first round MPCR at 0.95 and in the later rounds the MPCR is 0.5. I then simulate a repeated network public goods game for 10 rounds. Figure 3.12 reports mean contributions in 2-10 rounds. These mean contributions are significantly different (Mann-Whitney Test, $p < 0.0001$). Thus, understanding the determinants of behavior can help one to think about empirical mechanisms that can promote higher levels of contributions.

3.5 Conclusions

In this Chapter, I empirically investigated individual decision making in network public goods games using an incentivized experiment. I find no significant differences across contributions in network public goods games and public goods games with fixed non-overlapping groups indicating that presence of a network does not affect contribution choices. I also find that the distribution of cooperative types (free riders, conditional cooperators, and full cooperators) in network public goods games are identical to that of the distribution of types in non-network public goods games. I find that the estimated distribution of social preference parameters is reasonably successful in explaining the aggregate changes in the distribution of cooperative types when prices change. The model is able to predict an increasing proportion of full cooperators when the price of cooperation gets smaller and an increasing proportion of free riders when the price of cooperation gets larger. Finally, identical to the non-network case, I find that contribution choices in repeated network public goods games are best described by social preferences of participants affecting the choice of first round contributions and then subsequent contributions based on payoff-based reinforcement learning. I find that learning is identical across finitely and indefinitely repeated network public goods games. I deployed these empirical insights in an agent-based model. The "empirical" agent-based model thus created was deployed in thought-experiments to generate new insights about contribution dynamics in networked communities. Simulations with the empirical agent based model showed that network size, degree heterogeneity, average path length, and network density do not have any significant effect on contributions in repeated public goods games. These results stand in contrast with the results from agent-based simulations based on behavioral rules inspired from microbiology which argue that presence of a network structure significantly enhances cooperation in social dilemmas including public goods games. The empirical agent-based model shows that subsidizing cooperation in the first round can sustain higher levels of cooperation over later rounds of a repeated game. This is due to the way social preferences interact with the price of cooperation and the high inertia associated with reinforcement learning. These insights illustrate that empirical understanding of determinants of behavior in social dilemmas is instrumental in identifying effective mechanisms that can promote cooperation using agent-based models.

Agent-based models have gained a significant appreciation in explaining the relation between macro level outcomes and micro level dynamics in complex social systems. Agent-based models provide a rich representational freedom in terms of realistic or intuitive behavioral specification of agents, heterogeneity in agents' behavior, structure of interaction among agents, and the timing of interaction. This flexibility allows researchers to study real world social phenomena without making too many assumptions about behavior or interaction patterns so as to make models mathematically tractable. Because of the rich representational freedom agent-based models offer, they have been deployed to model a wide variety of real world social phenomena including neighborhood segregation (Schelling, 1998), cooperation in groups (Axelrod, 1997), cooperation in structured populations (Nowak & May, 1992; F. Santos et al., 2008), opinion dynamics (Deffuant, Neau, Amblard, & Weisbuch, 2000), organizational problem solving (Lazer & Friedman, 2007), evolution of culture (Axelrod, 1997; Epstein & Axtell, 1996), and emergence of conflict due to environmental drivers (Kennedy, Cotla, Gulden, Coletti, & Cioffi-Revilla, 2014). While representational freedom offered by agent-based models is one of their key strengths, it also opens up the possibility of arbitrariness in how one decides on which behavioral representations to use at the agent level. Earlier literature on agent-based models mainly concerned itself with providing alternative descriptive explanations of how complex macro level regularities emerge from seemingly intuitive and/or simple behavioral rules (Epstein & Axtell, 1996; Axelrod, 1997; Macy & Willer, 2002; Schelling, 1998). One of the main emphases was on the exposition that individual level rationality, which assumes unrealistic cognitive and computational capabilities of individuals, as theorized in standard economics is not necessary to generate rational aggregate outcomes. For example, Gode and Sunder (1993) showed that transaction prices converge to the equilibrium price in a continuous double auction market even with "zero-intelligent" agents that make random bids and offers. As long as zero intelligent agents do not trade for losses, the allocative efficiency in these markets reaches close to 100%. Earlier agent-based models were also not averse to the multiplicity of micro level explanations of a given macro level phenomena, in fact it was encouraged. An inadvertent side effect of this practice was that the methodology of agent-based models in its earlier stages was *not* suited to do comparative static analysis or to quantify the effect of a given intervention mechanism on a macro level outcome variable of interest or to forecast future outcomes.

Recent research in agent-based models set out to address this particular issue by using the so called empirical agent-based models (Janssen & Ostrom, 2006; Janssen & Ahn, 2006; Robinson et al., 2007; Heckbert, Baynes, & Reeson, 2010; Wunder et al., 2013; Smajgl & Barreteau, 2014). Empirical agent-based models typically use quantitative or qualitative data to derive behavioral specification of agents and/or their interaction patterns, for example as in Wunder et al. (2013). In some cases, empirical information is used to falsify and test a given agent-based model. This Chapter contributes to the methodology of empirical agent-based modeling. It addresses two challenges in empirical agent-based modeling paradigm: identification of empirically relevant determinants of behavior in a given domain and the implementation of behavioral heterogeneity. First, it demonstrates that the challenge of identifying pertinent behavioral determinants can be addressed using incentivized experiments. Second, it illustrate that econometric methods that allow for observed and unobserved parameter heterogeneity across individuals like that of Mixed Logit methods (Train, 2009) provide a systematic foundation to estimate the behavioral heterogeneity from empirical and/or experimental data and then to incorporate it into agent-based models. The insights in this Chapter have potential to standardize how agent-based modelers identify empirically relevant behavioral determinants and implement heterogeneity in the behavioral determinants in agent-based models.

3.6 Appendix

3.6.1 Experiment Instructions

.

Instructions

Welcome and thank you for participating in todays economic experiment. Please put away all your belongings and turn off your cell phones. You are not allowed to talk to any other participant during the experiment. If you have any questions, please raise your hand. We will come to you and answer your questions in private. The experiment will be run entirely on the computer and all interactions between yourself and others will take place via the computer terminal.

You have earned $5 just for showing up on time. This is yours to keep. In addition, depending upon the decisions you make, the decisions others make and random choice, you can earn more money. These instructions describe in detail the experiment and tasks you are asked to complete.

During the experiment, your earnings will be described in terms of tokens. At the end of the experiment, the total number of tokens you have earned will be converted to money at the following rate:

20 tokens = $1 (1 token = 5 cents)

The experiment consists of four tasks. You will receive instructions for each task prior to making decisions for that task. Your total earnings from the experiment will be the sum of your earnings in each task. At the end of the session, the total number of tokens earned across all four tasks will be converted to money and paid to you privately in cash, along with the $5 show-up fee.

Your Neighborhood

You will be placed in a network with 14 other participants as shown in the Figure 3.13.. The placement of participants in the network is random, and participants do not know who is connected to whom.

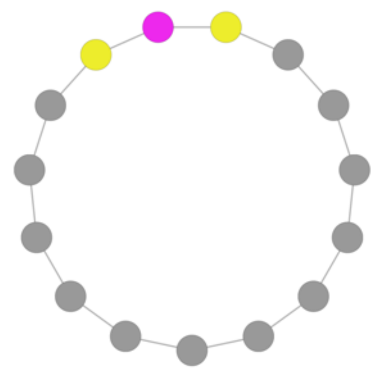

Figure 3.13: Network in the experiment

Each participant is placed at one position on the network and is connected to exactly two other participants. This placement and connection are fixed throughout each of the four tasks. You and the two other participants who are connected to you in the network define your neighborhood. Your neighborhood is always made up of the one participant to the right of you and the one participant to the left of you. This is illustrated in the Figure 1 for example, if you are placed in the position of the circle that is highlighted in pink then your neighbors are highlighted in yellow.

Neighborhoods overlap. Your neighbors are also members of other neighborhoods, and you are a member of other neighborhoods. This is illustrated in Figure 3.14. Your neighborhood is shown by the shaded area around the yellow, pink and yellow positions (numbered 1, 2, 3) in the network. The neighborhood of your neighbor to the right is shown by the shaded area around the pink, yellow and blue positions (numbered 2, 3, 4). In other words, you are in your neighborhood and also in the neighborhood of your neighbor to the right. Likewise, you are in the neighborhood of your neighbor to the left.

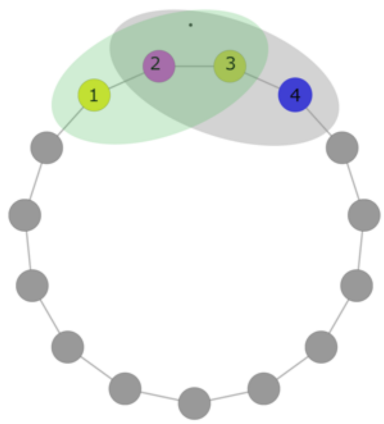

Figure 3.14: Neighborhoods on the network

The Decision Situation

Each participant is provided with 20 tokens and must decide how to allocate the tokens between a private account and a group project. You can choose to put none, all or some of your tokens into the group project. The tokens you choose not to contribute to the group project will remain in your private account. Everyone makes the same decision.

For each token put in the private account you earn exactly one token. You are the only one who earns tokens from your private account.

What you earn from the group project depends on the total number of tokens that you and your neighbors contribute to the group project. The more each member of the neighborhood contributes to the group project, the more each member earns. Remember that your neighborhood includes you and two other participants.

Your earnings from the group project are best explained by a number of examples.

Example 1: Suppose that you decided to contribute no tokens to the group project but the 2 other members of your neighborhood contribute a total of 36 tokens. Then your earnings from the group project would be 36 tokens $\times$ 0.4 = 14.4 tokens. Everyone else in your group would also earn 14.4 tokens.

Example 2: Suppose that you contribute 15 tokens to the group project and the 2 other members of your neighborhood invest a total of 36 tokens. This makes a group total of 51 tokens. Your earnings from the group project would be 51 tokens $\times$ 0.4 = 20.4 tokens. The other 2 members of the group would also earn 20.4 tokens.

Example 3: Suppose that you contribute 20 tokens in the group project but the other 2 members in your neighborhood invest nothing. Then you, and everyone else in the group, would earn from the group project 8 tokens (20 tokens $\times$ 0.4 = 8 tokens).

As you can see, every token contributed to the group project earns 0.4 tokens for every member of the neighborhood, not just the participant who puts it there. It does not matter who contributes tokens to the group project. Everyone will get a return from every token contributed therewhether they contributed tokens in the group project or not.

Your total earnings from the private account and group project will be:

Your total earnings = 20 - your tokens contributed to the group project + 0.4 $\times$ sum of tokens contributed to the group project by all members of your neighborhood

You will now complete some questions to make sure everyone understands how earnings are calculated.

Instructions for Task 1 – A-Task

In the A-task, you will make a choice for the decision situation described earlier. You will have 20 tokens and must decide how many to put into your private account and the group project. You will be randomly assigned to a 3-participant neighborhood.

Each participant has two decisions in this task: make an unconditional contribution and complete a contribution table. Details about these two decisions are as follows.

Unconditional Contribution: In this decision, you must decide how many of the 20 tokens you would like to put in the group project. You will make your decision on a screen such as the following.

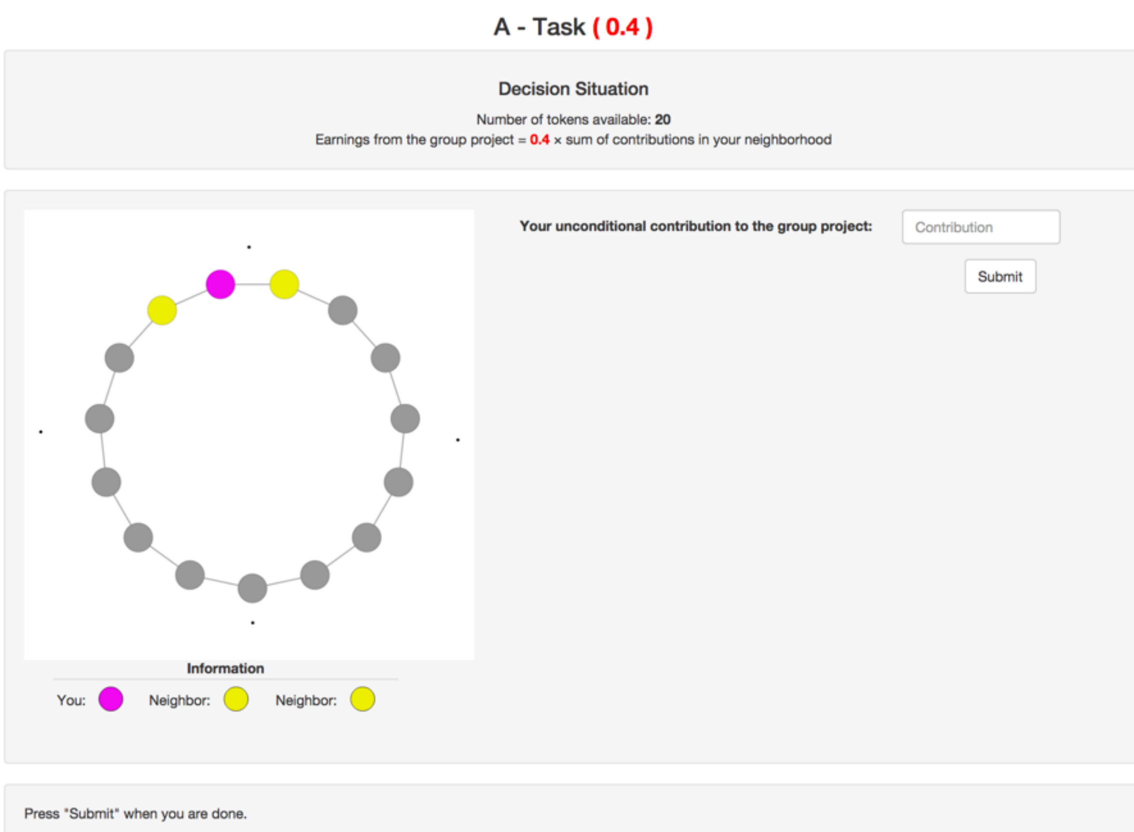

Figure 3.15: Decision screen for unconditional contribution choice

Contribution Table: In this decision, you must decide how many tokens you would like contribute to the group project for each possible average contribution of your neighbors (e.g. 0, 1, 2,..., 20). That is, if your neighbors contributed 0 tokens on average, how much would you contribute? If they contributed 1 token on average, how much would you contribute? If they contributed 2 tokens on average, how much would you contribute? And so on, up to 20 tokens on average.

This means that in total you have to give 21 responses. You will make your decisions on a screen such as the following

Once each participant has made the unconditional decision and completed the contribution table, the computer will randomly determine if the unconditional contribution or the contribution table will be used to determine your earnings. In each neighborhood, one of the three participants is randomly chosen to have the contribution table count to calculate earnings. For the other two participants in the neighborhood, the unconditional contribution counts to calculate earnings. How is this done? If the participant is chosen to have his contribution table count for earnings, first the unconditional contributions to the group project of his neighbors are averaged and rounded to the nearest integer (e.g. 0,1,2,...,20). Then, the contribution table of the participant is used to determine how many tokens the participant contributes to the group project. The number of tokens contributed is the amount he specified for the average contribution of his neighbors.

So, if the average contribution of his neighbors is 16 tokens, and he specified 10 tokens if the average contribution of his neighbors is 16, the total contributed to the group project by everyone in the neighborhood would be 42 ($16\times2+10$) tokens.

You will not know if advance which decision, the unconditional contribution or the contribution table, will count to determine your earnings, so you should make each decision as though it will count for your earnings.

The following examples should help make this procedure clear.

Example 1: Suppose the contribution table was randomly chosen to count for you. This means that the decisions you made in the contribution table determine your earnings.

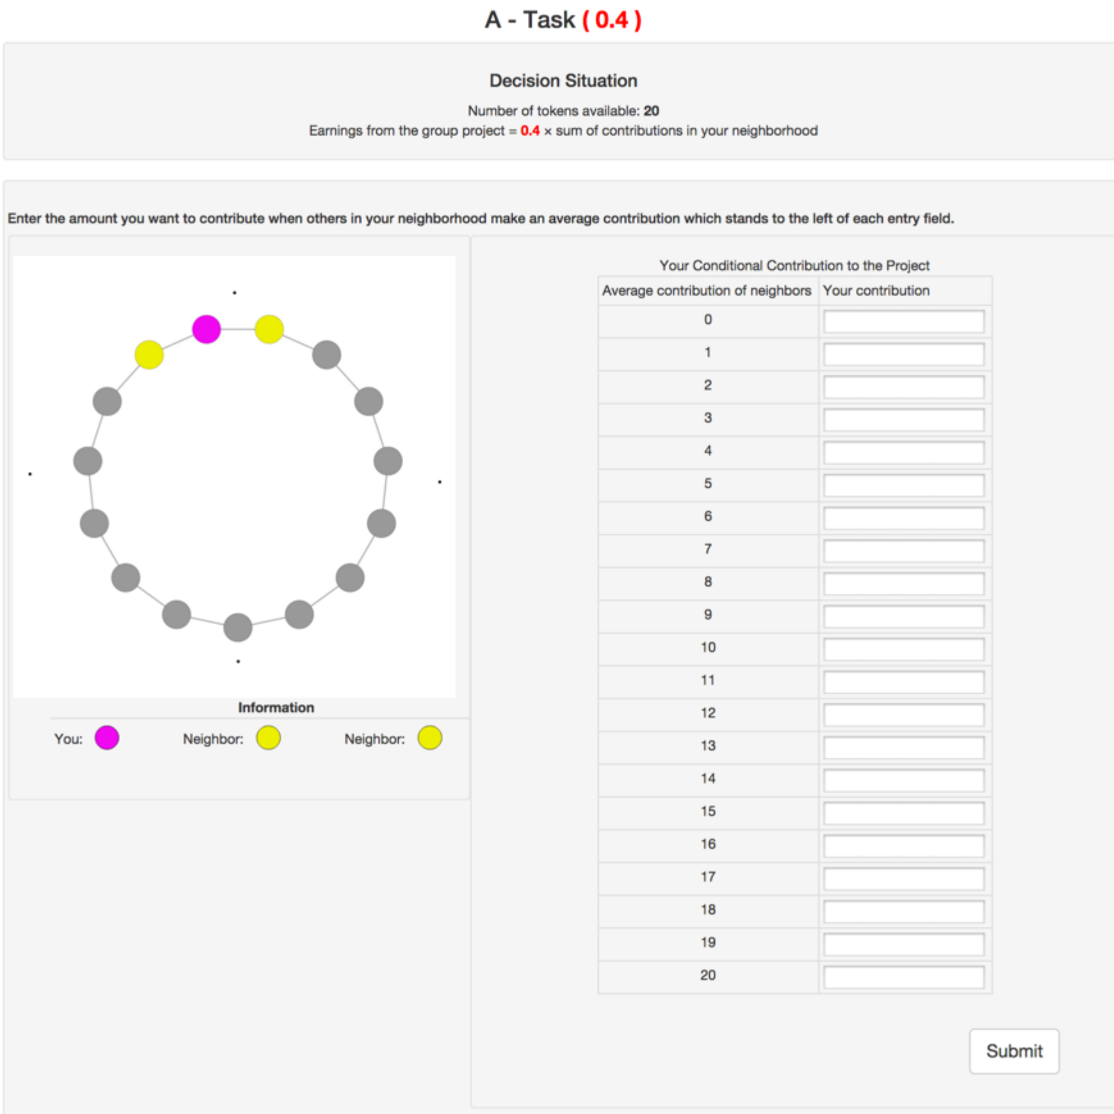

Figure 3.16: Decision screen for conditional contribution table

For the other two neighbors their unconditional contributions determine their earnings. Suppose that the total contributions of the other two neighbors are 26 tokens, and the average contribution 13 tokens (26 tokens/2). In your contribution table, suppose you chose to contribute 4 tokens if the average contribution of your neighbors is 13, then your earnings for Task 1 would be: $20 - 4 + 0.4 \times (4 + 26) = 28$. If instead you chose to contribute 14 if the average contribution of neighbors is 13, your payoff would be: 20 - 14 + 0.4 $\times$ (14 + 26) = 22.

Example 2: Suppose the unconditional contribution was randomly chosen to count for your earnings. Also, suppose that the unconditional contribution of the neighbor who was not selected for the contribution table to count is 12. If your unconditional contribution is 20, then the average unconditional contribution is 16 tokens ((20 + 12)/2). If the neighbor selected to have his contribution table count chose 18 tokens if the average contribution of his neighbors is 16, then your earnings are: 20 - 20 + 0.4 $\times$ (20 + 12 + 18) = 20.

Are there any questions before we begin?

Instructions for Task 2 – B-Task

In the B-Task, you will be randomly assigned to a 3-participant neighborhood as described earlier. Your neighbors in Task 2 may be different from your neighbors in Task 1, however, you will remain with the same neighbors for all decisions you make in Task 2.

The B-Task lasts for several rounds. The number of rounds is randomly determined. In each round, you face the basic decision situation described at the beginning of the experiment. After each round, there is an 85% probability that there will be one more round. So, for instance, if you are in round 2, the probability there will be a third round is 85% and if you are in round 9, the probability there will be another round is also 85%. How this works is as follows. After each round, the computer will randomly draw a number between 1 and 100 (e.g. 1, 2, 3,..., 100), where each number is equally likely to be chosen. If the chosen number is 85 or lower, there will be another round. If the chosen number is 86 or above, there will be no additional rounds, and the task will end. You will know there is another round if you see the decision screen again and are asked to make a decision. If the task ends, you will get a message saying the task is done. You will not know ahead of time for how many rounds you will make decisions.

In each round, you will be given 20 tokens and must decide how many tokens you would like to contribute to the group project and how many you would like to put in your private account. You will receive earnings only from the group project that involves participants in your neighborhood. Your earnings from your contribution decision in a given round are determined as:

Your total earnings in a round = 20 - your tokens contributed to the group project + 0.4 × sum of tokens contributed to the group project by all members of your neighborhood

You will participate in the decision situation repeatedly with the same neighbors, until it is randomly determined that there are no more rounds.

You will make decisions on a screen such as the following:

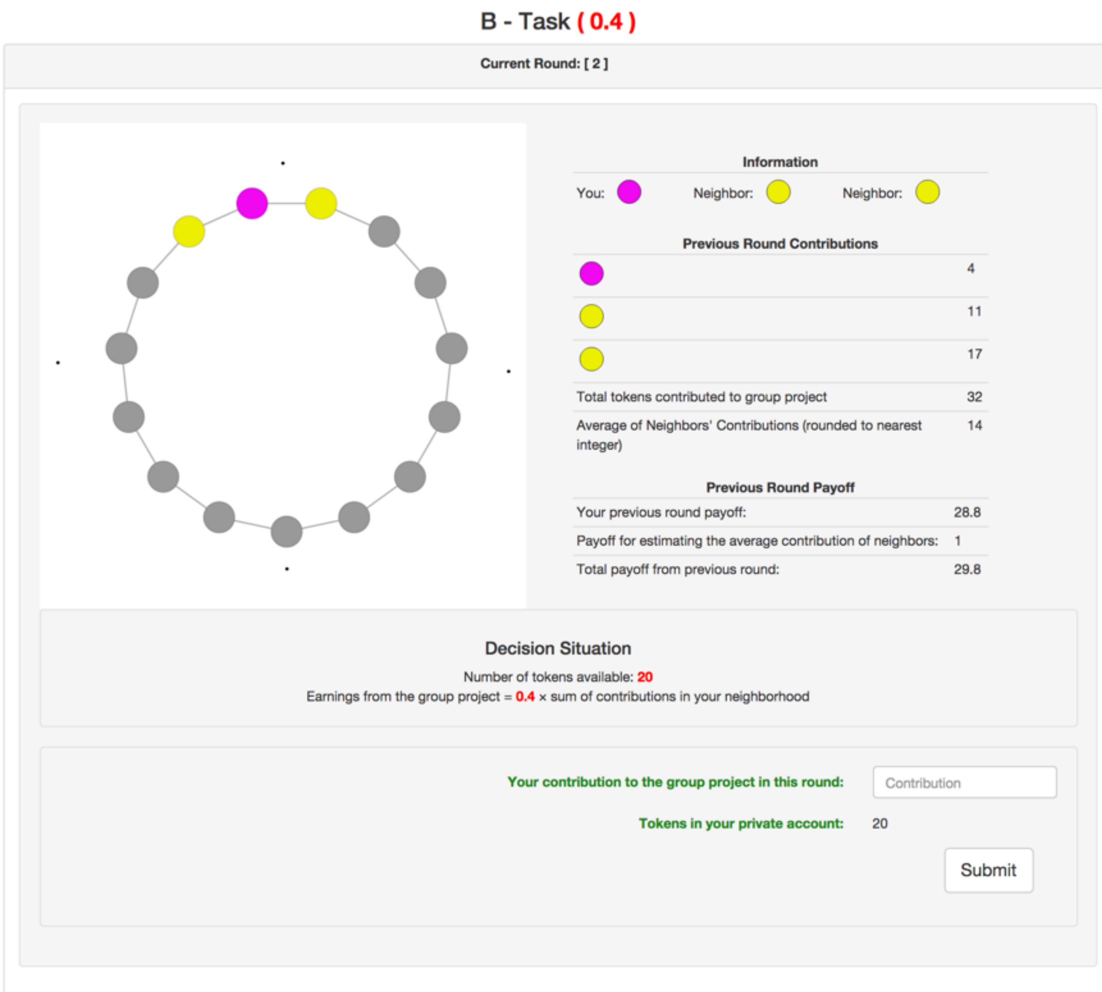

Figure 3.17: Decision Screen for contribution choice

Here is an example to explain how earnings are calculated in each round:

Example 1: Suppose you chose to contribute 12 tokens and your neighbors chose to contribute 30 tokens in total. Your earnings in that round would be: 20 - 12 + 0.4 × (12

+ 30) = 24.8 tokens.

In each round, after you decide how much to contribute to the group project, you will be asked to guess the average contribution to the project (rounded to the nearest integer) of your two neighbors. You will receive tokens for the accuracy of your estimate. If your guess is exactly equal to the average contribution of your neighbors you will receive 3 tokens in addition to your earnings for that round. If your guess was off by 1 token, you will get 2 additional tokens. If your guess was off by 2 tokens, you will get 1 additional token. And, if your guess was off by 3 or more tokens, you will get 0 additional tokens.

When everybody in your neighborhood has completed the two decisions, you will be shown each of their contributions, the total contributions to the group project, and the average contribution. You will only be informed of the contributions of those in your neighborhood. You will not be informed of contributions of participants in other neighborhoods. You will also be informed of your earnings for the current round.

Once all subjects in the experiment have completed the two decisions and are told their earnings and the contributions of their neighbors in the current round, the computer will randomly draw a number between 1 and 100 to see if everyone plays another round. If there is not another round, the task is done.

Are there any questions before we begin?

Instructions for Task 3 – A-Task

In Task 3, you will make two decisions again as you did in the Task 1 A-Task. The difference between this task and Task 1 is that for each token contributed to the group project you, and the other two neighbors, will get 0.8 tokens back.

Things to remember:

1. You will be randomly placed on the network at the beginning of the task and assigned to a 3-participant neighborhood. Your neighbors in this Task 3 may be different than in the previous two tasks.
2. You will make two decisions.
3. The first decision, the unconditional contribution, is how many of your 20 tokens you want to contribute to the group project.
4. The second decision, completing the contribution table, is how many of your 20 tokens you want to contribute to the group project for each possible average contribution of your neighbors (e.g. 0, 1, 2,..., 20).
5. Each token contributed to the group project will earn 0.8 tokens for each participant in the neighborhood.
6. One of the two decisions, the unconditional contribution or the contribution table, will be randomly chosen to determine earnings. You will not know ahead of time which decision will count.

Are there any questions before we begin?

Instructions for Task 4 – B-Task

In Task 4, you will make decisions again as you did in the Task 2 B-Task. The difference between this task and Task 2 is that for each token contributed to the group project you, and the other two neighbors, will get 0.8 tokens back.

Things to remember:

1. You will be randomly placed on the network at the beginning of the task and assigned to a 3-participant neighborhood. Your neighbors in Task 4 may be different than in the previous tasks, however, you will remain with the same neighbors for all rounds in this task.

2. You will face the same decision situation for several rounds. You must decide how many of your 20 tokens you want to contribute to the group project.

3. The number of rounds is randomly determined. After each round, there is an 85

4. If there is another round, you will see a decision screen to make another choice. If there is not another round, you will be a message saying the task is over.

Are there any questions before we begin?

References

Ambrus, A., & Pathak, P. A. (2011). Cooperation over finite horizons: A theory and experiments. *Journal of Public Economics*, *95*(7), 500–512.

Anderson, S. P., Goeree, J. K., & Holt, C. A. (2004). Noisy directional learning and the logit equilibrium. *The Scandinavian Journal of Economics*, *106*(3), 581–602.

Andreoni, J. (1988). Why free ride?: Strategies and learning in public goods experiments. *Journal of Public Economics*, *37*(3), 291–304.

Andreoni, J. (1995a). Cooperation in public-goods experiments: kindness or confusion? *The American Economic Review*, 891–904.

Andreoni, J. (1995b). Warm-glow versus cold-prickle: the effects of positive and negative framing on cooperation in experiments. *The Quarterly Journal of Economics*, 1–21.

Arifovic, J., & Ledyard, J. (2012). Individual evolutionary learning, other-regarding preferences, and the voluntary contributions mechanism. *Journal of Public Economics*, *96*(9), 808–823.

Arthur, W. B. (1991). Designing economic agents that act like human agents: A behavioral approach to bounded rationality. *The American Economic Review*, *81*(2), 353–359.

Axelrod, R. M. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration.* Princeton University Press.

Axtell, R. (2000a). *Effects of interaction topology and activation regime in several multi-agent systems.* Springer.

Axtell, R. (2000b). Why agents? on the varied motivation for agent computing in the social sciences. *Brookings Institution CSED Technical Report*.

Axtell, R. (2007). What economic agents do: How cognition and interaction lead to emergence and complexity. *The Review of Austrian Economics*, *20*(2), 105–122.

Barabási, A. L. (2009). Scale-free networks: A decade and beyond. *Science*, *325*(5939), 412–413. doi: 10.1126/science.1173299

Barabási, A. L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*(5439), 509.

Bardsley, N., & Moffatt, P. (2007). The experimetrics of public goods: Inferring motivations from contributions. *Theory and Decision*, *62*(2), 161–193.

Bhat, C. R. (2001). Quasi-random maximum simulated likelihood estimation of the mixed multinomial logit model. *Transportation Research Part B: Methodological*, *35*(7), 677–693.

Bolton, G., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American economic review*, 166–193.

Botelho, A., Harrison, G. W., Pinto, L. M. C., & Rutström, E. E. (2009). Testing static game theory with dynamic experiments: a case study of public goods. *Games and Economic behavior*, *67*(1), 253–265.

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly weather review*, *78*(1), 1–3.

Brosig, J., Riechmann, T., & Weimann, J. (2007). Selfish in the end?: An investigation of consistency and stability of individual behavior.

Burlando, R., & Guala, F. (2005). Heterogeneous agents in public goods experiments. *Experimental Economics*, *8*(1), 35–54.

Burton-Chellew, M. N., Nax, H. H., & West, S. A. (2015). Payoff-based learning explains the decline in cooperation in public goods games. *Proceedings of the Royal Society of London B: Biological Sciences*, *282*(1801), 20142678.

Burton-Chellew, M. N., & West, S. A. (2013). Prosocial preferences do not explain human cooperation in public-goods games. *Proceedings of the National Academy of Sciences*, *110*(1), 216–221.

Bush, R. R., & Mosteller, F. (1955). Stochastic models for learning.

Camerer, C., & Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*, 827–874.

Cartwright, E. J., & Lovett, D. (2014). Conditional cooperation and the marginal per capita return in public good games. *Games*, *5*(4), 234–256.

Castillo, D., & Saysel, A. K. (2005). Simulation of common pool resource field experiments: a behavioral model of collective action. *Ecological economics*, *55*(3), 420–436.

Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly journal of Economics*, 817–869.

Cherchi, E., & de Dios Ortúzar, J. (2008). Empirical identification in the mixed logit model: analysing the effect of data richness. *Networks and Spatial Economics*, *8*(2-3), 109–124.

Cheung, Y.-W., & Friedman, D. (1997). Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior*, *19*(1), 46–76.

Chmura, T., Goerg, S. J., & Selten, R. (2012). Learning in experimental 2× 2 games. *Games and Economic Behavior*, *76*(1), 44–73.

Clemens, C., & Riechmann, T. (2006). Evolutionary dynamics in public good games. *Computational Economics*, *28*(4), 399–420.

Cooper, D. J., Garvin, S., & Kagel, J. H. (1997). Signalling and adaptive learning in an entry limit pricing game. *The RAND Journal of Economics*, 662–683.

Cooper, D. J., & Stockman, C. K. (2002). Fairness and learning: an experimental examination. *Games and Economic Behavior*, *41*(1), 26–45.

Cox, J. C., Friedman, D., & Sadiraj, V. (2008). Revealed altruism1. *Econometrica*, *76*(1), 31–69.

Cox, J. C., & Sadiraj, V. (2007). On modeling voluntary contributions to public goods. *Public Finance Review*, *35*(2), 311–332.

Croson, R. T. (2007). Theories of commitment, altruism and reciprocity: Evidence from linear public goods games. *Economic Inquiry*, *45*(2), 199–216.

Cross, J. G., et al. (2008). A theory of adaptive economic behavior. *Cambridge Books*.

Deadman, P. J. (1999). Modelling individual behaviour and group performance in an intelligent agent-based simulation of the tragedy of the commons. *Journal of Environmental Management*, *56*(3), 159–172.

Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems*, *3*(01n04), 87–98.

Deza, M.-M., & Deza, E. (2006). *Dictionary of distances.* Elsevier.

Duffy, J., & Ochs, J. (2009). Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior*, *66*(2), 785–812.

Epstein, J., & Axtell, R. (1996). *Growing artificial societies: social science from the bottom up.* The MIT Press.

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*, *39*(1), 111–128.

Erev, I., Ert, E., & Roth, A. E. (2010). A choice prediction competition for market entry games: An introduction. *Games*, *1*(2), 117–136.

Erev, I., & Haruvy, E. (2005). Generality, repetition, and the role of descriptive learning models. *Journal of Mathematical Psychology*, *49*(5), 357–371.

Erev, I., & Haruvy, E. (2013). Learning and the economics of small decisions. *The handbook of experimental economics*, *2*.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, 848–881.

Erev, I., Roth, A. E., Slonim, R. L., & Barron, G. (2007). Learning and equilibrium as useful approximations: Accuracy of prediction on randomly selected constant sum games. *Economic Theory*, *33*(1), 29–51.

Ert, E., & Erev, I. (2007). Replicated alternatives and the role of confusion, chasing, and regret in decisions from experience. *Journal of Behavioral Decision Making*, *20*(3), 305–322.

Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, *90*(4), 980–994.

Fehr, E., & Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, *114*(3), 817–868. doi: 10.2307/2586885

Feltovich, N. (2000). Reinforcement-based vs. belief-based learning models in experimental asymmetric-information games. *Econometrica*, *68*(3), 605–641.

Fischbacher, U., & Gächter, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review*, *100*(1), 541-56. doi: 10.1257/aer.100.1.541

Fischbacher, U., Gächter, S., & Fehr, E. (2001). Are people conditionally cooperative? evidence from a public goods experiment. *Economics Letters*, *71*(3), 397–404.

Fudenberg, D., & Levine, D. K. (1998). *The theory of learning in games* (Vol. 2). MIT press.

Gächter, S., Renner, E., & Sefton, M. (2008). The long-run benefits of punishment. *Science*, *322*(5907), 1510–1510.

Gilbert, G. N. (2008). *Agent-based models* (No. 153). Sage.

Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of political economy*, 119–137.

Heckbert, S., Baynes, T., & Reeson, A. (2010). Agent-based modeling in ecological economics. *Annals of the New York Academy of Sciences*, *1185*(1), 39–53.

Herrmann, B., & Thöni, C. (2009, March). Measuring conditional cooperation: a replication study in russia. *Experimental Economics*, *12*(1), 87-92.

Herrmann, B., Thöni, C., & Gächter, S. (2008). Antisocial punishment across societies.

Science, *319*(5868), 1362–1367.

Hess, S., & Rose, J. M. (2012). Can scale and coefficient heterogeneity be separated in random coefficients models? *Transportation*, *39*(6), 1225–1239.

Ho, T. H., Camerer, C. F., & Chong, J.-K. (2007). Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory*, *133*(1), 177–198.

Hopkins, E. (2002). Two competing models of how people learn in games. *Econometrica*, 2141–2166.

Huberman, B. A., & Glance, N. S. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences*, *90*(16), 7716–7718.

Isaac, R. M., & Walker, J. M. (1988). Group size effects in public goods provision: The voluntary contributions mechanism. *The Quarterly Journal of Economics*, 179–199.

Isaac, R. M., Walker, J. M., & Thomas, S. H. (1984). Divergent evidence on free riding: An experimental examination of possible explanations. *Public choice*, *43*(2), 113–149.

Isaac, R. M., Walker, J. M., & Williams, A. W. (1994). Group size and the voluntary provision of public goods: experimental evidence utilizing large groups. *Journal of Public Economics*, *54*(1), 1–36.

Janssen, M. A., & Ahn, T.-K. (2006). Learning, signaling, and social preferences in public-good games. *Ecology and society*, *11*(2), 21.

Janssen, M. A., & Ostrom, E. (2006). Empirically based, agent-based models. *Ecology and Society*, *11*(2), 37.

Kennedy, W. G., Cotla, C. R., Gulden, T., Coletti, M., & Cioffi-Revilla, C. (2014). Towards validating a model of households and societies in east africa. In *Advances in computational social science* (pp. 315–328). Springer.

Keser, C., & Van Winden, F. (2000). Conditional cooperation and voluntary contributions to public goods. *The Scandinavian Journal of Economics*, *102*(1), 23–39.

Kocher, M. G., Cherry, T., Kroll, S., Netzer, R. J., & Sutter, M. (2008). Conditional cooperation on three continents. *Economics Letters*, *101*(3), 175–178.

Kosfeld, M., Okada, A., & Riedl, A. (2009). Institution formation in public goods games. *The American Economic Review*, 1335–1355.

Kümmerli, R., Burton-Chellew, M. N., Ross-Gillespie, A., & West, S. A. (2010). Resistance to extreme strategies, rather than prosocial preferences, can explain human cooperation in public goods games. *Proceedings of the National Academy of Sciences*, *107*(22), 10125–10130.

Kurzban, R., & Houser, D. (2001). Individual differences in cooperation in a circular public goods game. *European Journal of Personality*, *15*(S1), S37–S52.

Kurzban, R., & Houser, D. (2005). Experiments investigating cooperative types in humans: A complement to evolutionary theory and simulations. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(5), 1803.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, *52*(4), 667–694.

Ledyard, J. O. (1995). Public goods: some experimental results. In J. Kagel & A. Roth (Eds.), *Handbook of experimental economics.* Princeton, NJ: Princeton University Press.

Lin, J. (1991). Divergence measures based on the shannon entropy. *Information Theory, IEEE Transactions on*, *37*(1), 145–151.

Liu, Y.-H., & Mahmassani, H. S. (2000). Global maximum likelihood estimation procedure for multinomial probit (mnp) model parameters. *Transportation Research Part B: Methodological*, *34*(5), 419–449.

Loomes, G. (2005). Modelling the stochastic component of behaviour in experiments: Some issues for the interpretation of data. *Experimental Economics*, *8*(4), 301–323.

Luce, R. D. (2005). *Individual choice behavior: A theoretical analysis.* Courier Corporation.

Lugovskyy, V., Puzzello, D., Sorensen, A., Walker, J., Williams, A., & Hall, W. (2015). An experimental study of finitely and infinitely repeated linear public goods games.

Macy, M. W., & Willer, R. (2002). From factors to actors: Computational sociology and agent-based modeling. *Annual review of sociology*, 143–166.

Marchiori, D., & Warglien, M. (2008). Predicting human interactive learning by regret-driven neural networks. *Science*, *319*(5866), 1111–1113.

Miller, J. H., & Andreoni, J. (1991). Can evolutionary dynamics explain free riding in experiments? *Economics Letters*, *36*(1), 9–15.

Moffatt, P. G., & Peters, S. A. (2001). Testing for the presence of a tremble in economic experiments. *Experimental Economics*, *4*(3), 221–228.

Muller, L., Sefton, M., Steinberg, R., & Vesterlund, L. (2008). Strategic behavior and learning in repeated voluntary contribution experiments. *Journal of Economic Behavior & Organization*, *67*(3), 782–793.

Nikiforakis, N., & Normann, H.-T. (2008). A comparative statics analysis of punishment in public-good experiments. *Experimental Economics*, *11*(4), 358–369.

Nowak, M. (2006). Five rules for the evolution of cooperation. *Science*, *314*(5805), 1560–1563. doi: 10.1126/science.1133755

Nowak, M., Bonhoeffer, S., & May, R. (1994). Spatial games and the maintenance of cooperation. *Proceedings of the National Academy of Sciences of the United States of America*, *91*(11), 4877.

Nowak, M., & May, R. (1992). Evolutionary games and spatial chaos. *Nature*, *359*(6398), 826–829.

Poncela, J., Moreno, Y., et al. (2007). Robustness of cooperation in the evolutionary prisoner's dilemma on complex networks. *New Journal of Physics*, *9*(6), 184.

Rand, D., Arbesman, S., & Christakis, N. (2011). Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences*, *108*(48), 19193–19198. doi: 10.1073/pnas.1108243108

Robinson, D. T., Brown, D. G., Parker, D. C., Schreinemachers, P., Janssen, M. A., Huigen, M., ... others (2007). Comparison of empirical methods for building agent-based models in land use science. *Journal of Land Use Science*, *2*(1), 31–55.

Roca, C. P., Cuesta, J. A., & Sánchez, A. (2009). Effect of spatial structure on the evolution of cooperation. *Physical Review E*, *80*(4), 046106.

Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior*, *8*(1), 164–212.

Salmon, T. C. (2001). An evaluation of econometric models of adaptive learning. *Econometrica*, *69*(6), 1597–1628.

Santos, F., Santos, M., & Pacheco, J. (2008). Social diversity promotes the emergence of cooperation in public goods games. *Nature*, *454*(7201), 213–216.

Santos, F. C., & Pacheco, J. M. (2005). Scale-free networks provide a unifying framework for the emergence of cooperation. *Physical Review Letters*, *95*(9), 098104.

Sass, M., & Weimann, J. (2012). *The dynamics of individual preferences in repeated public good experiments.* Univ., Faculty of Economics and Management.

Schelling, T. (1998). *Micromotives and macrobehavior.* WW Norton & Company.

Sefton, M., Shupp, R., & Walker, J. (2007). The effect of rewards and sanctions in provision of public goods. *Economic Inquiry*, *45*(4), 671–690.

Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, *1*(1), 43–62.

Selten, R., & Chmura, T. (2008). Stationary concepts for experimental 2× 2-games. *The American Economic Review*, 938–966.

Selten, R., & Stoecker, R. (1986). End behavior in sequences of finite prisoner's dilemma supergames a learning theory approach. *Journal of Economic Behavior & Organization*, *7*(1), 47–70.

Simon, H. (1996). *The sciences of the artificial.* Cambridge, MA: The MIT Press.

Smajgl, A., & Barreteau, O. (2014). *Empirical agent-based modelling-challenges and solutions.* Springer.

Smith, V. L. (1962). An experimental study of competitive market behavior. *The Journal* *of Political Economy*, 111–137.

Smith, V. L. (1976). Experimental economics: Induced value theory. *The American Economic Review*, *66*(2), 274–279.

Smith, V. L. (1982). Microeconomic systems as an experimental science. *The American Economic Review*, *72*(5), 923–955.

Suri, S., & Watts, D. (2011). Cooperation and contagion in web-based, networked public goods experiments. *PloS one*, *6*(3), e16836. doi: 10.1371/journal.pone.0016836

Szabó, G., & Fáth, G. (2007). Evolutionary games on graphs. *Physics Reports*, *446*(4), 97–216.

Tan, L., & Wei, L. (2014). Voluntary contribution mechanism played over an infinite horizon. *Pacific Economic Review*, *19*(3), 313–331.

Thurstone, L. L. (1930). The learning function. *The Journal of General Psychology*, *3*(4), 469–493.

Train, K. E. (2009). *Discrete choice methods with simulation.* Cambridge university press.

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica: Journal of the Econometric Society*, 307–333.

Wang, J., Suri, S., & Watts, D. J. (2012). Cooperation and assortativity with dynamic partner updating. *Proceedings of the National Academy of Sciences*, *109*(36), 14363–14368.

Watts, D. (2003). *Small worlds: The dynamics of networks between order and randomness (princeton studies in complexity).* Princeton, NJ: Princeton University Press.

Watts, D., & Strogatz, S. (1998). Collective dynamics of 'small-world' networks. *Nature*, *393*(6684), 440–442. doi: 10.1038/30918

Weimann, J. (1994). Individual behaviour in a free riding experiment. *Journal of Public Economics*, *54*(2), 185–200.

Wendel, S., & Oppenheimer, J. (2010). An agent-based analysis of context-dependent preferences. *Journal of Economic Psychology*, *31*(3), 269–284.

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

Wunder, M., Suri, S., & Watts, D. J. (2013). Empirical agent based models of cooperation in public goods games. In *Proceedings of the fourteenth acm conference on electronic commerce* (pp. 891–908).

Curriculum Vitae

Chenna Reddy Cotla completed his secondary education in 2000 from Zilla Parishad High School, Polur and higher secondary education in 2002 from Ravindra Junior College, Kurnool in the state of Andhra Pradesh in India. He earned a Bachelor of Technology and a Master of Technology in Information Technology in 2008 from Indian Institute of Information Technology and Management, Gwalior, India. Before coming to George Mason for his PhD, he worked as a researcher in Machine Learning and Soft Computing group at Dipartimento di Informatica e Scienze dell'Informazione, University of Genoa, Italy.