Private Information in Common Pools
- A Repeated Game

MENGMENG DU (18914)

Stockholm School of Economics
Master of Science Thesis
International Economics 5110
Stockholm, Sweden, 2008

Presentation: Nov 26, 2008 (Room 328)
Supervisor: Magnus Johannesson
Examiner: Örjan Sjöberg
Opponent: Anders Eriksson Tempelmark

Abstract

In this study, we carry out multi-round two-player sequential common pool resource (CPR) experiments in order to see how players' behaviours develop over time. We focus on the differences between two experiment treatments, one where both players know the available resource size, and one where only the first mover knows the available sum.

We find that after repeated game rounds, there is a clear first mover advantage in the complete information sequential CPR game. First movers, however, lose ground despite having an information advantage in the private information game. Here, there is no significant difference in players' payoffs and they most commonly play according to the 50/50 split. Thus, with sequential CPR games, information asymmetry actually leads to a more equal sharing of available resources.

Our results show learning occurring in both experiment treatments. In the complete information case, players start off on a relatively equal level, but the first mover advantage increases over time. In the private information case, where we initially have a first mover advantage and high inefficiency, players adjust with repetition to make more equal claims on the resource, leading to significantly increased efficiency towards the end of the experiment.

However, in our ten-round experiment, some players still failed to coordinate efficiently in the last private information round. Despite that, all developments point to the 50/50 split being the equilibrium for private information sequential CPR games. We suspect that repeating the experiment in this study with a larger number of rounds would provide solid experimental results showing this equilibrium.

Acknowledgements

I would like to express my sincere gratitude to my supervisor Magnus Johannesson for his excellent advice and invaluable support throughout my thesis work, and Therese Lindahl, whose Ph.D. dissertation forms the basis for this thesis. I also want to thank Andreas Ehn for all his help with the experiments. Last but not least, I am very grateful to the Jan Wallander and Tom Hedelius Foundation for its financial support.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Problem Overview

A common pool resource (CPR) is a natural or man-made resource shared by rivaling users, which cannot be excluded from it without difficulty (Lindahl and Johannesson 2005). Being non-excludable, CPRs risk being overused. Hardin (1968) describes in his classic paper about the tragedy of the commons how a pasture open to all can be overgrazed. In his example, each herder receives a direct benefit from his own animals and suffers only a part of the costs from the deterioration of the pasture. Thus, the herders are motivated to add more cattle until the pasture is destroyed.

There are many other situations dealing with CPRs in real life, with examples ranging from divorce proceedings (Camerer et al. 2003) to international cooperation over fishing waters (Ostrom 1990). This, combined with the inefficiencies often associated with CPRs, has resulted in a great deal of research on how people behave in these situations.

Lindahl and Johannesson (2005) have carried out an experimental study of a CPR situation where two players make sequential claims on an amount in a single-shot game. The players receive their claims only if the sum of the claims does not exceed the available amount and otherwise nothing. The authors tested both the case where both players know the size of the resource, and the case where only the player who moves first is informed of the resource size, to reflect that in the real world, people often have different sources of information and can be differently skilled in processing available information.

Standard game theory predicts that if the players are selfish, the first player will claim the entire resource except for a small token amount. The second player, knowing that the first player will behave in this manner, will claim practically nothing, independent of whether she knows the resource size or not.

In their study, Lindahl and Johannesson show that the first player, contrary to standard game theoretical predictions, claims considerably less than the whole resource in both the complete information and the private information setting, indicating that fairness is important.

What makes their study even more interesting is that in the case of private information, the first player claims even less than when the resource size is publicly known. When only the first player knows the resource size, most pairs coordinate on the even split. Not only does the first player fail to gain from her informational advantage, but it might be the case that the second player wins from her lack of information.

However, there were inefficiencies in both information treatments, with pairs failing to coordinate. Also, there was more under-exploitation in the private information case, leading to an efficiency loss compared to the complete information treatment. Lindahl and Johannesson write that learning should be important for coordination and that in a repeated environment, it might increase efficiency in the private information case by driving the results towards coordination on the even split. This would mean that information asymmetry could actually lead to more equal sharing of resources.

1.2 Purpose

The purpose of this thesis is to study the validity of Lindahl and Johannesson (2005)'s hypothesis that learning will, in the private information treatment, lead to increased efficiency and coordination on the even split. We will do this by repeating Lindahl and Johannesson's experiment with more testing rounds.

More specifically, the questions we want to answer are: In a repetitive setting, when the first mover has more information about the resource size, will two players share a common resource more equally than when they both know the resource size? Will they learn to coordinate more efficiently over time? And is the even split the long-term equilibrium?

1.3 Definitions

We define an outcome to be efficient when the sum of both players' claims is exactly equal to the resource size. If their claims add up to less than the available amount, we have under-exploitation. If their claims together exceed the available amount, we have a resource breakdown. We also define learning as an observed change in a player's behaviour due to experience.

1.4 Limitations of Scope

In this experimental study, we will not take into consideration possible differences in behaviour because of gender, age, ethnicity, academic major, etc. The reason is insufficient sample size for these factors due to financial limitations restricting the number of participants.

1.5 Outline

In chapter 2, we dive into existing research regarding the CPR dilemma. We look first at standard game theory, and then we go into some experimental studies refuting standard theory. In particular, we describe Lindahl and Johannesson's study in more detail. After covering existing research, the experiment setup of this study is explained in chapter 3, along with the hypotheses of the experiment outcomes, which we base on the research overview given in chapter 2. We describe the statistics used to test these hypotheses. Next, in chapter 4, we present the experiment outcomes and the statistical test results. Finally, in chapter 5, we discuss the meaning of these results and conclude by summing up the main findings.

Chapter 2

Overview of Previous Research

In this chapter, we present an overview of existing research around CPR dilemmas, in order to provide a foundation on which to build the hypotheses for the experiments of this study. We first describe the sequential CPR game used by Lindahl and Johannesson (2005) and look at the model predictions for this game from a standard game theoretical perspective. We then go through their study, in which the experimental results contradict the standard model predictions. We also present briefly some other findings from experimental studies of sequential CPR (or equivalent) games. These studies all show that fairness considerations affect the real life outcome so that it differs from standard theory. Some of these studies also explore the effect of learning.

2.1 The CPR Dilemma Model[1]

In the sequential CPR game, there are two players, player A and player B, with exclusive access to a common resource of size x, for which they compete. These players make sequential discrete claims, r_A and r_B, on the resource. Player A makes her claim first. If the sum of both claims is equal or less than the total resource size, both players receive their respective claims. Else, if the sum of both claims exceeds the total size, they both receive nothing. Formally, player j's payoff is

$$p_j = \begin{cases} r_j & \text{if } r_A + r_B \leq x \\ 0 & \text{otherwise} \end{cases} \quad ; \quad j \in \{A, B\}\,. \tag{2.1}$$

2.1.1 Complete Information

Denote the game where both players have complete information about the resource size as $\Gamma_c(2)$. Under the conventional assumption that players only care about their

[1] This section is based on Lindahl and Johannesson (2005).

own material payoff, there are two subgame perfect equilibria (SPE)[2] outcomes of $\Gamma_c(2)$ (for proof, see Lindahl and Johannesson (2005)):

$$\begin{cases} p_A = r_A = x - \epsilon; & p_B = r_B = \epsilon \\ p_A = r_A = x; & p_B = r_B = 0 \end{cases} \tag{2.2}$$

In the first case, player A claims the entire resource but a very small amount, and player B claims this small amount. In the second case, player A claims the entire resource, and player B claims nothing. In both cases, the resource is divided efficiently and both players' payoffs equal their respective claims. Under complete information, the first mover, player A, has a clear advantage.

It is worth noting that the sequential CPR game under complete information is equivalent to the ultimatum game. The only difference is that in the ultimatum game, the second player accepts or rejects a proposed division of the resource instead of making her own claim on it. The equilibrium outcomes of the two games are the same.

2.1.2 Private Information

Denote the game where only the first player, player A, knows the resource size, and the second player, player B, is uncertain about it as $\Gamma_p(2)$. In this game, player B only knows that the resource size is distributed according to a discrete uniform probability distribution given by

$$f(x) = \frac{z}{\overline{k} - \underline{k} + z}, \tag{2.3}$$

where the discrete values that the resource size can take are

$$x \in \left\{\underline{k}, \underline{k} + z, \underline{k} + 2z, \ldots, \overline{k}\right\}, \tag{2.4}$$

and the mean resource size is

$$\mu(x) = \frac{\underline{k} + \overline{k}}{2}. \tag{2.5}$$

When there is uncertainty about the resource size, the SPE concept is inappropriate since not all players can determine for certain what their best response is to other players' strategies. Lindahl and Johannesson write that we need to use the perfect Bayesian equilibria (PBE) concept. A PBE is a set of strategies and beliefs such that at any stage of the game, players use optimal strategies given their beliefs. These beliefs are based on equilibrium strategies and updated with observed actions using Bayes' rule[3].

[2] A SPE is an equilibrium in an extensive-form game where players' strategies constitute a Nash equilibrium in every possible subgame. In a Nash equilibrium, no player has incentive to deviate from their strategy, if no other player will deviate from theirs (Camerer 2003).

[3] Bayes' rule states that

$$P(H_i \mid D) = \frac{P(H_i)P(D \mid H_i)}{\sum_k^n P(H_k)P(D \mid H_k)},$$

where H_i are the n possible hypotheses and D is the observed data (Camerer 2003).

According to Lindahl and Johannesson, by using the intuitive criterion of Cho and Kreps (1987), which puts restrictions on out-of-equilibrium beliefs based on the idea that certain types of players should not be expected to use certain strategies, we will only have two PBE outcomes of $\Gamma_p(2)$:

$$\begin{cases} p_A = r_A = x - \epsilon; & p_B = r_B = \epsilon; \quad \epsilon < z \\ p_A = r_A = x; & p_B = r_B = 0 \end{cases} \tag{2.6}$$

In both PBE outcomes, the resource is divided efficiently and both players' payoffs equal their respective claims. Also, player A, the first mover, maintains her advantage over player B.

Thus, from a standard game theoretical perspective, the outcome of the sequential CPR game is the same, independent on whether the second player knows the resource size or not. We have an efficient split of the resource, and there is a clear first mover advantage.

2.2 CPR Study by Lindahl and Johannesson (2005)

Lindahl and Johannesson (2005) used two experimental treatments to test the outcomes of the sequential CPR game as described above.

2.2.1 Experiment Design

The resource was represented by a sum of money x, randomly drawn from a uniform discrete distribution between SEK 100 and SEK 200, with SEK 10 increments, i.e.

$$x \in \{100, 110, \ldots, 200\}, \tag{2.7}$$

with the likelihood of getting a certain sum in the range being

$$f(x) = \frac{10}{200 - 100 + 10} = \frac{1}{11}, \tag{2.8}$$

and the mean resource size being

$$\mu(x) = \frac{100 + 200}{2} = 150, \tag{2.9}$$

corresponding to equations 2.3-2.5 on page 5.

In treatment 1, the game participants had complete information about the resource size. In treatment 2, only the first player, player A, knew the size of the resource, while player B only knew the distribution, just as in the private information game described in section 2.1.2.

The participants, students from Stockholm School of Economics (SSE), were paid a guaranteed fee of SEK 40. If the sum of their claims was less than or equal to the resource size, they also received their respective claims. Each participant only took part in one of the two treatments, and only played the game a single round.

The players were divided into two rooms, and each player was paired anonymously with another in the other room. Before the game, participants were given instructions and had opportunity to ask questions to the experimenter. During the game, player A wrote down her claim on a form, which the experimenter then handed out to the matching player B in the other room. Player B then wrote down her claim on the same form, which was handed back to player A.

2.2.2 Results

Complete Information

In treatment 1, with complete information, in over half of the 47 cases, player A claimed more than half of the pie. The average claim was 55% (see table 2.1). Most of the claims were 50-70%. Only in 3 out of 47 pairs (6%) did player A claim less than half, and there were only 2 cases (4%) where player A claimed more than 70%.

The difference in payoff between players A and B expressed as percentage of the resource size was significant at the 10% level when tested using both the independent samples t-test and the non-parametric Mann-Whitney test (see table 2.2)[4]. This shows that in the complete information sequential CPR game, there is some first mover advantage. However, the advantage is much smaller than what theory, as described in section 2.1.1, predicts.

Overall, there were some coordination failures, largely due to resource breakdowns (see table 2.3). Lindahl and Johannesson show that the larger claim player A makes, the higher is the probability of breakdown. This means that player B does not accept all divisions, once again refuting the theoretical predictions. The authors also show that neither the percentage claim nor the probability of breakdown were significantly related to the resource size for either players (see table 2.4.)

Private Information

In treatment 2, with private information, player A claimed exactly half of the pie in 62% of the 77 cases (see table 2.3). 30% of claims by player A were smaller than half, and only 10% were larger. The average claim was 49% (see table 2.1).

The payoff difference between players A and B, expressed as percentage of the resource size, was significant at the 10% level with both the t-test and the Mann-Whitney test (see table 2.2). The first mover has a significant advantage, albeit much smaller than what theory predicts (see section 2.1.2). Lindahl and Johannesson write that this first mover advantage is a result of under-exploitation, which happened in 49% of all cases. Total rate of coordination failures was 60%. Resource breakdowns most often coincided with higher claims by player A, further contradicting theoretical predictions of player B being willing to accept all resource divisions.

[4] The authors used both tests, since bargaining experiments typically lead to skewed, non-normal distributions, which means that the underlying normality assumption of the t-test does not hold. Their data was tested using the Kolmogorov-Smirnov normality test and almost all variables were found to be significantly skewed at the 5% level.

	Complete information		Private information		p-value of diff. (two-sided)	
	Result	Std. Dev.	Result	Std. Dev.	t-test	Mann-W.
Avg A claim (SEK)	88.38	20.58	73.80	23.62	0.036	0.012
Avg A claim (%)	55.09	8.25	48.93	9.34	<0.001	<0.001
Avg A payoff (SEK)	72.45	31.79	64.52	31.63	0.180	0.058
Avg A payoff (%)	48.01	18.00	42.38	17.74	0.093	<0.001
Avg B claim (SEK)	73.53	27.84	64.56	20.46	0.059	0.247
Avg B claim (%)	48.65	13.21	44.43	15.17	0.106	0.154
Avg B payoff (SEK)	62.43	28.33	55.34	27.86	0.177	0.197
Avg B payoff (%)	41.34	15.87	37.28	18.15	0.194	0.067

Table 2.1: Lindahl and Johannesson (2005): Average claims and payoffs.

	Complete information p-value (two-sided)		Private information p-value (two-sided)	
	t-test	Mann-W.	t-test	Mann-W.
Payoff diff. (SEK)	0.110	0.043	0.058	0.025
Payoff diff. (%)	0.060	<0.001	0.080	0.007

Table 2.2: Lindahl and Johannesson (2005): First mover advantage.

	Complete information	Private information	Pearson χ^2
50% claims by A	49%	62%	0.143
Imitations by B	51%	52%	0.924
Breakdowns	11%	12%	0.858

Table 2.3: Lindahl and Johannesson (2005): Proportions.

	Complete information		Private information	
	β	p-value	β	p-value
Linear regression:				
Claim A (%) – Sum	-0.024	0.520	0.042	0.217
Claim B (%) – Sum	0.018	0.767	-0.218	<0.001
Logistic regression:				
Breakdown – Sum	-0.009	0.561	-0.014	0.226

Table 2.4: Lindahl and Johannesson (2005): The importance of resource size.

52% of player Bs imitated player A's claim (see table 2.3). Overall, 57% of the participants played according to the even split, although this did not always lead to efficiency. However, in 97% of the cases where efficiency was reached, the players coordinated on the even split. This indicates that in the private information sequential CPR game, the even split could be the real life equilibrium outcome.

Lindahl and Johannesson show that the percentage claim was not significantly related to the resource size for player A. However, player B's percentage claims decreased significantly with the resource size. Despite this, logistic regression analysis shows that the probability of breakdown was not significantly related to the resource size. (See table 2.4.)

Comparison of Complete and Private Information

In both treatments, there is a first mover advantage which is far lower than what theory predicts. Both set of results also show clearly that player B is not willing to accept all resource divisions. Furthermore, the overall outcome is not strongly related to the resource size in either treatment.

As seen in table 2.1, there is less first mover advantage in the private treatment. Player A's payoff as percentage of the resource size in the complete information treatment is significantly higher at the 10% level using both the t-test and the Mann-Whitney test. On average, both the percentage claim and payoff of player B were lower with private information, due to under-exploitation. These results were, however, not quite significant (see table 2.1).

There was no significant difference in the proportion of breakdowns between the two treatments (see table 2.3). Lindahl and Johannesson write that in the private information case, most inefficiencies were due to under-exploitation when both players scaled down their claims, and as a result the probability of breakdown did not increase.

2.2.3 Implications

Lindahl and Johannesson state that the results from their experiment indicate that players' utilities are not only a function of their own material payoffs as standard game theory assumes, but that social preferences, such as fairness, are important. Players care both about not being inferior to others, not accepting resource divisions where they get much less than their opponent, and about not being too superior, with some examples of player A even claiming less than half of the available resource.

Under private information, a large share of all players coordinated on the even split. The authors show that the inequality-aversion model by Fehr and Schmidt (1999) together with some belief formation rules also predicts the even split outcome. They suggest that a possible reason for the coordination failures seen in their single-shot experiment is that people need to reiterate several rounds before learning to coordinate on the predicted outcome.

2.3 The Inequality-Aversion Model of Fairness

The inequality-aversion model by Fehr and Schmidt (1999) used by Lindahl and Johannesson (2005) to predict the even split outcome is also one of the most common models used for social preferences. The basic idea is that people dislike outcomes that they perceive inequitable or unfair. Also, relative payoffs matter. Player i's utility for the social allocation $X \equiv \{x_1, x_2, \ldots, x_n\}$ is

$$U_i(X) = x_i - \frac{\alpha_i}{n-1} \sum_{k \neq i} \max(x_k - x_i, 0) - \frac{\beta_i}{n-1} \sum_{k \neq i} \max(x_i - x_k, 0) \tag{2.10}$$

where $0 \leq \beta_i \leq 1$ and $\beta_i \leq \alpha_i$.

This means that players dislike both having lower and higher allocations than others. The factor α_i determines how much envy player i feels about others having more, and the factor β_i is her guilt weight. The model assumes that people dislike having more than others somewhat less than having less.

If player i is purely selfish, then $\alpha_i = \beta_i = 0$ and the utility function is solely based on material payoff as in standard game theory.

2.4 Other Experimental Studies on Fairness

Kahneman et al. (1986) provide some real life examples of fairness considerations playing a role in people's actions, e.g. when splitting resources. Around a hundred Canadian households were asked what they considered fair or unfair in different common life scenarios. The vast majority of households surveyed considered it unfair for a firm to reap all benefits from a production cost decrease, while an even split of the benefit was considered acceptable by approximately half of the respondents. Most households surveyed would also punish unfair behaviour even though it would cost them to do so. In one example, 68% of the respondents would switch to another drugstore further away than their regular, if their regular store raised prices after a competitor closed. When asked about how much tip they would leave in a restaurant they did not expect to visit again, the mean response was practically the same as for a restaurant that the respondents visited frequently. This indicates that people follow the "rules of fairness" without them being enforced, when they do not risk retaliation or embarassment or have anything to gain.

Experimental studies of the ultimatum game, which, as mentioned in section 2.1.1, is in standard game theory equivalent to the sequential CPR game, have also consistently shown that people deviate from the behaviour predicted by standard game theory. First movers claim significantly less than the entire resource, and resource breakdowns are common when they do claim most of what is available. These studies also show that fairness considerations have a significant influence on people's behaviour.

2.4.1 Complete Information Studies

Camerer (2003) writes that results from experiments involving ultimatum games are very regular. Judging from 15 different experimental studies, the first mover leaves on average 30-40% of the resource to the second mover. The median share offered to the second player is 40-50%, which is also the range where most offers fall in. Also, the first player seldom offers very little (0-10%) or more than half (51-100%) to her opponent. The second player, in turn, rejects offers less than 20% half of the time, and rarely turns down offers of 40-50%. Under-exploitation is not an issue.

According to Babcock and Loewenstein (1997), one explanation to why first movers claim more than half of the resource is self-serving bias. They write that "although psychologists debate the underlying cause of the self-serving bias, its existence is rarely questioned". The first movers, who in ultimatum games are responsible for allocating the resource, believe themselves to have a powerful role and thus deserve more than half of the pie. Hoffman et al. (1994) also show that first movers' feeling of entitlement leads them to make higher claims.

In regards to empirical similarity between the ultimatum and sequential CPR game, beyond the theoretical equivalence, Larrick and Blount (1997) write that the ultimatum game is strategically similar to the sequential CPR game, with only some slight differences. First movers claim slighly less in the CPR setting, and second movers reject the resource allocation less frequently. Their view is that "the language of claiming" makes both sides somewhat more generous by creating a sense of sharing and common ownership.

2.4.2 Private Information Studies

Under private information, the sequential CPR game differs somewhat in real life from the ultimatum game. In the CPR setting, second movers only learn the first movers' claims and have to make their own claims based on that. On the contrary, in the ultimatum setting, they choose to either accept or reject the offer given to them by the first movers. Most ultimatum game studies have shown that first movers take advantage of the second movers giving them benefit of the doubt, leading to higher first mover advantage (see e.g. Mitzkewitz and Nagel (1993), Straub and Murnighan (1995), Rapoport et al. (1996), and Croson (1996)). However, in the sequential CPR game, when first movers need to take into account the second movers' uncertainty of the resource size, the opposite seems to be the case.

As described in section 2.2.2, Lindahl and Johannesson (2005) found in their experiment that the first mover advantage was smaller with private information. Similar results were found in another single-shot sequential CPR game study by Gustafsson (1999). In his experiment, first movers made lower claims when their opponents lacked knowledge of the resource size. Also, second movers made larger claims when they did not know the available sum. Gustafsson believes that these results are due to an outcome-desirability bias, where second movers overestimate the resource size when it is uncertain[5]. First movers, who know the resource size, are aware of this bias, and therefore adapt their claims to avoid resource breakdown.

2.4.3 Pie Size Makes Little Difference

In regards to the effect the resource size has on the outcome, a number of studies have shown that there is limited difference in outcome between high and low stakes games, and that responders frequently reject substantial money offers.

Hoffman et al. (1996) compared the results of their high ($100) and low ($10) stakes ultimatum game experiments carried out in Arizona, US, and found no significant differences between the two. Others, e.g. Slonim and Roth (1998), List and Cherry (2000) and Cameron (1999), have shown that the rejection rate is somewhat lower in high stakes ultimatum games, but the difference is rather small. Second movers are slightly more likely in a high stakes game to accept percentage offers that they would have rejected if the stakes were lower. However, they still frequently reject substantial money offers. In List and Cherry's experiment, 25% of the second movers that were offered $100 of $400 available rejected the offer. Players in Cameron's Indonesian experiment rejected all offers in the 0-20% range, even when the total resource size corresponded to one month's wages.

2.5 Learning Theory

In theoretical texts, when various game equilibria are described, authors rarely mention how these are actually reached in real life. In practice, people often need time to figure out the optimal solution, or if that is not possible, to settle for a real life solution that is stable in the long term. One well-known example of this is p-beauty contests, where multiple players simultaneously pick a number between 0 and 100, and the player closest to the average of all players' choices muliplied by p wins a prize. The unique Nash equilibrium is 0, which can be devised through many steps of iterated dominance[6]. The game is difficult, since getting close to equilibrium requires players to think very far ahead and also trust that their co-players are doing the same. Ho et al. (1998) showed that players new to beauty contests at first end up far from equilibrium and make choices that are widely distributed[7]. Subsequent choices when the game is repeated, however, converge towards the equilibrium point, clearly showing how players learn from experience and each other[8].

[5]See the report for examples of studies demonstrating this bias.

[6]Take e.g. a 2/3 beauty contest. First-order iterated dominance excludes numbers larger than $(2/3) * 100 = 67$, since the average multiplied by 2/3 can impossibly be higher. Second-order iterated dominance then excludes all numbers larger than $(2/3) * 67 = 44$. Continuing in the same way will eventually lead to the unique Nash equilibrium 0.

[7]The median in their first experiment rounds was, depending on the value of p, only 1-2 steps of iterated dominance, with most players spread across 0-3 steps.

[8]After ten rounds, the median player in finite-threshold games ($p > 1$), where the equilibrium can be reached in a finite number of steps, chose the equilibrium, while the median player in infinite-threshold games ($p < 1$) iterated 6 steps or more.

There are several theories of learning[9], the two most common being reinforcement learning and belief learning. Reinforcement models assume that players' strategies are reinforced by their previous payoffs. However, unchosen strategies are not taken into account. Players thus learn from their successes but not their failures. Belief models assume players continuously update their beliefs about other players' actions based on history, and base their decisions on those beliefs. They ignore, however, information about their own past payoffs. Both of these approaches generally perform tolerably in predicting the direction players will take. They are, however, weak in different ways, since they both assume people ignore important information.

Camerer and Ho (1999) came up with the concept of experience-weighted attraction (EWA) learning, which combines the usefulness of reinforcement and belief learning. EWA assumes that people learn from both chosen and unchosen strategies, but may respond more strongly to their actual received payoffs than "what could have beens". It also considers that people might be affected more by recent than past experiences, i.e. people either forget or discard old information as time goes by. Applying the EWA learning approach to e.g. the p-beauty contest yields, according to Camerer et al. (2002), clearly better results than when using a reinforcement learning approach, and slightly better results than using a belief learning approach. In general, Camerer (2003) mentions that EWA predicts better than reinforcement and belief learning in 80-95% of the cases in which it has been studied.

2.6 Empirical Studies on Learning

There is ample empirical research on the effects of learning when it comes to multi-round complete information ultimatum game variants. Fewer studies, however, cover the private information case, even though it is the more interesting one seen from a learning perspective due to the coordination problems.

2.6.1 Complete Information

Many studies have shown that experience has limited effect on participants' decisions when they know the available resource size. Bolton and Zwick (1995) found no significant difference in player behaviour over multiple game rounds, while e.g. Slonim and Roth (1998) and List and Cherry (2000) found a slight tendency for first movers to claim more and second movers to reject less frequently over time. Roth et al. (1991) carried out the same experiment in four different countries. Participants repeated the ultimatum game for ten rounds, with different opponents in each round. The mean offer in the last round was in all countries approximately the same as the initial mean offer, with only a somewhat increased concentration of offers in the 30-40% bracket over time.

[9] See Chapter 6 in Camerer (2003) for an overview.

2.6.2 Private Information

One study on learning in a private information setting was carried out by Mitzkewitz and Nagel (1993). 80 participants played a demand variant of the ultimatum game, which closely resembles the sequential CPR game, in that first movers state their own claims on the pie instead of what they would like to offer their opponents. The game was repeated for eight rounds, with different player matching in each round. After each round, players were informed about their own and their opponent's payoffs. As in the study by Lindahl and Johannesson (2005), first movers had complete information, while second movers only knew the probability distribution. Resource size was determined in each case by rolling a dice, and was thus evenly distributed. The amounts involved were rather small, with a maximum of approximately $4.

Mitzkewitz and Nagel's results were somewhat inconclusive in regards to learning effects, with e.g. no coherent trend for player B's behaviour with increasing experience. They did, however, note that player A most often stayed on the same demand level if she managed to coordinate successfully in the previous round, while she was more likely to moderate her demand if the previous round was a failure. Player B's behaviour most often remained unchanged, independent on the outcome in the previous round. This would suggest that in time, the gap between the claims of player A and player B could decrease.

Chapter 3

Experiment Setup

In the following sections, we describe the experiment setup of this study and, based on the existing research presented previously in chapter 2, the expected outcomes. We also state the statistics needed to test these hypotheses using the experiment data we obtain. For a summary of the hypotheses of this study and the test statistics used, please see table 3.2 on page 20.

3.1 Design

As mentioned in section 1.2, the purpose of this study is to see if the results from the single-shot experiment on the sequential CPR game performed by Lindahl and Johannesson (2005) are strengthened in a repetitive setting. Therefore, it is important that the experiment used in this study mimick the one used by Lindahl and Johannesson as much as possible.

The experiment was carried out in September 2005 at SSE's computer labs. The participants were a mix of business students from SSE and engineering students from Royal Institute of Technology (KTH). According to the earlier experiment, described in section 2.2, two experimental treatments were used, one where the players had complete information, and one with private information, where only the first mover knew the available resource size.

In total four separate experiment sessions were held over the course of two days. The complete information treatment was used in one of these sessions, Session 1, while the private information treatment was used in the other three, Sessions 2-4. The players were only allowed to participate in a single session. The number of pairs in each session are shown in table 3.1.

While Lindahl and Johannesson carried out their experiment by passing paper forms between players, this experiment was, for practical purposes, carried out using a web application. In each of the four sessions, the players were, as in Lindahl and Johannesson's experiment, divided into two rooms, Room A and Room B, where they had access to computers on which the application ran.

Session	Participants	Pairs	Treatment
1	26	13	Complete
2	28	14	Private
3	24	12	Private
4	26	13	Private

Table 3.1: Number of experiment participants.

The web application randomly paired players in the two rooms in each round, while making sure that no two players ever met more than once. For each pair in each round, the application randomly drew a sum between SEK 100 and SEK 200 with the same characteristics as in the earlier experiment[1].

In each round, the application first showed player A the available sum. After player A submitted her claim, player B was notified and asked to submit her claim as well. In the complete information treatment, player B was also shown the available sum before having to submit her claim. The round ended after all players had made their claims and the application had shown them their results for the round[2].

Since players could only submit claims and see the results of their own play, and had no possibility to see or communicate with their opponents in the other room, total anonymity was guaranteed throughout the experiment. Furthermore, an experiment leader was always present in each of the two rooms, ensuring that there was no communication between players in the same room.

All players were given extensive instructions and had to correctly complete a quiz before the experiment rounds started, to ensure that everyone had understood the experiment setup[3].

Each experiment session took around one hour, with all players completing ten game rounds. Similar to Lindahl and Johannesson's experiment, all participants received a participation fee of SEK 50 at the end of the session, plus their payoffs from one of the ten game rounds, which was drawn randomly by the web application.

3.2 Hypotheses

Based on the results from the study made by Lindahl and Johannesson (2005) and other experimental studies mentioned in chapter 2, the hypothesis is that neither the complete information nor the private information experiment treatment will have the outcome predicted by standard game theory discussed in section 2.1 due to fairness considerations. Also, judging from the experiments described in section 2.4.3, the resource size should not have any significant effect on the outcome in either of the treatments.

[1] See equations 2.7-2.9 in section 2.2.1 for details.

[2] For screenshots of an example experiment round, see section B.3 in the Appendix.

[3] For full instructions and quiz, see Appendix B.

3.2.1 Complete Information

In the complete information treatment, just as in Lindahl and Johannesson's study and the studies mentioned in section 2.4.1, player A should start off with some first mover advantage over player B, albeit significantly less than what theory predicts.

When the game is repeated, we expect a similar outcome as in the studies described in section 2.6.1. Player A's payoff should either stay at the same level or show some tendency to increase over time, while the opposite applies to player B, as the two players discover each other's limits for what is considered acceptable. This means that the first mover advantage should either remain the same or increase when the game is repeated.

There should be no significant changes in efficiency. Under-exploitation is not a significant issue to begin with, and should not become one over time, as both players know the available resource size. Inefficiencies would then stem from resource breakdowns. These occur in the complete information case when people punish what they deem as unfair behaviour. Unless players start off very opportunistically (which should not be the case since every round could potentially count in terms of actual final payoff received), or what people consider fair changes rather quickly and drastically (also unlikely since the experiment lasts no longer than an hour), the same resource breakdowns would happen, regardless of how often a game is repeated.

3.2.2 Private Information

In the private information treatment, based on the results of Lindahl and Johannesson, we also expect an initial limited first mover advantage. With repeated rounds, however, the players are likely to split the resource evenly, as stated by Lindahl and Johannesson. This seems probable, considering that a significant share of players already chose the even split in the single-shot experiment. Also, as described in section 2.6.2, in the multi-round private information study by Mitzkewitz and Nagel (1993), first movers most commonly moderated their claims after a breakdown, while second movers typically remained firm, indicating that learning takes place towards the even split. Furthermore, as mentioned in section 2.2.3, the even split is the predicted equilibrium if using the inequality-aversion model of fairness by Fehr and Schmidt (1999).

The first mover advantage should decrease over time also because players learn to coordinate better. In e.g. Lindahl and Johannesson's experiment, 80% of the coordination failures were due to under-exploitation, primarily because player B was over-careful. These cases should decrease as player B gains more experience and starts demanding a fair share of the pie, instead of being over-careful. According to the EWA learning theory, described in section 2.5, players should learn from both their successes and failures to coordinate, and adapt their behaviour subsequently.

As a result, efficiency should increase over time, as under-exploitation decreases, while the level of breakdown should likely not increase, since most coordination failures were due to under-exploitation to begin with. The share of 50% claims made by player A should be high from the start of the experiment, as well as the share of imitation claims made by player B. These proportions should then increase as players learn and converge on the even split.

3.2.3 Comparison of Complete and Private Information

The first mover advantage should already from the start be smaller in the private information treatment than the complete information treatment, based on the single-shot experiments by Gustafsson (1999) and Lindahl and Johannesson. Player A, careful of avoiding resource breakdown, should make a smaller claim and receive less payoff in the private information treatment. We expect player B to initially also behave more carefully in the private information treatment, and claim and receive less payoff, just as in the experiment of Lindahl and Johannesson that we are mimicking. However, as the players learn, the difference between the treatments in regards to player B's payoff should decrease, as she starts claiming and receiving more.

The difference in first mover advantage between the treatments should further increase, with player A potentially gaining ground in the complete information treatment and most likely loosing ground in the private information treatment over time, as discussed in the previous sections. The share of player As claiming 50% of the pie and the share of player Bs imitating player A's claim should be higher in the private information treatment, as the even split is more likely there. This difference between the two treatments should increase with learning.

3.3 Statistics

Lindahl and Johannesson (2005) write that experiments similar to the one in this study and their own typically lead to results with non-normal distributions. In order to check if this is the case, we should use a normality test, like the Kolmogorov-Smirnov test, where the null hypothesis is that the sample tested is normally distributed. In the likely case of non-normal distributions, we should strictly use non-parametric tests to compare data within and between treatments, since these tests do not rely on the normality assumption.

Also, in the likely case that resource size does not matter for the outcome, we only need to analyse the players claims and payoffs expressed as a percentage of the sum, instead of both percentage and the absolute amount. To verify this, we should perform regressions controlling for the sum[4].

[4] For continuous variables such as claims and payoffs we need to use linear regression analysis, with equation form $dependent = \beta_0 + \beta_1 sum + \epsilon$. For binary variables such as breakdown and efficiency, we need to use logistic regression analysis, with equation form $logit(P) = \ln(P/(1-P)) = \beta_0 + \beta_1 sum + \epsilon$. If the null hypothesis of β_1 being significantly different from 0 can be rejected, then the absolute resource size has no effect on the outcome.

We will use the non-parametric Mann-Whitney independent samples test to compare the payoffs of players A and B in the first and last rounds of the experiment, in order to show any first mover advantage. The null hypothesis of this test is that the two samples compared are from the same distribution, and have equal means. If we can reject the null hypothesis, and player A has on average higher payoff than player B in a round, then there is a significant first mover advantage, else no significant first mover advantage exists.

In order to show the effects of learning, we can use the Wilcoxon related samples test to compare the first and last rounds, "before" and "after", in regards to the variables claim A, claim B, payoff A, payoff B, and difference in payoff between players A and B. The null hypothesis of the Wilcoxon test is that there is no difference between the "before" and "after" samples.

To test any differences within a treatment in the binary variables 50% claim by player A, imitation of player A's claim by player B, breakdown, and efficiency over time, we can use the related samples McNemar test on the results from round 1 and round 10. As with the Wilcoxon test, the null hypothesis is that there is no difference between the "before" and "after" samples.

In order to compare the two treatments, we need to use the Mann-Whitney independent samples test for the continuous variables, i.e. claims, payoffs, and the payoff difference between players A and B. We should analyse the data in both the first and last experiment rounds to see how players evolve.

The same applies for the binary variables, e.g. breakdown and efficiency. Here a suitable test is the Pearson χ^2 test of independence for frequencies, with the null hypothesis that the binary variable, i.e. the row variable, is unrelated to which treatment it is, i.e. the column variable. If the null hypothesis can be rejected, it means that there is a significant difference between the two treatments for the binary variable in question.

Hypothesis	Statistical test
Complete information:	
Payoff A $>$ Payoff B (round 1 and 10)	Mann-Whitney
Payoff A (round 10) $\geq$ Payoff A (round 1)	Wilcoxon
Payoff B (round 10) $\leq$ Payoff B (round 1)	Wilcoxon
Payoff diff. (round 10) $\geq$ Payoff diff. (round 1)	Wilcoxon
Breakdowns (round 10) $\approx$ Breakdowns (round 1)	McNemar
Efficiency (round 10) $\approx$ Efficiency (round 1)	McNemar
Private information:	
Payoff A $>$ Payoff B (round 1)	Mann-Whitney
Payoff A $\approx$ Payoff B (round 10)	Mann-Whitney
Payoff A (round 10) $<$ Payoff A (round 1)	Wilcoxon
Payoff B (round 10) $>$ Payoff B (round 1)	Wilcoxon
Payoff diff. (round 10) $<$ Payoff diff. (round 1)	Wilcoxon
50% splits by A (round 10) $>$ 50% splits by A (round 1)	McNemar
Imitations by B (round 10) $>$ Imitations by B (round 1)	McNemar
Breakdowns (round 10) $\leq$ Breakdowns (round 1)	McNemar
Efficiency (round 10) $>$ Efficiency (round 1)	McNemar
Complete vs. Private:	
Payoff diff. (Complete) $>$ Payoff diff. (Private)	Mann-Whitney
Payoff A (Complete) $>$ Payoff A (Private)	Mann-Whitney
Payoff B (Complete, round 1) $>$ Payoff B (Private, round 1)	Mann-Whitney
Payoff B (Complete, round 10) $\leq$ Payoff B (Private, round 10)	Mann-Whitney
50% splits by A (Complete) $<$ 50% splits by A (Private)	Pearson χ^2
Imitations by B (Complete) $<$ Imitations by B (Private)	Pearson χ^2

Table 3.2: Summary of this study's hypotheses and the test statistics used.

Chapter 4

Results

Below, we will look at the outcomes of the multi-round sequential CPR experiment, set up as described in chapter 3. As mentioned previously, the game was repeated in ten rounds, with randomly drawn pairs of players in each round. In total, we had in each round 13 pairs in the complete information treatment, and 39 pairs in the private information treatment. For an overview of the results corresponding to our hypotheses described in chapter 3, please see table 4.5 on page 29.

4.1 General Results

In order to reduce the risk of distorted results when carrying out statistical tests, we check the results and find that no obvious outliers exist, where players made negative claims or claims larger than the maximum resource size possible.

Testing the experiment results for normality, using the observations from round 1 in the private information treatment, we find that, according to the Kolmogorov-Smirnov test, we can reject the null hypothesis of normal distribution for the variables Claim A (%), Claim B (%), Payoff A (%), and Payoff B (%) ($\rho = 0.000$ for all variables tested). This means that the idea in section 3.3 to use non-parametric tests, which are not dependent on variables having a normal distribution, is valid. Also, linear and logistic regression analysis verify that the available resource size does not have any significant effect on the outcome[1].

With absolute resource size not significantly affecting the outcome, we continue our analysis using claims and payoffs of the players expressed as percentages of the sum. The average claims and payoffs in percentage of the resource size from the first and last rounds of both treatments are shown in table 4.1, together with the difference in payoff between players A and B, and the frequencies of 50% claims by player A, imitations by player B of the claim made by player A, resource breakdowns, and efficient splits of the available sum. For furter details on the development of the variables over time, please see Appendix A.

[1] The coefficient of the sum has $\rho > 0.1$ for all dependent variables analysed: round 1 and 10 – claim A (%), claim B (%), payoff A (%), payoff B (%), resource breakdown, efficiency.

	Complete		Private	
	Round 1	Round 10	Round 1	Round 10
Claim A (%)	53.3	61.1	54.0	49.4
Claim B (%)	53.6	38.9	46.6	48.1
Payoff A (%)	46.0	61.1	36.3	42.7
Payoff B (%)	46.3	38.9	28.7	40.9
Payoff diff. (p.p.)	-0.2	22.2	7.6	1.7
50% splits by A (%)	38.5	30.8	48.7	66.7
Imitations by B (%)	46.2	30.8	56.4	69.2
Breakdowns (%)	7.7	0	28.2	12.8
Efficiency (%)	92.3	100	28.2	61.5

Table 4.1: Average results for the first and last rounds.

4.2 Complete Information

As seen in figure 4.1a, at the start of the experiment, players A and B have similar claims and payoffs. Most claims by player A fall in the 50-70% range (8 of 13 cases, or 62%), the share being slightly lower than in the study by Lindahl and Johannesson (2005). The difference is mainly due to players making more modest claims. We have 4 claims (31%) under 50% and only 1 claim (8%) over 70%. Average claim by player A is 53%, compared to 55% in Lindahl and Johannesson's case. Player B's average claim is 54%, higher than Lindahl and Johannesson's 49%, as a result of player A making lower claims on average. Player B accepted all divisions of the pie efficiently, except in the single instance of player A claiming 95% of the pie. Average payoff for both players A and B in the first experiment round is 46%.

The Mann-Whitney independent samples test results in table 4.2 show that in round 1, there is no significant difference between player A and player B payoffs ($\rho = 0.812$), contrary to our hypothesis. However, in round 10, player A's payoff is significantly higher ($\rho = 0.000$), showing a clear first mover advantage. Here, player A and player B claim on average 61% and 39% respectively. Since we have efficiency in all cases, the players' payoffs equal their claims. As seen in figure 4.1b, player A's claims are predominantly in the 50-70% range (11 of 13 cases, or 92%), with no claims below 50% and only 2 cases (8%) over 70%. This is in line with Lindahl and Johannesson's single-shot experiment results, and also our hypothesis.

The results of the Wilcoxon test for related samples in table 4.3 show that the first mover advantage is significantly higher in round 10 than round 1 ($\rho = 0.013$). Player A's claims ($\rho = 0.075$) and payoffs ($\rho = 0.008$) increase over time, and the opposite goes for player B ($\rho = 0.008$ for claims, $\rho = 0.091$ for payoffs). This is in line with previous studies, where changes in players' behaviour, if any, are in player A's favour. Our observed increase in first mover advantage is, however, larger than in the studies in section 2.6.1, mainly due to a lower starting point, but the end results are similar.

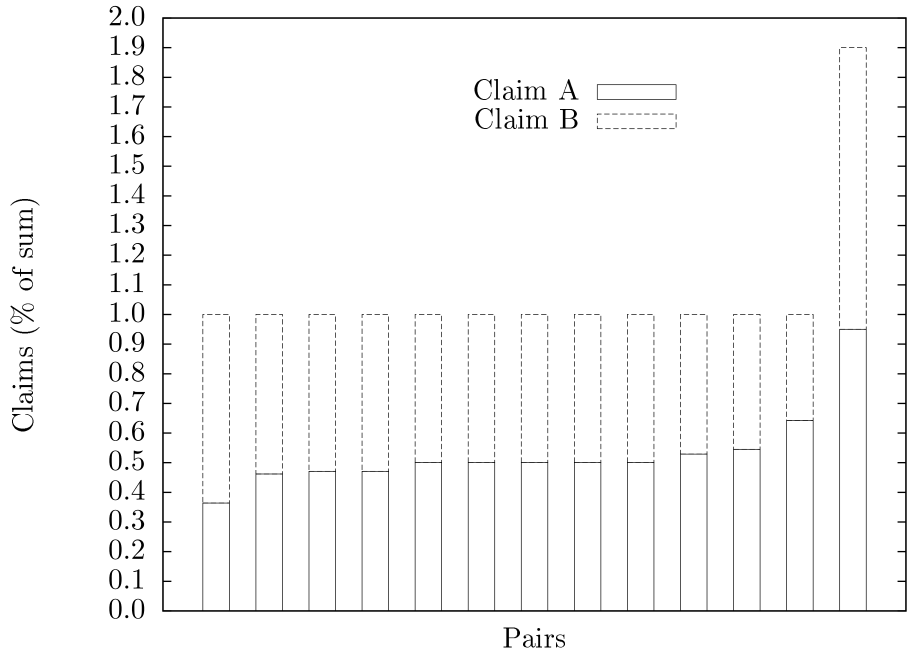

(a) Round 1

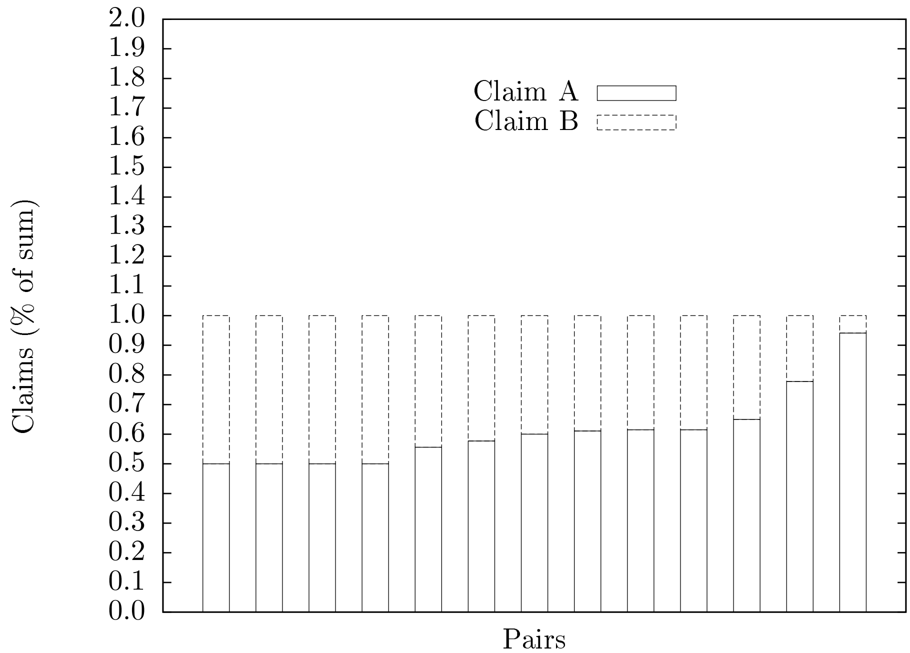

(b) Round 10

Figure 4.1: Claims made – Complete information.

	Complete			Private		
	Payoff A (%)	Payoff B (%)	ρ	Payoff A (%)	Payoff B (%)	ρ
Round 1	46.0	46.3	0.812	36.3	28.7	0.033
Round 10	61.1	38.9	0.000	42.7	40.9	0.313

Table 4.2: Payoff difference within treatments in the first and last rounds. (ρ values from the Mann-Whitney independent samples test.)

	Complete			Private		
	Round 1	Round 10	ρ	Round 1	Round 10	ρ
Claim A (%)	53.3	61.1	0.075	54.0	49.4	0.065
Claim B (%)	53.6	38.9	0.008	46.6	48.1	0.405
Payoff A (%)	46.0	61.1	0.008	36.3	42.7	0.171
Payoff B (%)	46.3	38.9	0.091	28.7	40.9	0.008
Payoff diff. (p.p.)	-0.2	22.2	0.013	7.6	1.7	0.097
50% splits by A (%)	38.5	30.8	1.000	48.7	66.7	0.118
Imitations by B (%)	46.2	30.8	0.625	56.4	69.2	0.383
Breakdowns (%)	7.7	0	1.000	28.2	12.8	0.146
Efficiency (%)	92.3	100	1.000	28.2	61.5	0.011

Table 4.3: Difference between the first and last rounds within treatments. (ρ values from the Wilcoxon related samples test for the continuous variables, and from the McNemar related samples test for the binary variables.)

All player As lowered their claims in the next round whenever they experienced resource breakdown, while 10 of 13 (77%) increased their claims if their previous round was successful. 3 of 13 players (23%) chose to stay firmly at 50%. Throughout the experiment, player B aimed for the efficient split. No significant difference can be seen in player B's likeliness to punish "unfair" splits. In 6 of the 10 rounds, we have resource breakdowns due to claims over 70% made by player A.

Using the McNemar test for related samples to compare our binary variables in round 1 and round 10, we see in table 4.3 that player As do not make more or less 50% claims as the experiment progresses ($\rho = 1.000$), neither do player Bs change the extent to which they copy player A's claim ($\rho = 0.625$). The frequency of resource breakdown does not change ($\rho = 1.000$), nor does efficiency increase as the game is repeated ($\rho = 1.000$), just as hypothesised.

4.3 Private Information

At the start of the experiment, players A and B claim on average 54% and 47% respectively, compared to 49% and 44% in the study by Lindahl and Johannesson (2005). Most commonly, as seen in figure 4.2a, player A claims exactly half of the

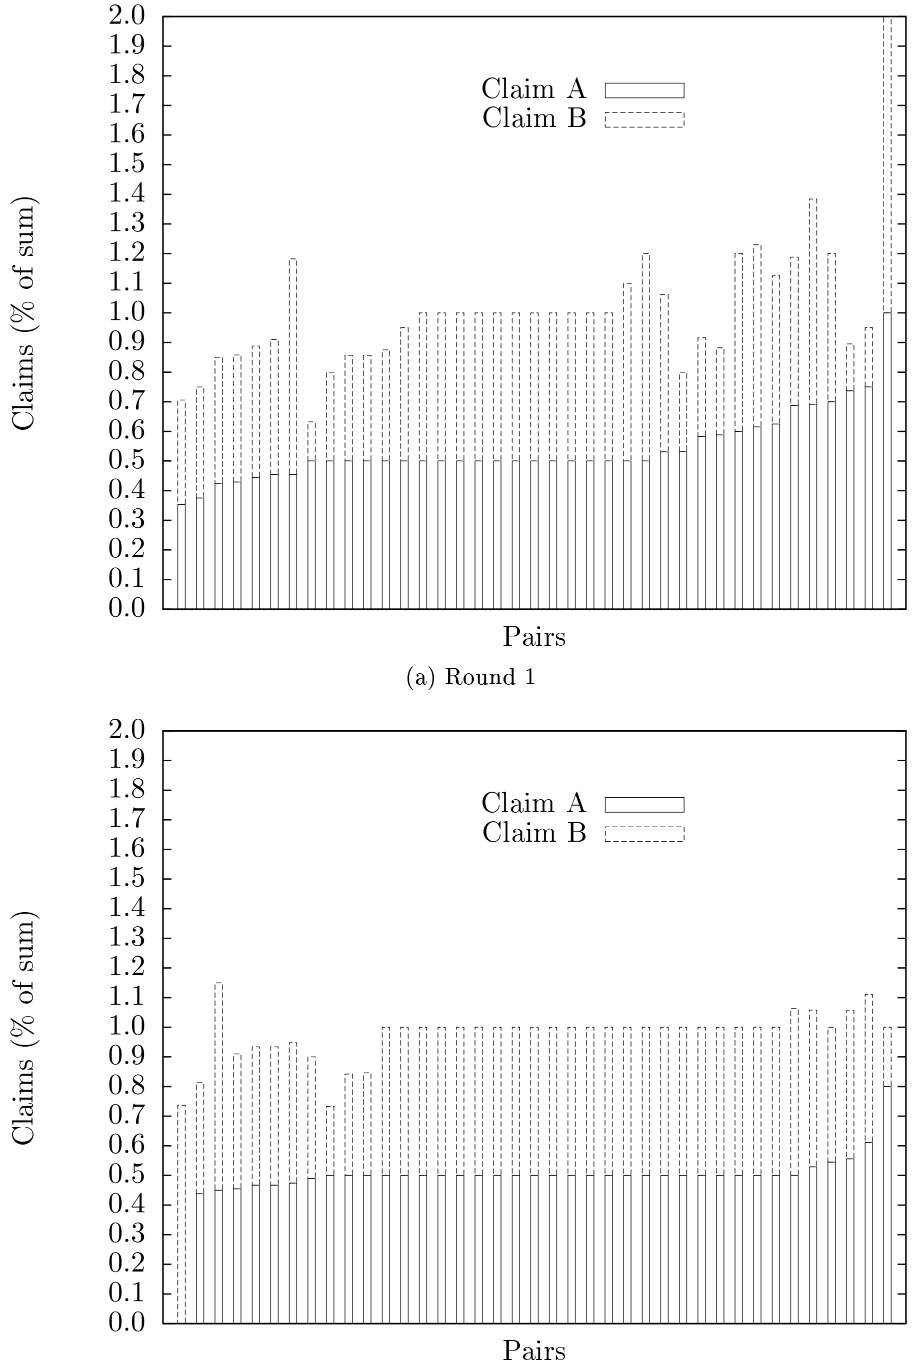

(a) Round 1

(b) Round 10

Figure 4.2: Claims made – Private information.

pie, similar to the previous study. Our result of 49% (19 of 39 cases) is, however, lower than the previous 62%, which fits with us having higher claims by player A on average. Player A's claim is lower than 50% in 7 cases (18%), and higher in 13 cases (33%). Player B mostly imitates player A's claim (22 of 39 cases, or 56%), in line with the previous study (50%). Player B's claim is higher than player A's in only 3 cases (8%), while in the remaining 14 cases (36%) it is lower. This means that $(19+22)/(39*2) = 53\%$ of all players started off aiming for an even split, the same figure being 57% in the previous single-shot experiment. Although not all cases led to efficiency, in all 11 efficient cases (28%), the players coordinated on the even split.

Of the 28 (72%) coordination failures, 11 cases (28%) were breakdowns, and 17 cases (44%) were under-exploitations. This is higher than the previous study's 60% failure rate. As a result, our first round average payoffs of 36% for player A and 29% for player B are lower than Lindahl and Johannesson's 42% and 37%. The Mann-Whitney test (see table 4.2) shows that player A's payoff is significantly higher than player B's in round 1 ($\rho = 0.033$). However, as hypothesised, this first mover advantage is no longer significant ($\rho = 0.313$) in round 10. The Wilcoxon test (see table 4.3) on the payoff difference between players A and B in round 1 and round 10 confirms that the first mover advantage is significantly lower in round 10 ($\rho = 0.097$), also in line with expectations.

In round 10, player A and player B's claims are on average almost the same, 49% and 48% respectively, and their average payoffs of 43% and 41% are also similar. According to the Wilcoxon test results in table 4.3, player A's claim is significantly lower at the end of the experiment compared to at the beginning ($\rho = 0.065$), and player B's payoff is, as predicted, significantly higher in the last round than the first ($\rho = 0.008$). However, player A's payoff does not change significantly during the experiment ($\rho = 0.405$), contrary to our hypothesis, and neither does player B's claims ($\rho = 0.171$). A likely reason to why player A's payoff does not decrease over time is that the initial, higher than expected claims made by some player As often led to breakdown, resulting in a relatively low average payoff to begin with.

The McNemar test results in table 4.3 show that the null hypothesis of no change over time cannot be rejected for the variables 50% claims by player A ($\rho = 0.118$), imitations by player B ($\rho = 0.383$), and breakdowns ($\rho = 0.136$). However, there are clear indications of less breakdowns, and more even splits by both players A and B as they gain more experience, which we hypothesised. In the final round, 26 of 39 player As (67%) claimed exactly 50% of the pie, and 27 of 39 player Bs (69%) copied player A's claim, meaning that $(26+27)/(39*2) = 68\%$ of the players aimed for the even split. Breakdowns were down from 28% to 13%, and efficiency was significantly up from 28% to 62% ($\rho = 0.011$).

In general, player A tended to make a lower claim in the next round after a breakdown (78% of all occurrences), claim the same share if the previous round resulted in an efficient split (77%), and either raise the claim (49%) or stay at the same level (30%) in the case of under-exploitation. Similarly, player B tended to lower the claim if the previous round resulted in a breakdown (90%), stay at the same level after an efficient split (55%), and make a higher claim after a case of

	Round 1			Round 10		
	Complete	Private	ρ	Complete	Private	ρ
Claim A (%)	53.1	54.0	0.593	61.1	49.4	0.000
Claim B (%)	53.6	46.6	0.100	38.9	48.1	0.004
Payoff A (%)	46.0	36.3	0.283	61.1	42.7	0.000
Payoff B (%)	46.3	28.7	0.001	38.9	40.9	0.151
Payoff diff. (p.p.)	-0.2	7.6	0.057	22.2	1.7	0.000
50% splits by A (%)	38.5	48.7	0.521	30.8	66.7	0.023
Imitations by B (%)	46.2	56.4	0.521	30.8	69.2	0.014
Breakdowns (%)	7.7	28.2	0.128	0	12.8	0.174
Efficiency (%)	92.3	28.2	0.000	100	61.5	0.008

Table 4.4: Difference between treatments in the first and last rounds. (ρ values from the Mann-Whitney independent samples test for the continuous variables, and from the Pearson χ^2 test for the binary variables.)

under-exploitation (76%). In the vast majority of cases where players chose to make the same percentage claim in the next round, the claim they stayed on was 50% (96% for player A, 98% for player B).

4.4 Comparison of Complete and Private Information

Table 4.4 shows the difference in results across treatments. The Mann-Whitney independent samples test results show that in round 1, the first mover advantage is significantly larger in the private information treatment than in the complete information treatment ($\rho = 0.057$). On average, the payoff difference between players A and B is near 0 with complete information, while player A has 8 p.p. higher payoff with private information. This is in contrast to our hypothesis, mainly due to the complete information resource splits in round 1 being surprisingly even (see section 4.2) compared to previous studies. In round 10, there is also a significant payoff difference ($\rho = 0.000$). However, here, the situation is as we predicted and reversed, so that the first mover advantage is now larger in the complete information treatment (22 p.p.) than in the private one (2 p.p.).

In regards to claims and payoffs, the Mann-Whitney test results show that player B is significantly better off in round 1 with complete information. Player B makes on average a higher claim (54% vs. 47%, $\rho = 0.100$) and receives a higher payoff (46% vs. 29%, $\rho = 0.001$). There is no significant difference in player A's claims and payoffs across the treatments in round 1. However, the results do show player A's payoff on average being higher in the complete information case (46% vs. 36%). This indicates that the initial difference in first mover advantage between the treatments is primarily due to player B being more careful under private information. In the last round, player A has on average both significantly higher claim (61% vs. 49%, $\rho = 0.000$) and payoff (61% vs. 43%, $\rho = 0.000$) in the complete information treatment than in the private one. Player B now makes significantly lower claims (39% vs. 48%, $\rho = 0.004$) in the complete information case, in opposite to the round 1 results. There is, however, no significant difference in player B's payoffs across the treatments (39% vs. 41%, $\rho = 0.151$). This is in line with our hypotheses of player B initially being better off with complete information, but with the difference disappearing over time, and player A being worse off with private information.

Results from the Pearson χ^2 test of independence show that efficiency is significantly higher in round 1 in the complete information treatment than in the private one (92% vs. 28%, $\rho = 0.000$). There are no significant differences for the other frequencies breakdowns, 50% claims by player A, and imitations by player B. However, in absolute values, we have in the private information case more 50% splits by player A (49% vs. 39%) and imitations by player B (56% vs. 46%). These results solidify over time, and in round 10, both the share of 50% claims by player A (67% vs. 31%, $\rho = 0.023$) and imitations by player B (69% vs. 31%, $\rho = 0.014$) are significantly higher with private information. There is at the end of the experiment little difference between the treatments in regards to breakdown ($\rho = 0.174$). However, the private information treatment remains significantly less efficient than the complete one (62% vs. 100%, $\rho = 0.008$).

Hypothesis	Result
Complete information:	
Payoff A $>$ Payoff B (round 1 and 10)	False (round 1), True (round 10)
Payoff A (round 10) $\geq$ Payoff A (round 1)	True ($>$)
Payoff B (round 10) $\leq$ Payoff B (round 1)	True ($<$)
Payoff diff. (round 10) $\geq$ Payoff diff. (round 1)	True ($>$)
Breakdowns (round 10) $\approx$ Breakdowns (round 1)	True
Efficiency (round 10) $\approx$ Efficiency (round 1)	True
Private information:	
Payoff A $>$ Payoff B (round 1)	True
Payoff A $\approx$ Payoff B (round 10)	True
Payoff A (round 10) $<$ Payoff A (round 1)	False
Payoff B (round 10) $>$ Payoff B (round 1)	True
Payoff diff. (round 10) $<$ Payoff diff. (round 1)	True
50% splits by A (round 10) $>$ 50% splits by A (round 1)	True in abs. terms ($\rho = 0.118$)
Imitations by B (round 10) $>$ Imitations by B (round 1)	True in abs. terms ($\rho = 0.383$)
Breakdowns (round 10) $\leq$ Breakdowns (round 1)	True in abs. terms ($\rho = 0.146$)
Efficiency (round 10) $>$ Efficiency (round 1)	True
Complete vs. Private:	
Payoff diff. (Complete) $>$ Payoff diff. (Private)	False (round 1), True (round 10)
Payoff A (Complete) $>$ Payoff A (Private)	True
Payoff B (Complete, round 1) $>$ Payoff B (Private, round 1)	True
Payoff B (Complete, round 10) $\leq$ Payoff B (Private, round 10)	True ($\approx$)
50% splits by A (Complete) $<$ 50% splits by A (Private)	True
Imitations by B (Complete) $<$ Imitations by B (Private)	True

Table 4.5: Summary of this study's hypotheses and the experimental results.

Chapter 5

Discussion and Conclusions

In this final chapter, we conclude this thesis by discussing the results presented in the previous chapter 4. We relate them to previous studies, both single-shot and multi-round, and look at them from learning and fairness perspectives. Finally, we summarise the main findings of the thesis.

5.1 General Comments on the Results

Considering the complexity of real life situations compared to the simple settings in experimental studies, one could argue about the relevancy of the results found in this study. In reality, we have issues with non-anonymity, leading to e.g. reputation building behaviour, possible confusion about e.g. the size of the other party's claim, etc. However, by distilling away all these issues, and keeping our experiment simple and anonymous, we can concentrate on people's underlying behaviour, allowing us to e.g. zero in on the learning effects, without other factors muddling the results.

5.2 Comparison with Previous CPR Studies

In the results chapter, we have already gone through the main differences between this study and the one by Lindahl and Johannesson (2005), and also where they deviate from our hypotheses based on findings from the previous studies described in chapter 2.

In general, the results correspond well with the expected outcomes, with the major exception being that we do not have a significant first mover advantage in the complete information treatment in the first experiment round. A possible key reason for this deviation could be our relatively small sample size of 26 participants in the complete information case, compared to e.g. the 94 participants Lindahl and Johannesson (2005) had in their study. The smaller sample size makes our results less solid than the very strong findings from previous complete information studies with larger samples, where there is a clear first mover advantage. However, our results as a whole still correspond well with previous studies. In our case, had player A on average only made a slightly higher claim, we most likely would have had the same results as Lindahl and Johannesson. Since player B in the first round of our study accepted all pie shares in the 35-65% range, they would likely have allowed the 10-20 p.p. first mover advantage seen elsewhere. Also, over time, our results develop into a very close match with those in previous studies, both single-shot and multi-round, making them still highly interesting to look at.

In the private information case, our sample size of 78 participants should be sufficiently large for all findings to be much more solid. It is similar to e.g. the 80 participants in the multi-round private information study carried out by Mitzkewitz and Nagel (1993) and not too far from the 154 participants in Lindahl and Johannesson's study. Our findings from the first experiment round are in general consistent with those of Lindahl and Johannesson. There are some slight differences due to player A in our case on average making a somewhat higher claim, leading to a higher share of breakdowns. However, as in Lindahl and Johannesson's case, we have that players most commonly aim for the even split, with player A most likely to claim 50% of the pie, and player B most often imitating player A's behaviour. In regards to our final results after repeated experiment rounds, we unfortunately have very little to compare with data-wise. It can, however, be noted that they seem to verify Lindahl and Johannesson's predictions of an efficient, even split gaining ground, although more rounds could be needed to further firm up these results.

5.3 Learning

In the private information sequential CPR experiment, our results show that learning does occur. Efficiency increases, the initial first mover advantage disappears, player A lowers her claim to approximately the same level as player B, and player B's payoff increases. Similar to Mitzkewitz and Nagel' study described in section 2.6.2, our players adjusted their claims based on the previous round. We see behaviour matching the EWA learning model (see section 2.5), where people learn from both chosen and unchosen strategies. Players tended to lower their claims after breakdowns, stay at the same level if achieving an efficient split, and either raise or stay at the same level after under-exploitations.

We also observe learning effects in the complete information treatment, somewhat unexpectedly considering previous studies (see section 2.6.1). The first mover advantage, which was not significant in the first experiment round, increases over time, with player A claiming and receiving more, and player B claiming and receiving less. The change in player B's claims and payoffs with repetition is primarily a result of player A getting bolder with repetition after a weak start. After a successful split, player A tended to raise her claim, while lowering it otherwise. Player B's underlying behaviour of accepting all reasonably “fair” pie divisions did not, however, change significantly, with player A being punished for excessive claims in most cases, independent on when it occurred during the experiment.

5.4 Fairness

For both information treatments, we can safely state that the outcomes are very different from what standard game theory predicts (see section 2.1), again confirming previous studies' findings that players are affected by fairness considerations. In both treatments, the first mover does not come close to claiming the entire pie, and the second mover receives a payoff significantly higher than a token amount or 0.

Mitzkewitz and Nagel (1993) write that in the complete information case, the even split is an obvious standard of fairness. Looking at our results over time, we see how players start off close to the even split, but move into a 60/40 split with repetition. Here, the self-serving bias we mentioned in section 2.4.1 likely plays a role, encouraging the first movers to claim a somewhat higher share of the pie. Still, very few claim more than 70% of the available resource, rather limiting the bias, and we do not end up far from Mitzkewitz and Nagel' "fairness standard".

In the private information setting, Lindahl and Johannesson (2005) state that applying the inequality-aversion fairness model of Fehr and Schmidt (1999) yields the even split as the predicted outcome. Although parts of the results at the end of our experiment are still indicative, due to significance values over the 10% threshold, all trends point towards this "fair" outcome. The payoff difference between players is close to 0 in the final experiment round, an increasing majority of players played according to the even split, and efficiency has increased significantly. Claiming exactly half of the pie is the most common player A behaviour throughout the experiment, and the same is true for player B imitating player A's claim. Also worth noting is that the majority of all player As who successfully avoided breakdown with a 50% claim stayed on that level for the remainder of the experiment, without trying to further improve their situation. This is in line with Fehr and Schmidt's assumption of people not only disliking being worse off, but also wanting to avoid being better off than others.

5.5 Conclusions

In chapter 1 we stated the purpose of this thesis as finding out if two players in a repetitive setting, when the first mover has more information about the resource size, will learn to efficiently share a common resource more equally than when they both know the resource size. We carried out multi-round experiments closely mimicking the single-shot sequential CPR experiment by Lindahl and Johannesson (2005) to see if their prediction of the even split being the long-term equilibrium in the private information setting becomes true.

In the complete information setting, after a careful start by player A leading initially to a more equal split than expected, our results quickly develop into what previous experimental research predicts, and we end up with a solid 60/40 split and complete efficiency.

In the private information setting, we start off in a similar fashion as Lindahl and Johannesson, with about half of the players aiming for the even split. And just as in the previous single-shot study, the vast majority of splits are inefficient, mainly due to under-exploitation as players are over-careful, and there is initially a first mover advantage. With repeated rounds, we see clearly how players learn and adjust their behaviour according to previous results. The initial first mover advantage shrinks and we no longer have a significant payoff difference between the players in the final round. We also have significantly higher efficiency, with the majority of players being successful in splitting up the resource exactly.

Comparing the two settings, at the end of our experiment, first movers receive significantly higher payoffs with complete information, while second movers are about as well off in both settings. The first mover advantage is clearly higher in the complete information case.

Thus, we conclude that players do learn in a repetitive setting, becoming more efficient over time and splitting the common pool resource more equally than if they had both known the available resource size. This conclusion is highly interesting, considering that, typically, the party with more information has the advantage. In our case, however, the first mover actually loses ground due to having access to more information.

Ideally, we would have liked to see players reaching the even split in term of payoffs, and not only in term of claims as in our case. This would have meant that not only is the player with more information worse off for it, but also that the player with less information actually gains from her ignorance. We see indications of this happening, with the share of players aiming for the even split increasing, and players tending to stay at the even split strategy, after once having discovered it to be successful. However, although most players had learned by the final experiment round, there were still cases of under-exploitation and breakdowns lowering players' average payoffs. We suspect increasing the number of rounds played from 10 to e.g. 30 would provide us with solid experimental results showing that the even split is the long-term equilibrium in private information CPR games.

References

L. Babcock and G. Loewenstein. Explaining bargaining impasses: the role of self-serving biases. *Journal of Economic Perspectives*, 11(1):109–126, Winter 1997.

G. E. Bolton and R. Zwick. Anonymity versus punishment in ultimatum bargaining. *Games and Economic Behavior*, 10(1):95–121, July 1995.

C.F. Camerer, editor. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press, 2003.

C.F. Camerer and T. Ho. Experience-weighted attraction learning in normal-form games. *Econometrica*, 67(4):827–874, July 1999.

C.F. Camerer, T. Ho, and K. Chong. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. *Journal of Economic Theory*, 104(1):137–188, May 2002.

C.F. Camerer, G. Loewenstein, and M. Rabin, editors. *Advances in Behavioral Economics*. Princeton: Princeton University Press, 2003.

L.A. Cameron. Raising the stakes in the ultimatum game: Experimental evidence from indonesia. *Economic Inquiry*, 27(1):47–59, January 1999.

I.-K. Cho and D.M. Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–222, May 1987.

R. Croson. Information in ultimatum games: An experimental study. *Journal of Economic Behavior and Organization*, 30(2):197–212, August 1996.

E. Fehr and K.M. Schmidt. A theory of fairness, competition and cooperation. *Quarterly Journal of Economics*, 114(3):817–868, August 1999.

M. Gustafsson. The role of asymmetric information in sequential resource dilemmas with unknown resource size. Technical Report 29:1, Department of Psychology, Göteborg University, 1999.

G. Hardin. The tragedy of the commons. *Science*, 162(5):1243–1248, December 1968.

T.-H. Ho, C. Camerer, and K. Weigelt. Iterated dominance and iterated best response in experimental "p-beauty contests". *American Economic Review*, 88(4): 947–969, September 1998.

E. Hoffman, K. McCabe, K. Shachat, and V.L. Smith. Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, 7(3):346–380, November 1994.

E. Hoffman, K. McCabe, and V.L. Smith. On expectations and monetary stakes in ultimatum games. *International Journal of Game Theory*, 25(3):289–301, September 1996.

D. Kahneman, J.L. Knetsch, and R.H. Thaler. Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review*, 76(4):728–741, September 1986.

R.P. Larrick and S. Blount. The claiming effect: Why players are more generous in social dilemmas than in ultimatum games. *Journal of Personality and Social Psychology*, 72(4):810–825, April 1997.

T. Lindahl and M. Johannesson. Private information in common pools. In T. Lindahl, Ph.D. thesis, *Strategic and Environmental Uncertainty in Social Dilemmas*, pages 13–42. Stockholm: Stockholm School of Economics, 2005.

J.A. List and T.L. Cherry. Learning to accept in ultimatum games: Evidence from an experimental design that generates low offers. *Experimental Economics*, 3(1): 11–31, June 2000.

M. Mitzkewitz and R. Nagel. Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory*, 22(2):171–198, June 1993.

E. Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press, 1990.

A. Rapoport, J.A. Sundali, and R.E. Potter. Ultimatums in two-person bargaining with one-sided uncertainty: Offer games. *International Journal of Game Theory*, 25(4):475–494, December 1996.

A.E. Roth, V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir. Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh and Tokyo: An experimental study. *American Economic Review*, 81(5):1068–1095, December 1991.

R.L. Slonim and A.E. Roth. Learning in high stakes ultimatum games: An experiment in the Slovak Republic. *Econometrica*, 66(3):569–596, May 1998.

P. Straub and K. Murnighan. An experimental investigation of ultimatum games: Information, fairness, expectations, and lowest acceptable offers. *Journal of Economic Behavior and Organization*, 27(3):345–364, August 1995.

Appendix A

Further Experiment Results

Figures A.1 and A.3 below show how the average claims and payoffs developed over time in both experiment treatments. In the complete information treatment, the gap between the two players increases, with player A's situation improving and player B's situation worsening over time. In the private information treatment, the two players' average claims converge towards the even split, and both are clearly better off towards the end of the experiment than at the beginning.

The development of efficiency, breakdowns, 50% claims by player A and imitations by player B are shown below in figures A.2 and A.4. In the complete information treatment, there are no clear trends in efficiency and breakdown over time. Statistics show no significant change in 50% claims by player A or imitations by player B, however the absolute values of these frequencies seem to decrease. In the private information treatment, we have increasing efficiency (statistically significant at the 10% level when comparing round 1 and round 10). Although not equally significant, breakdowns seem to decrease over time, and both players seem to increasingly aim for the even split.

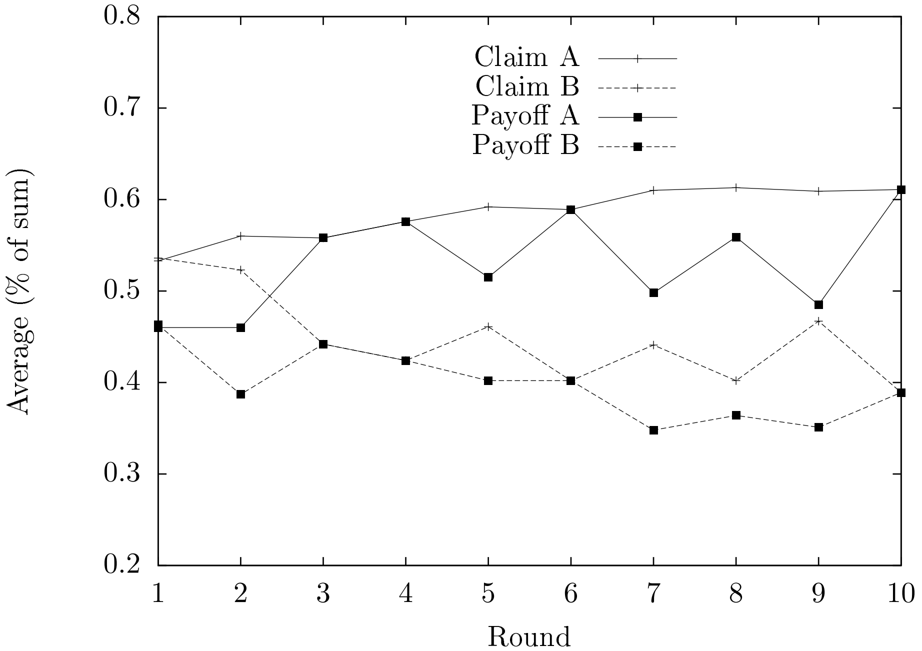

Figure A.1: Average claims and payoffs over time – Complete information.

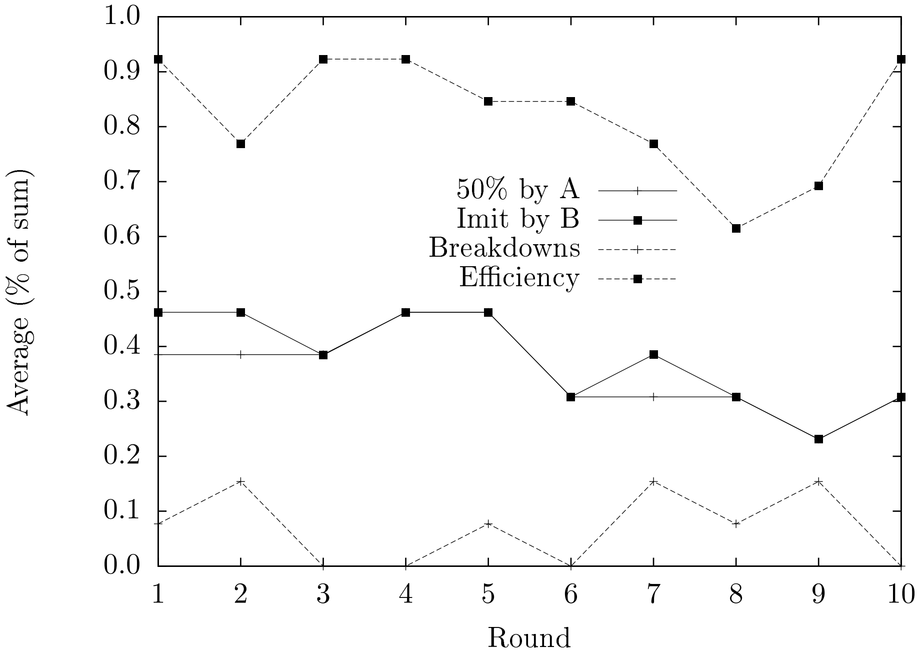

Figure A.2: Frequencies over time – Complete information.

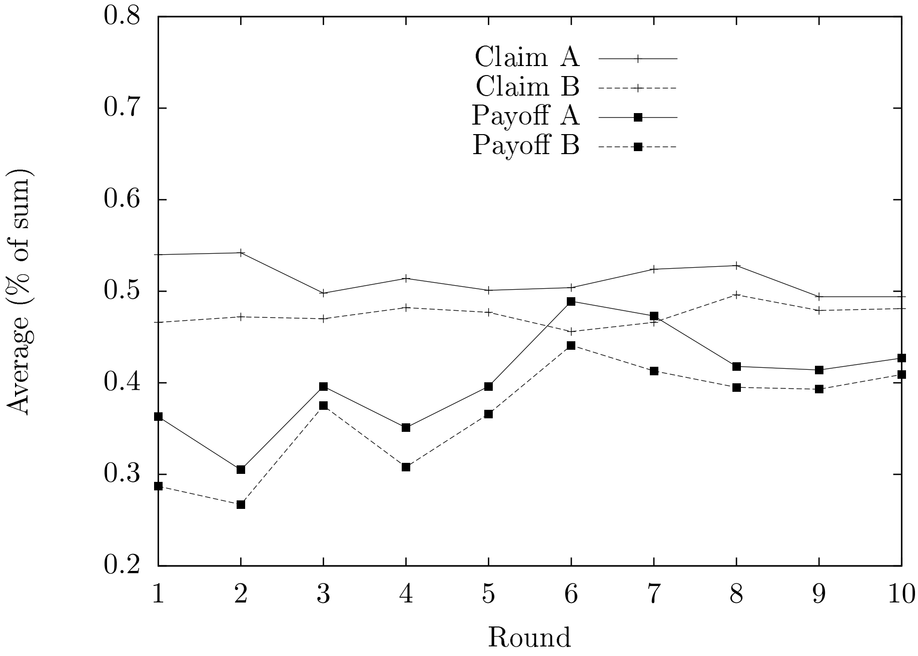

Figure A.3: Average claims and payoffs over time – Private information.

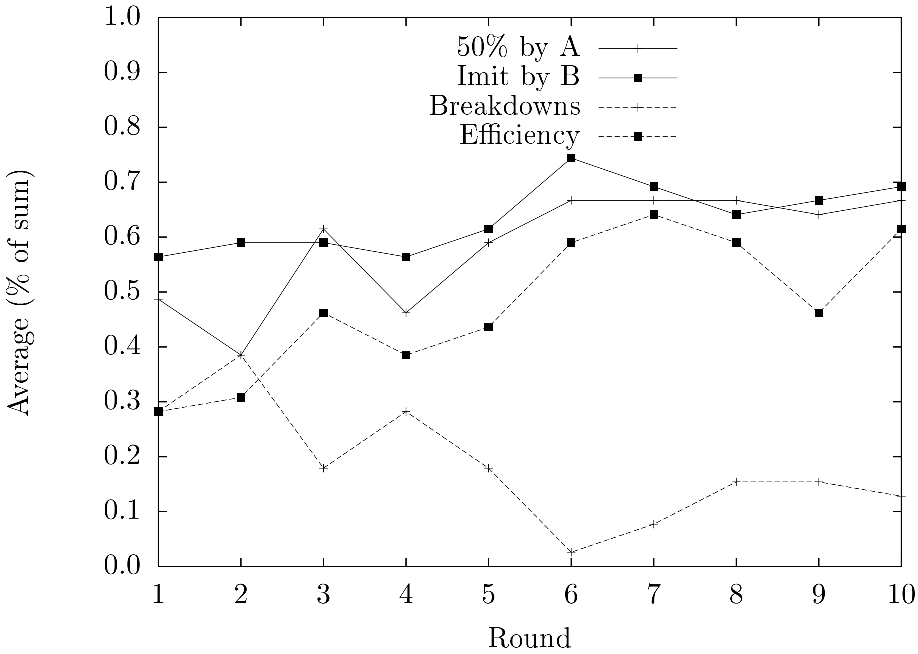

Figure A.4: Frequencies over time – Private information.

Appendix B

Experiment Setup Details

Below is the information shown to all players in room A in a private information experiment round, translated to English. The only difference in the instructions given to players in rooms A and B was in the part of the text describing which room they were in. The difference in the instructions given to players in the private and the complete information treatments was that in the private treatment, it was pointed out clearly that player B would not know the available sum before the claims were already made, while in the complete treatment it was stated that both player A and player B would know the available sum beforehand.

B.1 Instructions

Thank you for participating in this experiment. As a compensation for your participation, you (and the other participants) will receive SEK 50. This compensation, plus your potential payoff from this experiment, will be paid out at the end of the experiment.

Instructions

The experiment will last approximately 1.5 hours and consist of **10 different decision rounds, which are independent of each other**. In each round, you will be paired **randomly** with someone who is sitting in another room. You will never be matched to the same person twice. You will not find out who the people you were matched with were, neither during nor after the experiment. Both rooms (A and B) have the same number of people. **This is room A.** Each person in room A and room B have received these instructions.

In each round, each pair will decide how to distribute a given sum of money between the person in room A ("A") and the the person in room B ("B"). For each pair, the sum available will vary between SEK **100-200**. For each round and each pair, one of the following possible amounts will be assigned randomly (with equal probability for each amount): SEK 100, 110, 120, 130, 140, 150, 160, 170, 180, 190, or 200.

A will know which sum has been assigned to the pair. B will not know this sum. B will only know that the probability of the sum being any of the possible sums between SEK 100-200 is the same.

This experiment consists of 10 different rounds. At the end of the experiment, one of these rounds will be drawn at random. The probability for any given round to be drawn is the same. You will receive the payoff you made in the round drawn.

This is how a round works:

In step 1, A makes a claim on the available sum. B is then shown this claim, and makes her claim on the sum in step 2. (**Note that B will make her claim without knowing the size of the available sum.**) There are two possible outcomes to the round.

1. If the sum of your and your partner's claims exceeds the available sum, then neither of you receive anything (in the case this is the round drawn).
2. If the sum of your and your partner's claims is less than or equal to the available sum, then you will both receive your respective claims (in the case this is the round drawn).

This is how to play each round:

In step 1, the available sum is shown on A's screen. A fills out her claim in the text box shown and clicks "Submit". In step 2, A's claim is shown on B's screen (**B will not be shown the available sum**). B fills out her claim in the text box on her screen and clicks "Submit". Then, both A and B will be shown the available sum in the current round, both claims, and their respective payoffs.

Your own results from earlier rounds are shown on the right side of the screen.

When everyone has finished 10 rounds, each participant will receive her payoff from the round drawn, plus the SEK 50 participation fee.

You must avoid the following:

- Reloading the page. The experiment page will refresh automatically.
- Going to another page than the one shown at any given time.
- Going back to an earlier page.
- Interrupting the experiment before finishing all rounds.
- Speaking to other participants.

Feel free to contact the experiment leader if you have any questions or thoughts.

B.2 Quiz

The following online quiz was given after the experiment participants had read all instructions (see figure B.1 for a screenshot). The participants were allowed to move on to the experiment rounds only after passing all the quiz questions.

1. What are the payoffs of A and B, if the available sum is 160, A claims 60 and B claims 90?
2. What are the payoffs of A and B, if the available sum is 180, A claims 120 and B claims 60?
3. What are the payoffs of A and B, if the available sum is 130, A claims 100 and B claims 40?
4. What information does A and B have at the beginning of each round?
   - A knows the available sum, but not B
   - B knows the available sum, but not A
   - Both A and B know the available sum
   - Neither A nor B knows the available sum
5. Which of the following is a possible sum?
   - 130
   - 155
   - 210
6. My results from previous rounds affect later rounds.
   - True
   - False
7. In each round, I will meet a new, randomly drawn person in the other room.
   - True
   - False
8. Which of the following am I allowed to do during the experiment?
   - Speak to other participants
   - Contact the experiment leader
   - Reload the page
   - Go to another page than the one currently shown

B.3 Screenshots of an Example Round

Figures B.2-B.4 show an example experiment round for player A and player B.

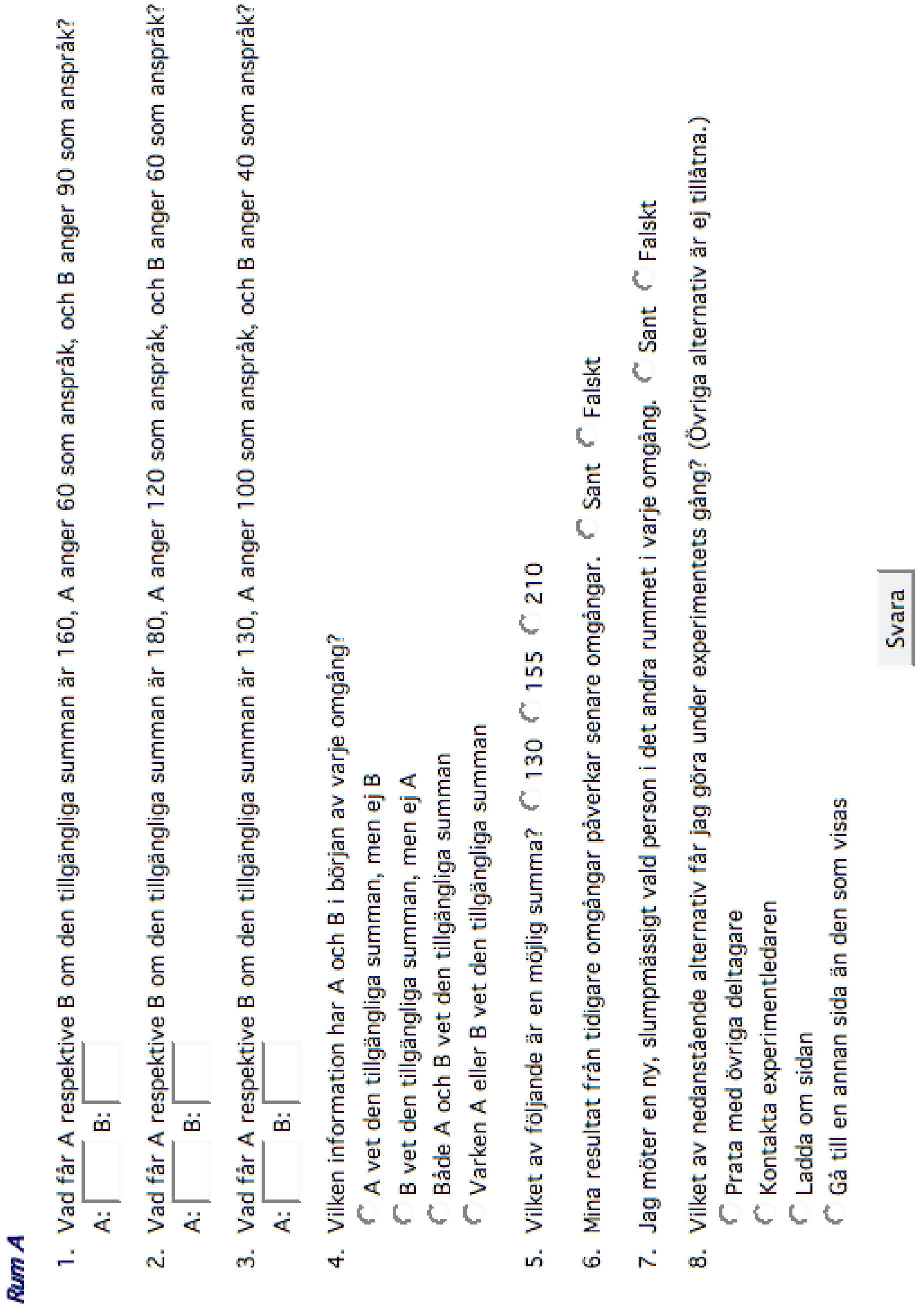

Figure B.1: Control quiz.

Figure B.2: Waiting for A to make her claim. Available sum 130.

Figure B.3: Waiting for B to make her claim. B sees A's claim of 60.

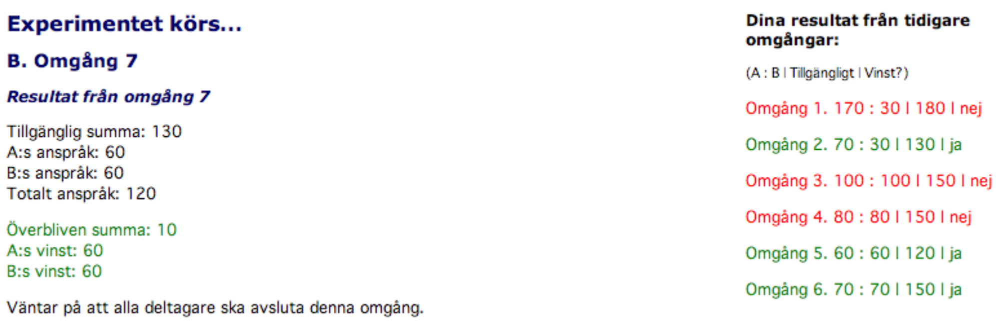

Figure B.4: Both players have made their claims and are shown the round results.