Essays on Cooperation and Reciprocity

Thesis by

Nilanjan Roy

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2013

(Defended 14 May 2013)

To my teachers, family and friends

Acknowledgements

I thank everyone in the Division of Humanities and Social Sciences at the California Institute of Technology for providing me with a productive and enjoyable research environment. I would like to express my deepest gratitude to my adviser, Thomas Palfrey, for his excellent guidance, care, patience and belief. This thesis would not have been possible without his unconditional and undying support. I am also grateful to Peter Bossaerts, Federico Echenique, John Ledyard, Charles Plott, Howard Rosenthal and Matthew Shum for their care and encouragement. I have learnt a lot from their knowledge and experience.

A special thanks to Laurel Auchampaugh for always being there, Chris Crabbe for developing the experimental software, both Walter Yuan and Estela Hopenhayn at UCLA for helping me with the experiments, Victoria Mason for making sure that I have enough petty cash to run my experiments on time and Eric Bax for the enjoyable time while I was at Yahoo! The funding from the Gordon and Betty Moore Foundation, National Science Foundation and the California Institute of Technology is gratefully acknowledged.

It would be unjust if I did not thank all my teachers as well as faculty members at my previous institutions, especially Arunava Sen at the Indian Statistical Institute. A lot of credit must go to all of them for selflessly helping me in my endeavors and believing in me throughout. Finally, I would also like to thank my family and friends for being there in this journey of happiness, sadness, anger and loneliness. They have been kind enough to forgive

me for all faults of mine and embrace me the way I am.

Abstract

This dissertation comprises three essays that use theory-based experiments to gain understanding of how cooperation and efficiency is affected by certain variables and institutions in different types of strategic interactions prevalent in our society.

Chapter 2 analyzes indefinite horizon two-person dynamic favor exchange games with private information in the laboratory. Using a novel experimental design to implement a dynamic game with a stochastic jump signal process, this study provides insights into a relation where cooperation is without immediate reciprocity. The primary finding is that favor provision under these conditions is considerably less than under the most efficient equilibrium. Also, individuals do not engage in exact score-keeping of net favors, rather, the time since the last favor was provided affects decisions to stop or restart providing favors.

Evidence from experiments in Cournot duopolies is presented in Chapter 3 where players indulge in a form of pre-play communication, termed as revision phase, before playing the "one-shot" game. During this revision phase individuals announce their tentative quantities, which are publicly observed, and revisions are costless. The payoffs are determined only by the quantities selected at the end under real time revision, whereas in a Poisson revision game, opportunities to revise arrive according to a synchronous Poisson process and the tentative quantity corresponding to the last revision opportunity is implemented.

Contrasting results emerge. While real time revision of quantities results in choices that are more competitive than the static Cournot-Nash, significantly lower quantities are implemented in the Poisson revision games. This shows that partial cooperation can be sustained even when individuals interact only once.

Chapter 4 investigates the effect of varying the message space in a public good game with pre-play communication where player endowments are private information. We find that neither binary communication nor a larger finite numerical message space results in any efficiency gain relative to the situation without any form of communication. Payoffs and public good provision are higher only when participants are provided with a discussion period through unrestricted text chat.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis comprises three chapters that use non-cooperative game theory-based experiments to gain understanding of how cooperation and efficiency is affected by various variables and institutions in different types of strategic interactions prevalent in our society. Three types of "games" are analyzed to capture the notion of "cooperation" in different contexts: indefinite horizon dynamic favor-exchange games with private information, finite horizon dynamic Cournot revision games, and one-shot threshold public good games with private information.

Experimental evidence concerned with indefinite horizon two-person dynamic favor exchange games with private information is presented in Chapter 2. This chapter contributes to a growing literature on experiments on repeated games with imperfect public monitoring or with private information. Using a novel experimental design to implement a dynamic game with a stochastic jump signal process, this study provides insights into a relation where cooperation is without immediate reciprocity. Individuals interact in pairs in continuous time and occasionally one of them receives a privately observed opportunity to provide a favor to her partner. Only one player is in a position to do a favor at a given instant. The cost of doing a favor is strictly less than the benefit to the recipient, so that, it is socially optimal to always provide a favor.

The payoffs realized by the participants are considerably lower than what they could have achieved under the most efficient perfect public equilibrium. This is a robust finding, even under the situation where the benefit is very high compared to the cost of providing favors and when opportunities to provide favors arrive very fast. Next, individuals do not engage in exact score-keeping of cumulative favors, as proposed in the literature. Rather, their focus is more on the time since last favor provided by her and her partner suggesting that behavior is explained more by the notion of "What have you done for me lately?". Finally, efficiency is enhanced by either letting individuals perfectly observe the opportunities received by their partners or by allowing them to communicate before the start of a bilateral relation.

Chapter 3 implements Cournot duopolies in the laboratory where players indulge in a form of pre-play communication termed as a revision phase before playing the "one-shot" game. During this revision phase individuals announce their tentative quantities, which are publicly observed, and revisions are costless. The payoffs are determined only by the quantities selected at the end in a real time revision game; while in a Poisson revision game, opportunities to revise arrive according to a synchronous Poisson process and the tentative quantity corresponding to the last revision opportunity is implemented.

The act of revising actions before the play of the underlying "one-shot" game is prevalent in many areas. It is very common to announce plans for the production of cars few months ahead of actual production in the US motor vehicle and aircraft industries. The trade journal *Ward's Automotive Reports* publishes the firm's announcements of their plans for monthly U.S. production of cars as early as six months before actual production and these plans could be continuously revised until the end of the target month. Revision of actions is a phenomenon that is practiced in some of the financial markets such as Nasdaq and Euronext. Prior to opening of the market, participants are allowed to submit orders which can be continuously withdrawn and changed until the opening time. During the entire pre-opening phase these orders and resulting prices are publicly posted but only the orders that are still posted at the opening time are binding and are executed. Traders do not always manage to withdraw and submit new orders simultaneously due to technological and other reasons. Other examples would include situations where communication and implementation occur at different times, possibly due to delays. The Poisson revision game implemented in this study is a stylized representation of these types of situations. It models these inefficiencies or imperfections in implementing the intended choices. At the other extreme, real time revision games implement intended actions without any imperfection.

We find that when individuals play the Cournot duopoly game with real time revision but without observing the rival firm's revisions then play converges to the Cournot-Nash outcome. Quantity choices are even more competitive than Cournot-Nash in the presence of a real time revision phase with perfect observability of rival's revisions. In contrast, when revision plans are mutually observable and revision opportunities arrive according to a Poisson process, significantly lower quantities are selected than in the choices implemented under the real time revision games. This shows that partial cooperation can be sustained even when individuals interact only once.

The dynamics of revisions gives rise to interesting observations. First, while a real time revision game is characterized only by late upward quantity adjustments, the Poisson revision games are characterized both by initial downward adjustments and also by the late upward revisions. Second, the quantity adjustments during the initial period of the revision phase show that individuals imitate their opponent's desired quantity choices, whereas, behavior towards the end of the revision phase can be explained by the best response to the competitor's desired output.

In Chapter 4, we investigate the effect of varying the message space in a public good game with cheap talk where player endowments are private information. The game involves the production of a discrete public good by soliciting voluntary contributions from group members. The production technology is such that one unit of input is required from at least some fraction of the group. The minimum number of inputs needed to provide the good is referred to as the threshold. The inputs may be thought of as being a fixed contribution of time, effort, or money to some common goal. The cost of the input is privately observed. No side-payments are allowed. Inputs are non-refundable; consequently, it is not possible for a player to reclaim a contribution if either the threshold is not reached (under-contribution) or if the threshold is exceeded (over-contribution).

Threshold games present severe coordination problems. There is a weak incentive to free ride in the sense that each member of the group would prefer other members to contribute to the public good provision. The no-refund aspect of the game reinforces this free-riding incentive. There is another aspect of the coordination problem which arises because of the presence of incomplete information and heterogeneous preferences. The ex-ante efficient solution is for just enough of the members of a group to contribute their unit of input as is needed, and for the contributors to be the ones with relatively low opportunity costs of inputs while the free riders are the ones who have relatively high valuations for inputs. Without communication, it is not possible for players to know each others' relative valuations. However, with sufficient communication, it becomes feasible to coordinate decisions in a way that produces the desired outcome. The problem is then that the members of a group will communicate strategically to induce other members to bear the cost of producing the good.

We experimentally implement three different communication structures prior to the decision move: (a) a simultaneous exchange of binary messages, (b) larger finite numerical message space and (c) unrestricted text chat. We find that neither the one-time simultaneous exchange of binary messages meant to reveal one's intention to contribute, nor the one-time display of a numerical message is enough to create any efficiency gain relative to the situation without any form of communication. Only when participants are provided with the opportunity to engage in text chat in an unrestricted fashion are earnings and public good provision significantly higher. Thus, it seems that without a 'common language' there is no 'obvious' way to interpret the binary or the numerical messages among subjects. Unrestricted chat gives them the opportunity to understand and interpret each others' intentions and messages. Also, gains relative to the situation with no communication are higher when it is common knowledge that everyone has costs less than the 'public' benefit. This is intuitive because in this situation the question is not whether the good "should" be provided but rather which of the two people in a group are going to provide it. In contrast, in the sessions where it is possible for individuals to have higher costs than the 'public' benefit, first one needs to figure out whether a good "should" be provided and hence, there are more chances of mis-coordination and mis-representation.

While all the three chapters are linked through the usage of data from laboratory experiments based on non-cooperative game theory, there are other aspects that tie them (or at least a subset of them) together. First, a very crucial aspect of the models considered in Chapters 2 and 4 is the presence of private information. The opportunities to provide a favor are privately observed in Chapter 2 while the opportunity costs of contributing to a public good production are private information in Chapter 4. Second, Chapters 2 and 3 are dynamic analyses and both involve subjects interacting in continuous time. Both chapters use a methodological innovation to study dynamic games with a jump-signal process in the laboratory. This type of experimental design is different from the usual “discrete period” design as well as “flow payoff” design and has not been previously implemented in the experimental economics literature. Finally, at least one treatment in each chapter implements some form of pre-play communication. While Chapter 2 introduces unrestricted chat through discussion in a natural language, Chapter 3 has real-time as well as Poisson revision phases. In addition to the unrestricted text chat, Chapter 4 has binary communication and another treatment with a finite numerical message space.

Chapter 2

Cooperation Without Immediate Reciprocity: An Experiment in Favor Exchange

"You scratch my back and I'll scratch yours."- an old saying

2.1 Introduction

Over the last three decades there has been considerable advancement in the theoretical literature on repeated games with imperfect public monitoring[1]. However, less is known about how people behave and how much efficiency can be attained under these settings. While empirical work has been limited, some experimental work in this area has given insights into the behavior of individuals in public goods games and repeated prisoners' dilemma games with noise or private information[2]. This chapter contributes to this growing literature by providing a laboratory investigation of favor exchange relationships where a cooperative action can only be rewarded at an uncertain future date. Favor exchange captures the es-

[1]See Green and Porter (1984), Abreu et al. (1990), Fudenberg et al. (1994), Athey and Bagwell (2001), Athey et al. (2004), Mobius (2001), Nayyar (2009), Kalla (2010), Hauser and Hopenhayn (2011) and Abdulkadiroglu and Bagwell (2012a, 2012b).

[2]For example, see Palfrey and Rosenthal (1994), Cason and Khan (1999), Aoyagi and Frechette (2009) and Fudenberg et al. (2012). These studies are discussed in the section on related literature.

sential features of relations where cooperation is without immediate reciprocity and is also a leading example of a repeated game with private information. An analysis of these relationships could provide more insights with applications to industrial organization, international relations, political theory and day-to-day interactions.

A favor is an action that an individual can take which benefits another person but has an immediate net cost to herself. In theory, an individual may perform a favor in the expectation that her partner will reciprocate in the future. Exchange of favors is the basis of much of the informal cooperation that goes on in a society. An enormous part of economic activities that remain fairly uncaptured by market transactions are carried out through long term cooperation with friends, relatives and other individuals without any explicit agreements. Friends help each other move. Workers regularly cover for each other in emergencies. Politicians might support legislation introduced by other politicians, in the expectation of similar support in the future. These behaviors generate mutual gains for self-interested individuals, but reciprocity is not immediate[3].

In this chapter, the favor exchange game is as follows: two players interact in continuous time for an indefinite length; occasionally, a player receives a privately observed opportunity to provide a favor to the other player. Only one player is in a position to do a favor at a given instant. The cost of doing a favor is strictly less than the benefit to the recipient, so that, it is socially optimal to always provide a favor. Consider the following example[4]. There are two individuals (say, a plumber and an electrician) performing two separate tasks, task 1 and task 2. While both can perform each of the tasks, one of them is better at performing task 1 and the other is better in task 2. Whenever a customer arrives to these individuals,

[3]The incidence of favor exchanges may be even higher if markets are less well-developed. At the very extreme, the favor exchange relationship would take a prominent position in a moneyless society.

[4]This example is also discussed in Garicano and Santos (2004) and Abdulkadiroglu and Bagwell (2012b).

efficiency is maximized if they refer the customers looking for a certain task to the person who specializes in it. The arrival of customers is private information and referral certainly is costly because an individual loses out on the money that he could have earned if he did not refer to the more specialized person. But he would be willing to incur this cost provided he gets some of the customers from the other person in return. Another example could be a situation where two firms are involved in a partnership. At random times one of them finds a privately observed discovery which, if disclosed, could make total payoffs higher. But, disclosure comes at the cost of lowering the discovering firm's payoffs. A firm would be willing to incur the cost of disclosure only if the other firm could somehow return the favor by disclosing its privately observed discoveries in the future.

As it is very difficult to gather private information on the arrival of opportunities, as well as on the costs and benefits of favors in the field, it is advantageous to examine favor exchanges in a laboratory. In the laboratory, one can not only observe and control key variables including the information available to players, but also replicate a given scenario multiple times and make causal inferences. It turns out that the usual discrete-time design is not suitable for implementing a two-person favor exchange model in the laboratory in which opportunities to perform a favor arise stochastically over time. The key methodological innovation of this chapter is to let participants switch between the action choices asynchronously in continuous time. The continuous time design conveniently implements a dynamic game with a jump signal process where the payoff function can be updated only at specific discrete instants[5]. Thus, the design is in between that of the usual discrete-time repeated games and the continuous-time games with flow payoff structure. This type of 'hybrid' design has not been previously implemented in the laboratory in the experimental

[5]The favor exchange game is an example of this class of games.

economics literature. Also, as is usual in the repeated game literature, there are multiple equilibria and a large set of feasible and individually rational payoffs can be supported as equilibria. Experiments are ideal to investigate what payoffs are realized by the individuals and which variables are more important from a decision-making perspective.

The contributions of the present study are the following. First, it contributes to the growing empirical literature on cooperation in dynamic games in the presence of private information. Specifically, it is concerned with analyzing how cooperation is sustained in a long term relationship where cooperation is without immediate reciprocity and there is imperfect observability of stochastic events. Second, it adds to the understanding of the effects of communication and monitoring on cooperation in dynamic games. Third, at a more general methodological level, it implements a dynamic game with a jump signal process in the laboratory where the payoff function can be updated only at specific points in time.

There are three major questions asked in this research. The first question asks how the favor provision rate observed in the data with private information compares to the rate under the most efficient equilibrium. Next, what strategies are used by subjects and which variables influence the decision-making? Third, can lifting the informational constraint or providing a pre-play communication phase help subjects achieve a higher rate of favor provision?

To answer the first question, we need to be more specific about the notion of equilibrium and efficiency. The set of equilibrium payoffs is restricted to symmetric perfect public equilibrium payoffs when the players' repeated game strategies are public in the sense that today's actions are determined only by past public signals, and perfect in the sense that they form a Nash equilibrium after every public history. Different scenarios are considered by changing the benefit of receiving a favor and the arrival rate of opportunities. The primary finding is that the overall favor provision level is significantly lower than the level that could have been attained by the most efficient perfect public equilibrium. This is true for all the scenarios considered, including situations with very high benefit and very high arrival rate of opportunities, showing that it is a robust finding. The favor exchange game with private information thus results in considerably low efficiency levels.

One class of equilibria that has been proposed in the favor exchange literature is known as the chips mechanism. These are threshold strategies conditional on net favors as the state variable. Under these strategies people keep providing favors until the net favors provided by them hit an upper bound. The data rejects these types of score-keeping behavior and reveals that individuals do not employ any of the threshold strategies based on net favors provided. Rather, the decision to stop granting favors primarily depends on whether the individual has just provided a favor in the recent past. Similarly, the decision to restart providing favors depends mainly on whether the partner has provided a favor recently. Thus, the focus is more on these 'time since last favor' variables showing that behavior is explained more by the notion of "What have you done for me lately?" than individuals engaging in exact score-keeping of cumulative favors.

There are certain situations where the opportunities to provide a favor are publicly observed by both parties involved. For example, this public observability of opportunities seems to be a reasonable assumption in the context of village economies in developing countries[6]. The question posed then is that whether letting individuals observe the opportunities received by their partners helps them achieve a higher efficiency. In the absence of such

[6]People living in village economies have a low and highly volatile income. In the absence of insurance and credit markets, individuals engage in risk sharing through informal institutions. People transfer a significant part of their income in order to assist those who have received a low income. Also think of a gift exchange relation where an individual keeps sending and receiving gifts on important occasions like birthdays, anniversaries etc. See Camerer (1988), Landa (1994), Carmichael and MacLeod (1997), and Charness and Haruvy (2002) for studies and references on gift-exchange relationships.

information, people have to condition their behavior on publicly observed state variables, e.g., net favors granted and time since last favor provided by partner. However, when the informational constraint is absent, individuals can simply condition on whether or not their partner provided them a favor at the last received opportunity. Findings show that the overall rate of favor provision is significantly higher when individuals can perfectly monitor the opportunities received by their partners.

We live in a world where language and communication plays an important part in our day-to-day activities. We often talk our way through a problem. As such it becomes imperative to ask the question concerning whether individuals can improve the efficiency levels attained by them when they can talk among themselves before starting out with a favor exchange relation with privately observed opportunities. In line with other studies on effect of pre-play communication in one-shot games, cooperation is indeed higher when individuals can talk before the start of the dynamic favor exchange game. Thus, both communication and information are channels through which individuals gain some understanding of their partners' intentions and this seems to help them achieve higher levels of efficiency.

There are important differences between the repeated prisoner's dilemma and the dynamic favor exchange game. The primary distinction is in the strategic nature of the games. A costly cooperative action cannot be immediately rewarded; one must wait some time before a favor is returned. Neither is this time deterministic nor is there any guarantee that the partner is going to reciprocate. The time horizon is also stochastic and adds to the uncertainty. The introduction of private information makes the environment completely different from that of a noisy version of the repeated prisoner's dilemma. This is because events are asynchronous and stochastic in the favor exchange game. Chance plays a much bigger role in these games as opposed to in a prisoner's dilemma game. The stochastic opportunities basically imply how frequently individuals interact among themselves and are taken as a control variable in this study. Even the situation where individuals can observe the opportunities received by their partners is different from a prisoner's dilemma with perfect monitoring. Since opportunities are stochastic, an individual can switch off cooperation not just because she is not being "nice" but simply because she got too many opportunities by chance and first wants the net favors to even out. Thus, there might be asymmetries involved in the gains from cooperation depending on the realized timings of the opportunities even though the game is ex-ante symmetric.

The remainder of the chapter is organized as follows. The related literature is reviewed in Section 2.2. Section 2.3 presents the model of favor exchanges on which the experiment is based and discusses the theoretical benchmarks. Section 2.4 describes the experimental design and procedures as well as the central hypotheses. The results are discussed in detail in Section 2.5. Section 2.6 concludes.

2.2 Related Literature

This section has four parts. The first part presents a brief overview of the theoretical models developed in the favor exchange literature. Since the present chapter contributes to the budding literature on repeated game experiments in a continuous time framework, a short description of the studies included in this literature is provided in the second part of this section. Experiments on repeated games with imperfect monitoring are discussed in the third part and the last part lists some conceptually related studies.

2.2.1 Theoretical Literature on Favor Exchange Models

A two person model of favor exchange with private information was first studied by Mobius (2001) who analyzed the equilibria in threshold strategies conditioned on net favors provided[7]. Using a similar model as in Mobius and allowing individuals to grant partial favors, Hauser and Hopenhayn (2011) characterize the Pareto frontier of the equilibrium outcomes via an integro-differential equation and show that it is self-generating, that is, the frontier of the equilibrium set is supported by payoffs also on the frontier. They also find that the first-best cannot be supported.

The study of favor exchanges began with Calvert (1989). He analyzes a simple two-player model where in stage one, nature selects at random whether each player will have the opportunity to 'receive a favor' from the other. Learning the result of stage one, the player who can provide the favor chooses (or both players choose if both can provide a favor) a level of effort to devote to providing the favor. Strategies necessary to realize the full gains from cooperation are examined. He identifies the conditions for implementation of full and partial "favor exchange" in equilibrium. Although Calvert introduces uncertainty and asymmetry in his model, there is no private information.

Neilson (1999) examines the favors people perform for each other using a stochastic version of the infinitely repeated prisoner's dilemma game. He shows that the socially efficient outcome may not be consistent with an equilibrium because some favors are too expensive to perform. He also finds that although it may be impossible to support some exceptionally costly favors in equilibrium, there do exist equilibria in which inefficient favors are performed.

Recently there has been a renowned interest in repeated favor exchange games with

[7]The next section discusses these equilibria in greater detail.

private information. Nayyar (2009) develops a discrete-time model of favor exchange that allows for asymmetric opportunities for doing favors and finds that cooperation among players can be approximately supported in equilibrium even under those asymmetries. She also shows the monotonicity of the equilibrium set in the discount factor as well as the likelihood that someone will be in a position to do a favor. Lastly, using results from Azevedo and Moreira (2007) she finds that an exact folk theorem does not hold in her model. For any discount factor less than one, there exists feasible and individually rational payoffs that are not in the equilibrium set. Kalla (2010) studies favor trading when players have privately known types. By introducing low and high types[8], he finds conditions under which the high type players are almost always able to separate themselves from the low type players.

Another related study is Abdulkadiroglu and Bagwell (2012b) who analyze a two person repeated trust game with private information having a similar flavor to that of a favor exchange game. They find that players are willing to exhibit trust and thereby facilitate cooperative gains only if such behavior is regarded as a favor that must be reciprocated either immediately or in the future. Thus, they have both immediate reciprocity and dynamic reciprocity in their model. They capture different forms of trust relations by considering different self-generating sets[9]. For any given form of trust relation, they then construct and interpret the optimal cooperative strategies of players with bounded patience levels. Athey and Bagwell (2001) develop a theory of optimal collusion among privately informed and impatient firms in discrete time. They consider an infinitely repeated Bertrand game, in

[8]He varies the discount factor of the players to generate "low type" players who do not find cooperation beneficial and "high type" players who have benefits from cooperation.

[9]They implement the symmetric self-generating line (SSGL) of payoffs where the self-generating set of payoffs is a line along which total payoffs sum to a constant value and this line is symmetric around the 45 ° line. They also study strongly symmetric equilibria (SSE) and hybrid equilibria building from the optimal SSGL and SSE constructions.

which prices are publicly observed and each firm receives a privately observed i.i.d. cost shock each period. In every period, a firm's cost can be high or low, so productive efficiency calls for only the low-cost firm to be producing when the other firm has high costs. They characterize the set of perfect public equilibrium values and find a discount factor strictly less than one with which the first-best payoff is achieved as an equilibrium. When the firms have the same cost level, they may allocate market share asymmetrically and thereby achieve transfers without sacrificing efficiency. This is to be contrasted with the result in Abdulkadiroglu and Bagwell (2012b) where the first-best payoff is not achievable.

2.2.2 Experimental Literature on Games in Continuous Time

The experimental set-up in this chapter is continuous in the sense that there are no discrete periods and individuals interact continuously over time. The literature on continuous time experiments on repeated/dynamic games is rather limited. Following is a modest review of the experimental literature on games in continuous time.

Friedman and Oprea (2012) study prisoner's dilemma games played in continuous time with flow payoffs accumulated over 60 seconds. They compare cooperation rates from these continuous time sessions to rates under control grid and one-shot sessions. They find a smooth negative relationship between the cooperation rate and the length of the stage game. Bigoni et al. (2011) compare behavior in continuous time prisoner's dilemma games under deterministic and stochastic durations and report that cooperation rates under a deterministic duration are similar to or higher than under a stochastic one, which is in contrast to the results obtained from games in discrete time (see Dal Bo (2005)). Feeley et al. (1997) implement infinite-choice continuous time prisoner's dilemma games in the laboratory to observe what sorts of behavioral patterns emerge[10]. They also analyze the dynamics of cooperation using phase space and individual over-time plots with subjects interacting for a known time horizon of 5 minutes.

Apart from analyzing experiments that involve studying cooperation in continuous time, there have been a few other related experiments in a continuous time setting. Continuous coordination laboratory games are examined by Brunnermeier and Morgan (2010) and Cheung and Friedman (2009). Camerer et al. (2012) explore the empirical nature of temporal anxiety in economics decisions by implementing experimental timing games. Oprea et al. (2011) implement a series of war of attrition games in a near-continuous time setting. A similar experimental setting is also used by Horisch and Kirchkamp (2010) who study an all-pay auction where the subjects have to decide when to stop bidding. Also, Friedman et al. (2011) implement the standard Hawk-Dove bimatrix game in continuous time in the laboratory[11].

2.2.3 Experimental Literature on Repeated Games with Imperfect Monitoring

Most of the experimental studies on cooperation in repeated games implement a complete information setting with perfect monitoring, with a few exceptions. Palfrey and Rosenthal

[10] The incentive structure is however different in their study. Participants were told that a random drawing for a free compact disc of their choice would be held after the completion of the experiment. Their chances of winning were directly proportional to their score. Hence, the higher their final score, the better their chances to win the free disc.

[11] Other than the already mentioned analyses, there are few early experimental studies worth discussing. Kahan and Rapoport (1974) report the results from an experiment about games of timing presented by a computer-controlled, two-person, infinite game that simulates the Western-style duel. Games of timing constitute a sub-class of two-person, constant-sum, infinite games, where the problem facing each player is not what action to take, but rather when he should take action. In a duel, each of the two players has a gun with one bullet and, starting at opposite ends of town (time 0), they slowly walk towards each other. The closer they are, the more accurate their fire. The decision each must make is when to draw his gun and fire at the other duelist. Rapoport et al. (1976) extend this analysis to probabilistic duels, where a player does not know with certainty whether or not his opponent is armed, but only knows the probability of such armament. Both these studies implement the duel experiments on a continuous platform.

(1994) compare the rate of contribution to a public good under private information about contribution costs when players meet only once with the case when they meet repeatedly over an indefinite horizon. They find that repetition leads to more cooperation than one-shot games. Cason and Khan (1999) study a repeated public good experiment and compare standard perfect monitoring with perfect, but delayed monitoring of past actions. They do not find any significant difference in the levels of contributions between the two treatments. Feinberg and Snyder (2002) consider a version of the repeated prisoner's dilemma where each subject observes his own payoff in each period. They introduce imperfection by occasionally manipulating those payoff numbers, and compare the treatments with and without the ex post revelation of such manipulation. Less collusive behavior is found in the latter treatment. Holcomb and Nelson (1997) analyze a repeated duopoly model in which information about the opponent's quantity choice is randomly changed 50% of the time. They find significant effect of such manipulation on market outcomes.

Aoyagi and Frechette (2009) study infinitely repeated prisoner's dilemma games under imperfect monitoring and find that cooperation declines with the level of noise of the public signal. Fudenberg et al. (2012) study the experimental play of the repeated prisoner's dilemma when intended actions are implemented with noise, and identify the sort of strategies used by the subjects. They find substantial heterogeneity in strategies across players and also show that forgiving and lenient strategies are more successful. Another recent paper by Fudenberg et al. (2013) studies reciprocal altruism in the setting of infinitely repeated games where the intended actions are implemented with error.

The environment of the current study is, however, different from these papers. There is asynchronicity and a "cooperative" action cannot be immediately rewarded. The monitoring structure is different too. An informative signal here necessarily means a "cooperative" action by the opponent but the incidence of no signal does not guarantee that the opponent is not being “cooperative”.

2.2.4 Other Related Studies

Finally, there are a few studies that are somewhat conceptually related to the exchange of favors. Many of these studies investigate behavior when turn-taking is an efficient way to cooperate. While Luce and Raiffa (1957) introduced the notion of single-period alternation, Lau and Mui (2008) show that among different subgame-perfect equilibria with “multiple-period alternation”, the equilibria with single-period alternation give the players the highest payoffs. Lau and Mui (2012) investigate turn-taking behavior in many 2x2 games and determine conditions under which there exists a unique “turn-taking with independent randomizations” strategy profile that can be supported as a subgame-perfect equilibrium. Cason et al. (2012) observe turn-taking in a two-player common-pool resource assignment game and demonstrate that subjects are able to teach such behavior to future opponents. Using simulations, Neill (2003) shows that turn-taking can be achieved in a noisy environment, even when agents use limited memory strategies. Kuzmics et al. (2012) find that turn-taking forms the predominant type of continuation play in repeated allocation games with three players as well as with two players. Kaplan and Ruffle (2011) study a repeated entry game with private information where each of two players makes a binary choice of entry or exit after observing privately a randomly drawn integer between 101 and 105 each period. Exit yields a payoff of zero, while entry gives the player her drawn number if her opponent exits, but one-third of the drawn number if both players enter. They show that turn taking can be supported as an equilibrium in this repeated game, and find that many subjects take turns to enter to avoid efficiency loss due to simultaneous entry.

The dynamic favor exchange game with private information is also related to the literature on dynamic incentives for risk sharing. One could reinterpret favor as a consumption good with the benefit as the utility for the receiver and cost as the forgone utility for the individual providing the favor. In that case, the utilitarian solution allocates the good to the person with highest marginal utility of consumption. Green (1987) characterizes the incentive-compatible optimal allocation in an environment with a continuum of players with privately observed endowments and shows that this allocation can be supported by the constrained exchange of infinite-lived bonds between the traders and the intermediary. Also see Atkeson and Lucas (1992) and Hertel (2004). Thomas and Worrall (1990) characterize the Pareto frontier of equilibrium outcomes in a model where a risk neutral lender offers insurance to a risk-averse borrower who receives privately observed i.i.d. income shocks each period. Thus, there is an incentive for the borrower to report low income each period.

Jackson et al. (2011) provide a game theoretic foundation for social enforcement of informal favor exchange, and also examine network patterns of favor exchange from 75 villages in rural India. They consider settings where simple bilateral quid-pro-quo enforcement is insufficient to sustain favor exchange, but when these bilateral interactions are embedded in a network of interactions whose functioning can be tied to each other, individuals find it in their interest to cooperate given threats of ostracism or loss of multiple relationships for failure to behave well in any given relationship.

2.3 The Model

Our model is based on Mobius (2001). This framework is also considered in Hauser and Hopenhayn (2011) and Nayyar (2009). For simplicity, the model is presented here in a discrete-time setting as follows: Two risk-neutral individuals, 1 and 2, interact indefinitely and each period one of them receives an opportunity to provide a favor to the other person with probability p. A granted favor has benefit b to the receiver and it costs $c < b$ to provide a favor. Thus, it is socially optimal to always provide favors. There is a common discount factor $\delta \in (0, 1)$ and players seek to maximize the present discounted values of their utilities. A similar model in continuous time would have the random opportunities arriving according to a Poisson process[12].

At any time period t, one of three states, F_1, F_2, F_ϕ occurs. In state $F_1(F_2)$, player 1(2) is in a position to do the other player a favor. In state F_ϕ, neither player is in a position to perform a favor. The states are independently and identically distributed in each period, F_1 with probability p, F_2 with probability p, and F_ϕ with probability $1 - 2p$, where $p < \frac{1}{2}$. Depending on the assumptions on the observability of the opportunities to provide a favor, one can analyze the situation differently. The situation where each person can observe the opportunities received by both players is presented first and then the more complicated case is analyzed where the ability to do a favor is private information.

2.3.1 Public Observability of Opportunities

In the repeated game with perfect observability of the opportunities received by both players, at time t, every player observes the history of: (i) when each of them had an opportunity to provide the other a favor, and (ii) whether a favor was provided when there was an opportunity to do so. Let us define a variable Y^t that records the identity of the player who does a favor in period t. It takes a value of one (two) if player 1 (player 2) provides a favor at time t and takes a value of zero if no favors are observed in period t. Denote the realized

[12]In the experiments, a discrete period is implemented at 1 second intervals but players can asynchronously switch between two modes showing the intent to provide a favor and hence their decision is to switch between 'the intent to provide a favor' and 'no intent'. More is discussed about the implementation of the model in the next section.

value of Y^t at time t by y^t.

Let $h^t = (\gamma^t, \xi^t)$ be the public history of the realized states and that of when and who provided favors in the past at time t, where $\gamma^t = (s^0, s^1, ..., s^{t-1})$ and $\xi^t = (y^0, y^1, ..., y^{t-1})$. Denote the set of all possible public histories at time t by H^t. Then, a (pure) strategy for player i at time t can be defined as $x_i^t : H^t \times \Omega \to \{0, 1\}$; and let x_i denote a sequence of such strategies for $t = 0, 1, ...\infty$, such that,

$$x_i^t(h^t, s^t) = \begin{cases} x_i^t(h^t) \in \{0,1\} & \text{if } s^t = F_i \\ 0 & \text{if } s^t \neq F_i \end{cases} \quad .$$

Player i's expected payoff in the repeated game from a strategy profile x is given by

$$\pi_i(x) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} p(bx_j^t - cx_i^t) \quad .$$

The strategy profile x is a (pure) equilibrium if for every i, $\pi_i(x) \geq \pi_i(x_i^{'}, x_j)$ for any strategy $x_i^{'}$. An equilibrium is subgame-perfect if x_i is a best response to x_j for each i after every public history, h^t. It is easy to see that in this case, where the opportunities to provide favors are observable by both individuals, a subgame-perfect equilibrium would exist that achieves the social optimum payoff through the following strategy: a person grants favors whenever possible, as long as her partner has done so in the past, and stops granting favors whenever she sees her partner defect (i.e. not help when she can). This equilibrium can be supported for:

$$-c + \delta p(b - c) + \delta^2 p(b - c) + ... + \delta^t p(b - c) + ... > 0 \quad .$$

The left hand side gives the expected payoffs when the prescribed strategy is followed and a favor granted at time t when there is an opportunity to do so and the right hand side shows the payoff from foregoing an opportunity when there is one. Thus, the condition is:

$$\delta p(b-c)(1+\delta+\delta^2+...+\delta^t+...) > c$$
$$\iff \quad \frac{p(b-c)}{r} > c,$$

where $r = \frac{(1-\delta)}{\delta}$ is the discount rate.

2.3.2 Private Observability of Opportunities

However, when the opportunities to provide favors are privately observed, individuals have to infer from the history of favor exchange the extent to which their partner provided favors for them. Now, player i's information set is $\{(F_i), (F_{i \neq j}, F_\phi)\}$ and a player who is not in a position to do a favor does not know whether or not her opponent is in a position to do a favor. Thus, using the notations from the previous model with public observability of opportunities, now only the history, $\xi^t = (y^0, y^1, ..., y^{t-1})$, of when and who provided favors in the past is publicly accessible. The entire history of the realized states is no longer available publicly.

Attention is restricted to sequential equilibria where players condition only on public histories and their current type but not on their private history of types. Such strategies are called public strategies and such sequential equilibria are termed as perfect public equilibria (PPE)[13]. Define a "public strategy" for player i at time t as $\sigma_i^t : \Xi^t \times \Omega \to \{0,1\}$ (which is a function of the public history, $\xi^t \in \Xi^t$, alone) and let σ_i denote a sequence of such strategies for $t = 0, 1, ...\infty$, such that,

[13]The Markov perfect equilibrium (MPE) in this game is uninteresting. It is unique and involves no exchange of favor.

$$\sigma_i^t(\xi^t, s^t) = \begin{cases} \sigma_i^t(\xi^t) \in \{0,1\} & \text{if } s^t = F_i \\ \\ 0 & \text{if } s^t \neq F_i \end{cases} \quad .$$

Player i's expected payoff in the repeated game from a strategy profile σ is given by

$$\pi_i(\sigma) = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} p(b\sigma_j^t - c\sigma_i^t) \quad .$$

The strategy profile σ is a (pure) equilibrium if for every i, $\pi_i(\sigma) \geq \pi_i(\sigma_i^{'}, \sigma_j)$ for any strategy $\sigma_i^{'}$. An equilibrium $\sigma = (\sigma_1, \sigma_2)$ is "public" if σ_i is a public strategy for each i. A public equilibrium is "perfect" if σ_i is a best response to σ_j for each i after every public history, ξ^t. It is "symmetric" if $\sigma_1 = \sigma_2$. Let $E(\delta)$ be the set of perfect public equilibrium payoffs, given δ.

Denote the set of feasible and individually rational payoffs of the repeated game by U and the continuation values for the players by $(v_c, w_c) : Y^t \to U$. Depending on the value of Y^t at time t, one can denote (v_c, w_c) as:

$$(v_c, w_c) = \begin{cases} (v_1, w_1) & \text{if } Y^t = 1 \\ \\ (v_2, w_2) & \text{if } Y^t = 2 \\ \\ (v_\phi, w_\phi) & \text{if } Y^t = 0 \end{cases} \quad .$$

The Pareto frontier, F of $E(\delta)$ is defined as: $F = \{(v,w) \in E(\delta) \text{ such that } w = g(v)\} \cap \{(v,w) \in E(\delta) \text{ such that } v = h(w)\}$, where

$$g(v) = \max_{((v_1,w_1),(v_2,w_2),(v_\phi,w_\phi))} \{w : (v,w) \in E(\delta)\}$$

$$\text{and}\, h(w) = \max_{((v_1,w_1),(v_2,w_2),(v_\phi,w_\phi))} \{v : (v,w) \in E(\delta)\}$$

Contrary to the situation with public observability of opportunities, the first-best payoff cannot be supported in any equilibrium when opportunities are privately observed. To achieve the first-best payoff, both players have to provide favors at every possible opportunity. This is regardless of whether or not the other player has reciprocated and hence, regardless of any history. But, then each player has a profitable deviation to not do a favor, which the other person cannot observe as opportunities are privately observed. Thus, in equilibrium, the first-best outcome cannot be enforced.

While the first-best is not achievable, there is also no tractable way to compute the Pareto frontier of the equilibrium payoffs so as to understand how close players can get to the first-best payoff. It turns out that one can focus on a class of equilibria which is not only simple and intuitive but also analytically tractable. These equilibria were first studied by Mobius (2001) where he restricted attention to the difference k_i in the past number of favors provided by person i to person j and the number of favors she received from person j. This state variable becomes a natural choice to condition individual's strategies as it captures the common notion that someone 'owes' favors (negative k) or is owed favors (positive k). Thus, one can focus on equilibria where players' strategies only depend on the state variable k_i. For $-k^* \leq k \leq k^*$, individuals follow the following strategy: grant favors if $k < k^*$ and stop doing favors if $k = k^*$. As the equilibria are symmetric, her partner grants favors if $k > -k^*$ and stops doing favors when $k = -k^*$. Formally, a strategy for an individual can be defined as follows:

$$\sigma_i^t(k_i^t) = \begin{cases} 1 & k_i^t < k^* \\ \\ 0 & k_i^t = k^* \end{cases}$$

where 1 means granting a favor and 0 means not granting a favor when the opportunity arrives.

There are at least two reasons to restrict attention to these equilibria. First, the most efficient equilibrium in this class of equilibria can attain high payoffs and can provide a useful "lower bound" on the maximum efficiency that can be achieved. That is, the most efficient PPE can attain payoffs that are at least as large as those achievable under the most efficient equilibrium in the class of equilibria considered by Mobius (2001). Second, it gives a clear set of strategies that can be supported as an equilibrium of the game which is helpful because one can "test" for these strategies in the data. Often, the theory of repeated games with private information is silent about what strategies could achieve the payoffs in the equilibrium set.

These threshold strategies are also known as chips mechanisms (CMs), the idea being the following. One can think of both players as starting with k^* poker chips each. Whenever one player receives a favor, she gives the other player a chip. If one player runs out of chips, she receives no favor until she grants a favor to the other individual and obtains a chip in exchange. Under a CM, people keep track of the imbalances between each other and engage in score-keeping, with the restoration of balance being a primary aspiration. The mechanism that corresponds to playing the equilibrium with lowest k^* other than zero[14], that is, with a $k^* = 1$, can be termed the simple chips mechanism (SCM). Under an SCM, an individual

[14] $k^* = 0$ also constitutes an equilibrium where nobody provides a favor ever, thus, resulting in zero payoffs for both the players. It is also the unique Markov perfect equilibrium. Hence, the set of CM equilibria is non-empty.

can provide at most two consecutive favors[15]. The mechanism that corresponds to playing the equilibrium with k^{max} =maximal k^* is more complex and allows for further efficiency gains (to be explained later in the section) as compared to the SCM. This mechanism is called the best chips mechanism (BCM)[16].

Let $V_{k^*}(k)$ denote the value of a person in state k in the favor exchange relationship corresponding to the equilibrium where individuals stop granting favors just as $k = k^*$. This is without loss of generality because the equilibrium is symmetric. The Bellman equation and boundary conditions then are as follows:

$$rV_{k^*}(k) = p(V_{k^*}(k+1) - V_{k^*}(k) - c) + p(V_{k^*}(k-1) - V_{k^*}(k) + b) \quad ,$$

$$rV_{k^*}(k^*) = p(V_{k^*}(k^*-1) - V_{k^*}(k^*) + b) \quad ,$$

$$rV_{k^*}(-k^*) = p(V_{k^*}(-k^*+1) - V_{k^*}(-k^*) - c) \quad .$$

The equation system defines a perfect public equilibrium if the incentive compatibility (IC) constraints and the individual rationality (IR) constraints are satisfied. Not only do individuals have to gain from doing favors for $k < k^*$, but the relationship between two people must have positive value for $k = -k^*$. The incentive compatibility constraint provides an incentive to a player to perform a favor when she is in a position to do so. The incentive is provided in the form of future benefit to this player. When a player performs a favor in exchange for a chip, she is entitled to a higher expected payoff next period than

[15]Thus, this behavior under an SCM is more subtle than the turn-taking behavior in Lau and Mui (2008) and Kuzmics et al. (2012).

[16]One can also call this mechanism the optimal CM/the most efficient CM.

she was at the beginning of the current period. Thus, we must have the following:

$$V_{k^*}(k+1) - V_{k^*}(k) \geq c \quad for - k^* \leq k < k^* \quad (ICC)$$

$$V_{k^*}(-k^*) \geq 0 \qquad\qquad\qquad\qquad (IRC) \quad .$$

Using these equations and solving for the resulting second order difference equation gives the following form of the value function:

$$V_{k^*}(k) = \frac{p(b-c)}{r} + x\alpha^k + y\beta^k \quad ,$$

where α, β, x, and y are given by

$$\alpha = 1 + \frac{r}{2p} - \sqrt{\frac{r^2}{4p^2} + \frac{r}{p}} \quad ,$$

$$\beta = 1 + \frac{r}{2p} + \sqrt{\frac{r^2}{4p^2} + \frac{r}{p}} \quad ,$$

$$\alpha < 1 < \beta, \alpha\beta = 1 \quad ,$$

$$x = \frac{(b - c\alpha^{2k^*+1})\beta^{k^*}}{(\alpha - 1)(\beta^{2k^*+1} - \alpha^{2k^*+1})} \quad ,$$

$$y = -\frac{(b - c\beta^{2k^*+1})\alpha^{k^*}}{(\beta - 1)(\beta^{2k^*+1} - \alpha^{2k^*+1})} \quad .$$

The magnitude of k^* is important in determining the expected payoffs of individuals. As it is fully efficient to grant favors whenever possible, an efficiency loss occurs when individuals reach the boundaries k^* and $-k^*$. By increasing k^*, individuals' utilities from the relationship increase because they spend less time at the boundary of the state space. However, since one favor done today is rewarded by the promise of exactly one favor in the future, k^* cannot be infinitely large. The k^{max} corresponding to the best CM can be calculated using the following equality[17]:

$$\left(1+\frac{r}{2p}-\sqrt{\frac{r^2}{4p^2}+\frac{r}{p}}\right)^{2k^{max}+1} = \frac{b}{c}-\sqrt{\frac{b^2}{c^2}-1} \quad .$$

It follows immediately that k^{max} increases as b/c goes up and as r/p declines. A higher benefit of receiving a favor makes individuals cooperate more as the relationship has become more valuable. On the other hand, a faster arrival rate of favors makes individuals more willing to cooperate as cheating becomes harder.

Apart from being intuitive and easy to implement, a chips mechanism has several desirable properties that make it an obvious choice as the benchmark in the private information environment. Chips mechanisms, also known as equality matching mechanisms, are based on a model of even balance and people keep track of how far out of balance (depending on the metric of balance) the relationship is. The idea is that each person is entitled to the same amount as each other person in the relationship, and that the direction and magnitude of an imbalance are meaningful. People think about how much they have to give to reciprocate or compensate others or come out even with them. CM always entails a simple additive metric of who owes what and who is entitled to what.

While the best CM is asymptotically efficient[18] and will also serve as the benchmark in this current study, it is important to note that there are two special features of this mechanism which, if relaxed, can lead to higher expected payoffs provided "partial favors"

[17] For a detailed proof, see Mobius (2001). Also see Abdulkadiroglu and Bagwell (2012a) where a simple algorithm is discussed to find the BCM.

[18] To see this, note that in the long run the distribution over the states $k = \{-k^*, -k^*+1, ..., 0, 1, ..., k^*-1, k^*\}$ is uniform, so the probability that a favor is not granted is $1/k^*$. It is easy to verify that $k^* \to \infty$ as $p/r \to \infty$, so that this probability goes to zero.

are also allowed instead of just the binary choice of "full favor" or "no favor", as emphasized by Hauser and Hopenhayn (2011) and Abdulkadiroglu and Bagwell (2012b). The first special feature is that the rate of exchange of current for future favors is the same regardless of entitlements. Letting the relative price of favors decline with a person's entitlement could reduce the region of inefficiency. The second feature is that individuals' continuation values do not change unless an individual does a favor. This is restrictive in that it rules out the possibilities of appreciation or depreciation of entitlements and punishment in case "not enough" cooperation is observed. They find "debt forgiveness", that is favors owed depreciate over time, which gives players an incentive to continue doing small favors even when reciprocation has not been received for previous favors. Obviously, this feature cannot be exhibited in the current model without "partial favors".

2.4 Experimental Design, Procedures and Hypotheses

This section is divided into two parts. The first part discusses the design and procedures of the experiment. The second part lists the central hypotheses that are investigated in this study.

2.4.1 Design and Procedures

In the experiment, the cost of providing a favor, c is fixed at five points and the discount rate, r at 0.01. The parameters, opportunity arrival rate, p, and benefit of receiving a favor, b are varied to give rise to a set of four treatments (for each private and public information environment). These are summarized in Table 2.1, along with the k^{max} corresponding to the BCM in each case for the private information environment. Also displayed is the initial expected value that a player obtains if she and her partner follow a SCM, BCM, and the "always grant favors" strategy. Obviously there are significant gains from following the BCM as opposed to the SCM. The parameters are chosen such as to give rise to different k^{max} in more than two treatments and same k^{max} in exactly two of the (intermediate) treatments. The first best is achievable under these parameters in the public information treatments. The expected length of a match is 100 seconds. A shorter expected duration is avoided purposefully to ensure games with "long" durations. This is important because this is a game without immediate reciprocity as opposed to a flow payoff design and hence individuals must be given enough time to "reciprocate".

Treatment	b	p	c	r	k^{max}	V^{SCM}	V^{BCM}	V^{FB}
Low b-Slow p (Low-Slow)	10	0.1	5	0.01	2	33.87	40.93	50
Low b-Fast p (Low-Fast)	10	0.3	5	0.01	4	100.56	135.05	150
High b-Slow p (High-Slow)	25	0.1	5	0.01	4	135.50	183.72	200
High b-Fast p (High-Fast)	25	0.3	5	0.01	6	402.24	563.06	600

b = benefit of receiving a favor; p = opportunity arrival rate; c = cost of providing a favor; r = discount rate; k^{max} =maximal k^* predicted in the private information case; V^{SCM} = initial expected value if a SCM is followed; V^{BCM} = initial expected value if a BCM is followed; V^{FB} = initial expected value if all favors are granted.

Table 2.1: Treatments, parameters and the k^{max}

The experiments were all conducted at the Social Science Experimental Laboratory (SSEL), California Institute of Technology (Caltech) using the Multistage software package[19]. Subjects were recruited from a pool of volunteer subjects, maintained by the SSEL. A total of eight sessions were run, using a total of 108 subjects. No subject participated in more than one session. The treatment variables are the benefit of receiving a favor (low and high), arrival rate of opportunities to provide a favor (slow and fast), the information about the partner's opportunities (private and public) and the existence of a discussion period prior to play (no communication and pre-play communication). Table 2.2 summarizes the

[19]Please visit http://software.ssel.caltech.edu/ for more details on the multistage software project.

characteristics of each session. On arrival, instructions[20] were read aloud. Subjects interacted anonymously with each other through computer terminals. There was no possibility of any kind of communication between the subjects, except in the communication sessions. There were two practice matches, followed by the paid matches with each match requiring two subjects to be paired together and play the continuous time favor exchange game. In the communication sessions each subject was allowed to send unlimited messages[21] to her partner for a length of 60 seconds before the start of a match. A subject had two modes named 'Do favor' and 'Do not do favor' to choose from, and her choice was to switch from one mode to the other whenever she wished to do so.

Figures 2.1-2.3 show the user interface. A match begins by matching two participants from the room to play the dynamic favor exchange game. Figure 2.1 displays the initial choice screen. Each subject had to choose between the two options labeled 'INVEST' and 'DO NOT INVEST', respectively. To keep the experimental language neutral[22], the favor exchange game was presented as an investment game. The subjects were told that they would receive random opportunities to undertake an investment that would cost them c points but pay an immediate return of b to the person they are matched with. Thus, 'INVEST' is synonymous with 'Do favor' and 'DO NOT INVEST' is synonymous with 'Do not do favor'. Once all the subjects in the room made an initial decision, the game began in real time. For the treatments with communication, there was an additional 60 seconds discussion period in the beginning before the initial choice screen was displayed.

[20] A copy of the instructions is given in appendix A.

[21] Subjects were asked to send messages relevant to the experiment and were told not to reveal their identities and also not to use any offensive language.

[22] Using the word favor may have suggested that doing a favor is ethical or morally right. Also, the person with whom a subject was matched with was not termed as partner/counterpart but rather as 'Other'.

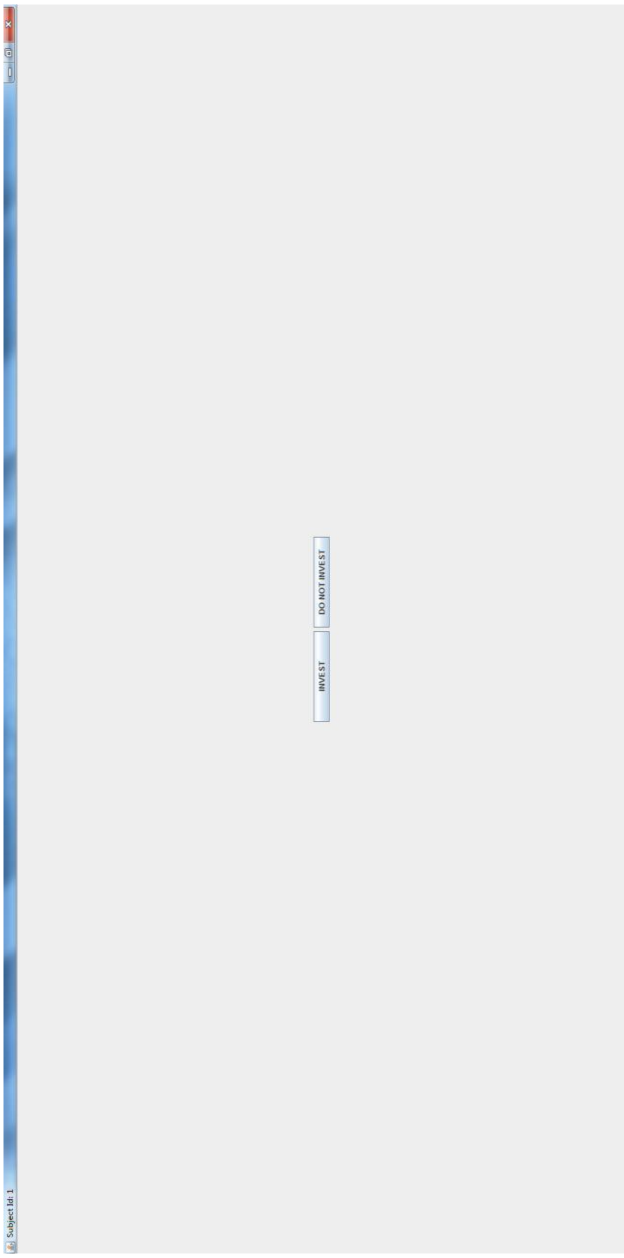

Figure 2.1: Initial choice screen

Figure 2.2 (Figure 2.3) displays a sample screen from the play of a game in the private information treatment (public information treatment). Each subject could freely switch between the two options by clicking on the respective buttons as many time as they wished[23]. The current option selected by the subject was also clearly mentioned in the top left corner of the screen. Having the 'INVEST' option activated means that the subject automatically invests whenever an opportunity arrives. Similarly, having the 'DO NOT INVEST' option activated means that the subject automatically foregoes the opportunity to invest whenever such an opportunity arrives. Note that this is like choosing the intent to invest rather than making a decision after the arrival of an opportunity, which is exogenously specified.

[23]Since clicking makes a sound, it would be best to have touch screen computers so that subjects can switch between options by touching the screen with a finger instead of a click. In fact, Bigoni et al. (2011) use touch screen PCs in their experiments.

Figure 2.2: Subject screen for “Private Information” treatments

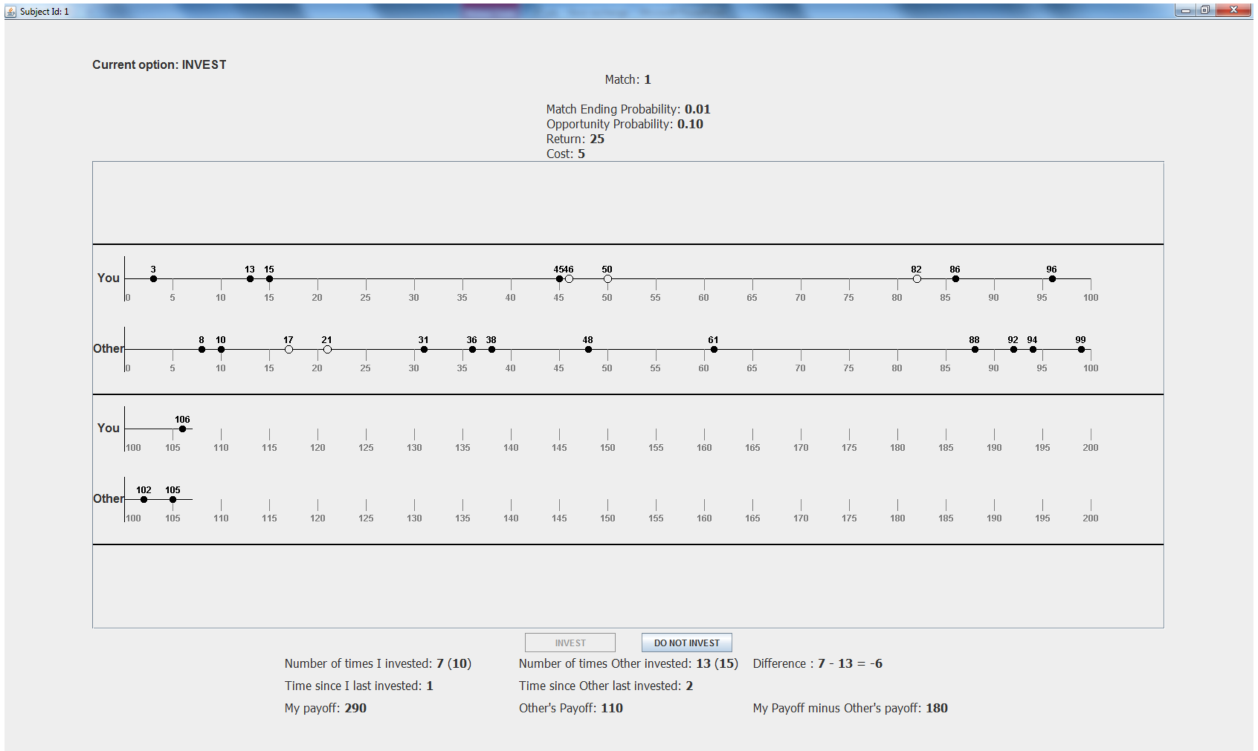

Figure 2.3: Subject screen for “Public Information” treatments

The screen also had the basic parameter information along with the match number at the top center of the screen (see Figures 2.2 and 2.3). There were two tickers corresponding to the labels 'You' and 'Other' which progressed in the course of a match. Opportunities are shown as either a black or a white dot on the ticker. A black dot represents an investment along with the time (in seconds) at which it was undertaken, whereas a white dot stands for an investment opportunity that was not undertaken along with the time of the opportunity arrival. In the private information treatments (Figure 2.2), subjects were told that the opportunities they receive are observed only by them and not by the person they are matched with. The person they are matched with only gets to see when the subject actually invested.

A white dot can only occur on the ticker corresponding to 'You' in the private information treatments because individual opportunities are not observed by the partner. This is however not true in case of the public information treatments where each individual can observe the opportunities received by her partner too. However, the information on the current selected option by the partner was still not displayed anywhere on the screen. The only difference in the private and public information treatments is over the observability of the opportunities and not over the current selected option by partner.

A session without communication lasted on average 1 hour and 5 minutes, whereas the communication sessions lasted 20 minutes longer. Subjects earned on average US$15 in the private information treatments without communication, US$25 in the private information treatments with communication and US$26 in the public information treatments[24].

In each match, the arrival of opportunities for each pair of participants was decided using a random draw of an integer from an uniform distribution once per second[25]. The

[24] Payoffs ranged from US$6 to US$23 in the private information treatments without communication, from US$13 to US$36 in the private information treatments with communication and from US$14 to US$31 in the public information treatments. Each subject also earned an additional US$5 show-up fee.

[25] This implies that the inter-arrival time between opportunities follows a geometric distribution. This is an approximation to the exponential distribution in continuous time.

Session	Information	Communication	Subjects	Treatment*	Matches	Games
1	Private	No	12	Low-Slow	4	24
1	Private	No	12	Low-Fast	4	24
1	Private	No	12	High-Fast	4	24
1	Private	No	12	High-Slow	4	24
2	Private	No	16	High-Slow	4	32
2	Private	No	16	High-Fast	4	32
2	Private	No	16	Low-Fast	4	32
2	Private	No	16	Low-Slow	4	32
3	Private	No	16	High-Slow	4	32
3	Private	No	16	Low-Slow	4	32
3	Private	No	16	Low-Fast	4	32
3	Private	No	16	High-Fast	4	32
4	Private	No	12	Low-Fast	4	24
4	Private	No	12	High-Fast	4	24
4	Private	No	12	High-Slow	4	24
4	Private	No	12	Low-Slow	4	24
5	Public	No	12	Low-Slow	4	24
5	Public	No	12	Low-Fast	4	24
5	Public	No	12	High-Fast	4	24
5	Public	No	12	High-Slow	4	24
6	Public	No	16	High-Slow	4	32
6	Public	No	16	High-Fast	4	32
6	Public	No	16	Low-Fast	4	32
6	Public	No	16	Low-Slow	4	32
7	Private	Yes	12	Low-Slow	4	24
7	Private	Yes	12	Low-Fast	4	24
7	Private	Yes	12	High-Fast	4	24
7	Private	Yes	12	High-Slow	4	24
8	Private	Yes	12	High-Slow	4	24
8	Private	Yes	12	High-Fast	4	24
8	Private	Yes	12	Low-Fast	4	24
8	Private	Yes	12	Low-Slow	4	24

Exchange Rate : US$ 1 was worth 160 points earned in the experiment.
Subjects (n) : Number of subjects in a session.
Matches (m) : Number of dynamic games played per subject.
Games : Total number of dynamic games played (summing across all subjects). As a game is played between two persons, there are $\frac{m \times n}{2}$ dyadic games in total.
*: in the order in which it was presented to participants.

Table 2.2: Treatments and sessions

discounted payoffs were induced by a random termination rule with the draw of an integer from a uniform distribution over the range [1,100]. If the draw was 100, the match was ended. Thus, the probability that a match continues is 0.99 and regardless of how much time has already elapsed, the match is still expected to last another $\frac{1}{1-0.99} = 100$ seconds. This is equivalent to an infinite horizon where the discount factor attached to future payoffs is 0.99 per second. Thus, 1 second is equivalent to a "period". The key methodological innovation of this study is to let participants switch between the two action choices asynchronously in continuous time, and hence they don't have to make a decision every period. This allows the implementation of a dynamic game with two essential features of the situation without immediate reciprocity. First, interactions are infrequent, and hence the design is different from the "flow payoff" design in continuous time. This is controlled by the probability of opportunity arrival for each participant. Second, the time horizon is long enough so that individuals are given the chance to reciprocate. The discount factor in this game is very high and the horizon is also significantly high.

Finally, the screen (see Figures 2.2 and 2.3) also showed in real time the number of times the subject invested out of the total opportunities he received, number of times his partner invested (and also the number of opportunities his partner received in the public information treatments), the difference in the number of investments, the time since last investment by the subject and his partner, the subject's payoff, his partner's payoff and the payoff difference. It is easy to see that payoffs are updated only if an opportunity arrives and the receiver of that opportunity was in an 'INVEST' mode when she received it. Information on these variables other than the 'net favors' variable discussed in Section 2.3.2 was provided because 'time since last favor' seems to be an interesting variable to condition decisions. As subjects knew the average number of times one would get opportunities in a specified time interval, time since last the favor provided becomes an immediate choice of variable whose information should be given to the players. Also, this reduces the experimenter demand effects in the sense that only informing the subjects on 'net favors' might actually induce them to follow a strategy based on that variable. Providing subjects information on multiple variables seems to be better than giving information on only one variable.

Subjects were randomly rematched with a new partner every match. Random rematching protocol allows for a larger number of dynamic games in a session than a protocol where subjects are not paired with each other in more than one match (such as the turnpike protocol as in Dal Bo (2005)). Although the probability of a pair of subjects interacting together in more than one dynamic game is high, this is not likely to cause any problems. Dal Bo and Frechette (2011) mention that their results suggest that the matching protocol does not introduce additional repeated game effects. Also, Dal Bo (2005) uses a turnpike protocol with results consistent with other studies that have used random matching protocols.

2.4.2 Hypotheses

The central hypotheses are listed in this section. To simplify the discussion, whenever a reference is made to "private information" in this study, it necessarily means the situation without pre-play communication unless otherwise mentioned. While the most efficient perfect public equilibrium cannot attain the first-best payoff, it can achieve very high efficiency levels under most parameter configurations. We simply define "efficiency" in terms of the maximum payoffs possible, i.e., as a percentage of the first-best payoffs. In at least three out of four situations considered in this study, more than 90% efficiency could be achieved. The first hypothesis is concerned with what levels of favor provision can one expect to observe in the data and how close the levels would be to the level under the most efficient PPE. Performing a favor can be thought of as taking a risky action by trusting the other person and this risk is rewarded when the other person returns the favor. While obviously a dynamic favor exchange game is more complicated than just a repeated trust game, one can form the hypothesis partially based on the results from some trust games implemented in the laboratory. In these games, high levels of reciprocity and "favor granting" is observed. While in non-repeated trust games with random rematching, Berg et al. (1995) show that more than 90% of the subjects show trust and a lot of that is repaid, and high levels of trust are also displayed in repeated settings, as found in Engle-Warnick and Slonim (2006)[26]. Also, the parameters of the experiment are such that if no one provides any favors then both subjects earn zero. That is, to obtain positive payoffs one must receive favors and to receive favors it is expected that one must also return the favors at some point. Thus, one can expect substantial exchange of favors and the first hypothesis can be formed as below.

Hypothesis 1 ***(Efficiency Hypothesis)***. *When the opportunities to provide a favor are privately observed, favor provision is approximately equal to the level under the most efficient PPE.*

The next set of two hypotheses are concerned with the comparative statics of changing the benefit of receiving a favor and the arrival rate of opportunities in the favor exchange game with private information and no communication. With high benefit, the gains from cooperation are higher and hence, the favor exchange relationship is more valuable. Thus, intuition suggests that a higher benefit should drive individuals to do more favors compared to a low benefit scenario. While the presence of multiple equilibria does not allow us to tie this intuitive hypothesis to any specific equilibrium, the best chips mechanism does

[26]However, "reciprocity" is not an "equilibrium" in those games. In contrast, in favor exchange games, there is positive amount of reciprocity in many equilibria.

predict a similar effect, while the simple chips mechanism does not predict any change in the efficiency with respect to a change in the benefit. Based on intuition and insights from the BCM, the next hypothesis is as follows.

Hypothesis 2a ***(Benefit Hypothesis)***. *When the opportunities to provide a favor are privately observed, a higher benefit of receiving a favor results in a rate of favor provision which is higher than the rate under the situation with low benefit.*

Next, a higher arrival rate of opportunities means more interaction between the individuals that might induce higher cooperation in the form of favor provision as foregoing opportunities become difficult. At the very extreme, an infinite rate of arrival would make the informational constraint vanish and detecting whether a partner had foregone an opportunity to provide a favor becomes possible with certainty. The BCM predicts an increase in efficiency when the arrival rate is increased. It can also be conjectured that the equilibrium set is weakly monotonic in p, as shown in Nayyar (2009) for a similar model with partial favors allowed. Thus, the following hypothesis is obtained.

Hypothesis 2b ***(Frequency Hypothesis)***. *When the opportunities to provide a favor are privately observed, a faster arrival rate of opportunities results in a higher rate of favor provision than the rate under the situation with a slower arrival rate.*

Two other variations in the basic environment are examined. Under the first variation, individuals are allowed to observe the opportunities received by their partners. This situation can be termed as the ‘public information’ or ‘perfect monitoring’ treatment. The main interest is in investigating whether the perfect monitoring of partner’s actions has any effect on the favor provision among individuals. While the first-best payoff cannot be supported under any PPE in the scenario with private information, it is possible to achieve 100% efficiency under public observability of opportunities. Also, some experiments in the past have shown that perfect monitoring of partner's actions enhances efficiency compared to the situation with imperfect monitoring in prisoners' dilemma and related games (Feinberg and Snyder (2002), Aoyagi and Frechette (2009)). Thus, one can expect higher rates of favor provision when individuals can observe the opportunities received by their partners. This implies the following hypothesis.

Hypothesis 3 ***(Monitoring Hypothesis)***. *For each parameter configuration, the rate of favor provision under public information is higher compared to the rate under private information.*

Next, comparison is made between the situation with and without a pre-play communication phase when the opportunities are privately observed. While pre-play communication can certainly help subjects signal their strategies and understand partners' strategies, it might also aid in deceiving others by lying. In one-shot games, non-binding cheap talk has been found to facilitate both coordination (Cooper et al. (1989), Cooper et al. (1992)) and cooperation (Duffy and Feltovich (2002)). However, it is not clear that one can generalize these results to the dynamic game under consideration. On the other hand, there is no specific reason to assume that cheap talk would "hurt" the favor exchange relation and also no experimental evidence exists where cheap talk results in lower efficiency in a repeated game setting. Thus, one can conjecture that the inclusion of a cheap-talk phase does not significantly affect the exchange of favors among individuals. Hence, we have the following hypothesis.

Hypothesis 4 ***(Communication Hypothesis)***. *For each parameter configuration, the rate of favor provision under private information with pre-play communication is similar to the rate under private information without any communication.*

As more favor provision translates into higher payoffs for subjects, all the hypotheses concerned with efficiency can also be re-cast in terms of subjects' payoffs with respect to the highest attainable payoffs.

The last hypothesis is regarding the strategy employed by the participants when opportunities are privately observed and there is no pre-play communication. As discussed earlier, the literature on favor exchanges proposes chips mechanisms as "specific strategies". But, even in that class of strategies, there are multiple equilibria. Consistent with the *Efficiency Hypothesis*, the next hypothesis conjectures that subjects follow the best chips mechanism.

Hypothesis 5 ***(Behavioral Hypothesis)***. *When the opportunities to provide a favor are privately observed, participants follow the BCM.*

2.5 Results

The findings of this study are reported in this section. The overall results about efficiency are presented in Subsection 2.5.1 while some findings about player strategies are documented in Subsection 2.5.2. Subsection 2.5.3 collects other results on the inequality in payoffs and the volatility of play. Appendix B contains additional data analysis.

2.5.1 Rate of Favor Provision: Efficiency

The first objective is to study what fraction of the favors is performed by the subjects in the environment when the opportunities to provide a favor are privately observed and there is no pre-play communication. Each treatment has observations on 112 dynamic games under this private information environment and each game differs in the length for which it was played. The reported data are taken at one second intervals. This is the most obvious way to report data as the random termination rule is implemented after each second, and also opportunities can't arrive more than once during a second. Hence, using a second as the unit of time keeps things consistent and thus, the action of a subject each second is to either be in a 'Do favor' mode or in a 'Do not do favor' mode.

Treatment	Private Info. Data	Simple CM	Best CM	Number of Obs.
Low-Slow	62.42	68.90	83.24	18406
Low-Fast	61.47	67.53	90.46	16686
High-Slow	71.66	67.08	91.22	23824
High-Fast	74.99	67.92	94.27	17440

The unit of observation is a subject per second. The mean favor provision rate for a treatment is the average across all seconds and all subjects in the treatment.

Table 2.3: Mean favor provision rates

Table 2.3 shows the time spent by an average subject in the 'Do favor' mode as a proportion of the total time spent in both modes. Thus, it gives the percentage of favor provision[27] by treatment, aggregating over time and all subjects across sessions. Figure 2.4 plots the instantaneous probability of granting a favor over time in the four treatments[28]. The graphs collect the overall average behavior over time. Few observations are immediate regarding the characteristics of the average behavior over time in all treatments. First, although not all individuals start out in the 'Do favor' mode, the initial probability of starting in the 'Do favor' mode is high and above 0.8 in all treatments (0.8437 in Low-Slow, 0.8482 in Low-Fast, 0.8348 in High-Slow and 0.8973 in the High-Fast treatment). Second, favor provision decays over time in all of the four treatments. This decline is, however, much steeper in the low benefit treatments as compared to the high benefit ones.

While individuals fail to achieve full efficiency by a remarkable margin, they also attain

[27]Strictly speaking, the percentage of favor provision should equal the number of favors provided divided by the total number of opportunities. However, this measure is almost same as the measure where the time spent by an average subject in the 'Do favor' mode is divided by the total time spent in both modes. Using any measure gives similar results.

[28]Note that although at time 0 (initial time) there are observations on 224 individuals in each of the treatments, observations go on decreasing over time as matches keep on ending probabilistically. Some of the matches were very short, even less than 10 seconds, and some were very lengthy, even more than 200 seconds. The graphs show only up to 100 seconds.

efficiency levels that are significantly lower than the levels predicted under the situation where everyone follows the BCM. This is true for all four treatments. Comparing the overall rate of favor provision observed in the data to the situation (hypothetical, yet possible) where everyone follows the BCM shows that favor provision is lower than the predicted level by 20.82 percentage points in the Low-Slow, 28.99 in the Low-Fast, 19.56 in the High-Slow and by 19.28 in the High-Fast treatments. These differences are statistically significant with p-values of less than 0.001[29]. This is also confirmed from the comparison of the time path of the probability of doing a favor observed in the data to the time path that would have been generated if everyone followed the BCM, as shown in Figure 2.4. Since, the most efficient PPE can attain efficiency levels as large as the ones under the BCM, we have the following result.

Result 1. *Favor provision is significantly less than the level under the most efficient perfect public equilibrium. This rejects the Efficiency Hypothesis (H1).*

Figure 2.4 also shows that both the simple and the best chips mechanisms predict that the probability of granting a favor starts out at one[30], declines over the early few seconds, and then generates a fairly "stationary play". As expected, the decline of the likelihood of granting a favor is steeper in the SCM as compared to the BCM. It is also observed that the initial decline in the probability as seen in the data is far less steep than the SCM. Table 2.3 also displays the favor provision rates that would be seen if everyone followed the SCM. Favor provision is lower than that for the SCM in the low benefit treatments, whereas

[29]These p-values are obtained from a two-tailed unpaired Mann-Whitney U test with each observation being the time spent by each subject in the 'Do favor' mode as a proportion of the total time spent in both modes in a treatment. Similar results are obtained, if instead, a difference in means t-test is used with clustering at the pair/subject level and controlled for a time trend. Since the correlation structure of the data allows the usage of tests depending on which type of correlation is assumed, these tests are performed throughout this study for robustness issues and mentioned wherever they show different results.

[30]As net favors are zero at time zero, all chips mechanisms predict that each player should start out in the 'Do favor' mode.

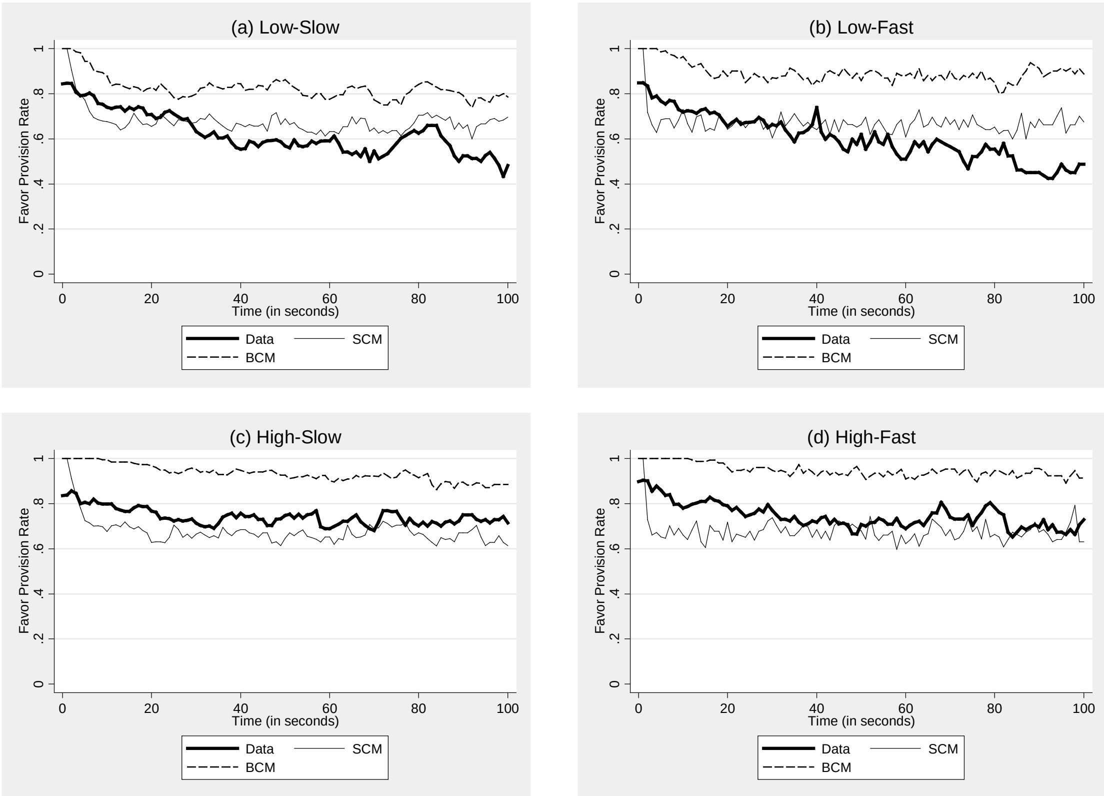

Figure 2.4: Rate of favor provision over time: comparison with the best and simple chips mechanisms

it is higher in the high benefit treatments. However, these differences are not statistically significant with p-values > 0.1[31], except in the High-Fast treatment where the p-value is < 0.01. Thus, there is no clear ranking of efficiency under the SCM and what is observed in the data.

The four treatments considered give rise to two sets of binary comparative statics exercises, one in benefit and the other in the arrival rate of opportunities. It can be readily inferred from Table 2.3 that the levels of favor provision are higher in the high benefit

[31] These p-values are obtained from a two-tailed unpaired Mann-Whitney U test with each observation being the time spent by each subject in the 'Do favor' mode as a proportion of the total time spent in both modes in a treatment. Similar results are obtained, if instead, a difference in means t-test is used with clustering at the subject level and controlled for a time trend. When clustering is done at the pair level, however, the differences are then significant with p-values < 0.05 in Low-Fast and High-Fast treatments. For the treatments with slow arrival rate, the differences are still indistinguishable statistically.

treatments as compared to the low benefit treatments (71.66% vs. 62.42% and 74.99% vs. 61.47%). These differences are statistically significant with p-value < 0.1 for the slow arrival rate treatment and with p-value < 0.01 for the fast arrival rate treatment. Further evidence is documented in panels (a) and (b) of Figure 2.5 that plot the instantaneous probability of granting a favor over time in these treatments. The time path for the high benefit scenario is always above the time path for low benefit one.

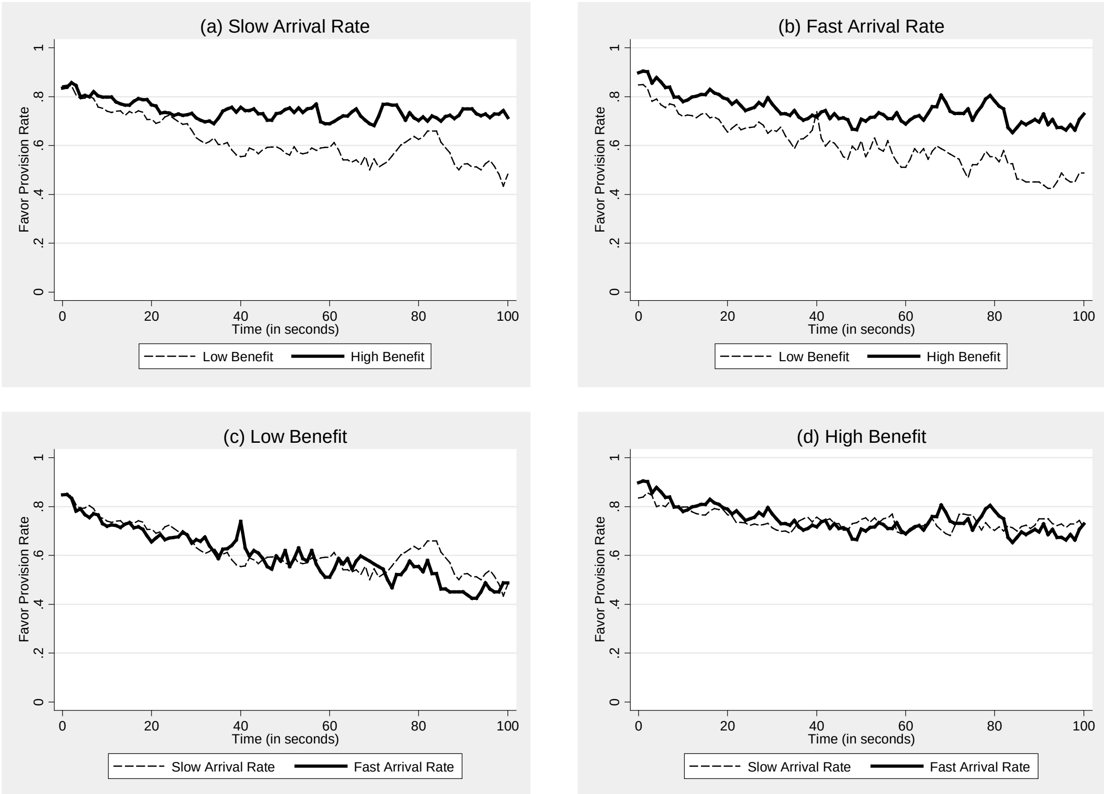

Figure 2.5: Rate of favor provision over time: comparative statics

Table 2.3 also shows that changing the arrival rate of opportunities does not have any aggregate effect on the level of favor provision. The differences (0.95 points in the low benefit treatment and 3.33 in the high benefit treatment) are statistically indistinguishable with p-values > 0.1[32]. Panels (c) and (d) of Figure 2.5 give further evidence of this; the time paths under slow and fast arrival rate cannot be distinguished. However, it is obvious that the overall payoffs realized by the individuals are higher in the faster arrival rate treatment than the slower arrival rate treatment for a given level of benefit. This is because a greater absolute number of favors are exchanged even though the efficiency level is similar. Thus, supporting the *Benefit Hypothesis (H2a)* and rejecting the *Frequency Hypothesis (H2b)*, we have the following result.

Result 2. *When the opportunities to provide a favor are privately observed, a higher benefit of receiving a favor results in a higher rate of favor provision than when benefit is low (support for H2a) but a faster arrival rate of opportunities generates similar rate of favor provision as when the arrival rate is slow (reject H2b).*

While the positive effect of a change in benefit on favor provision is intuitive, less obvious is the result that the overall rate of favor provision is invariant to the change in arrival rate of opportunities. These are obtained for intermediate values of arrival rates (that is, far from the extreme values of zero and infinity). But, when the arrival rate is very high such that people often meet with each other, one might expect to see high levels of favor provision as then it would be more likely that the partner is foregoing opportunities to provide a favor when no favors are received. In fact, at the very extreme, when individuals meet too often (approximated by an 'infinite' arrival rate) then this favor exchange game collapses into a prisoners' dilemma game, where the payoffs from 'not cooperating' are $(0, 0)$. But, as the aim of this study is to analyze situations where cooperation is without immediate

[32] All the p-values in the analysis of comparative statics of changing the benefit/arrival rate of opportunities are obtained from a two-tailed paired Wilcoxon rank sum test. An observation is the time spent by each subject in the 'Do favor' mode as a proportion of the total time spent in both modes in a treatment. The reason for using a paired test is that it acknowledges the fact that each subject is in both the treatments for which the test is performed. Similar results are obtained if instead a difference in means t-test is used with clustering at the pair/subject level and controlled for any time trend.

reciprocity, only intermediate values of the arrival rate are considered.

The next question posed is the following: Is there any efficiency gain from letting individuals observe the opportunities received by their partners? This situation is also referred to as the 'public information' environment. Focusing on the differences in the initial rate of favor provision[33] across the private and public information environments show that the initial favor provision rate is higher for public information treatments. It is 84.37 versus 88.39 in the Low-Slow, 84.82 versus 92.86 in the Low-Fast, 83.48 versus 99.11 in the High-Slow and 89.73 versus 94.64 in the High-Fast treatment. Table 2.4 shows the average percentage of favor provision by treatments in both the private and the public information environments. Rate of favor provision is significantly higher with monitoring than without monitoring in each of the four treatments. Among the low benefit treatments, the difference is 14.03 percentage points when the arrival rate of opportunities is slow and 21.38 points when the arrival rate is fast. Among the high benefit treatments, this difference is 17.52 percentage points when the arrival rate of opportunities is slow and 12.00 points when the arrival rate is fast. All differences are significant with p-values < 0.05[34].

Figure 2.6 displays the overall rate of favor provision over time for the private information and public information cases for each of the four treatments. The time path of average rate of favor provision under public information is always above the time path under private information. Thus, we see that the informational constraint results in significant efficiency losses. Supporting the *Monitoring Hypothesis (H3)*, Table 2.4 together with Figure 2.6 provide the following result.

Result 3. *Favor provision is significantly higher if individuals can observe the oppor-*

[33] It is defined as the proportion of time zero 'Do favor' choices to the total number of decisions.

[34] The p-values are reported from a two-tailed unpaired Mann-Whitney U test with each observation being the time spent by each subject in the 'Do favor' mode as a proportion of the total time spent in both modes in a treatment. The p-values change slightly when the differences in means t-tests are used with clustering at the pair/subject level and include a time trend. But all values are < 0.05.

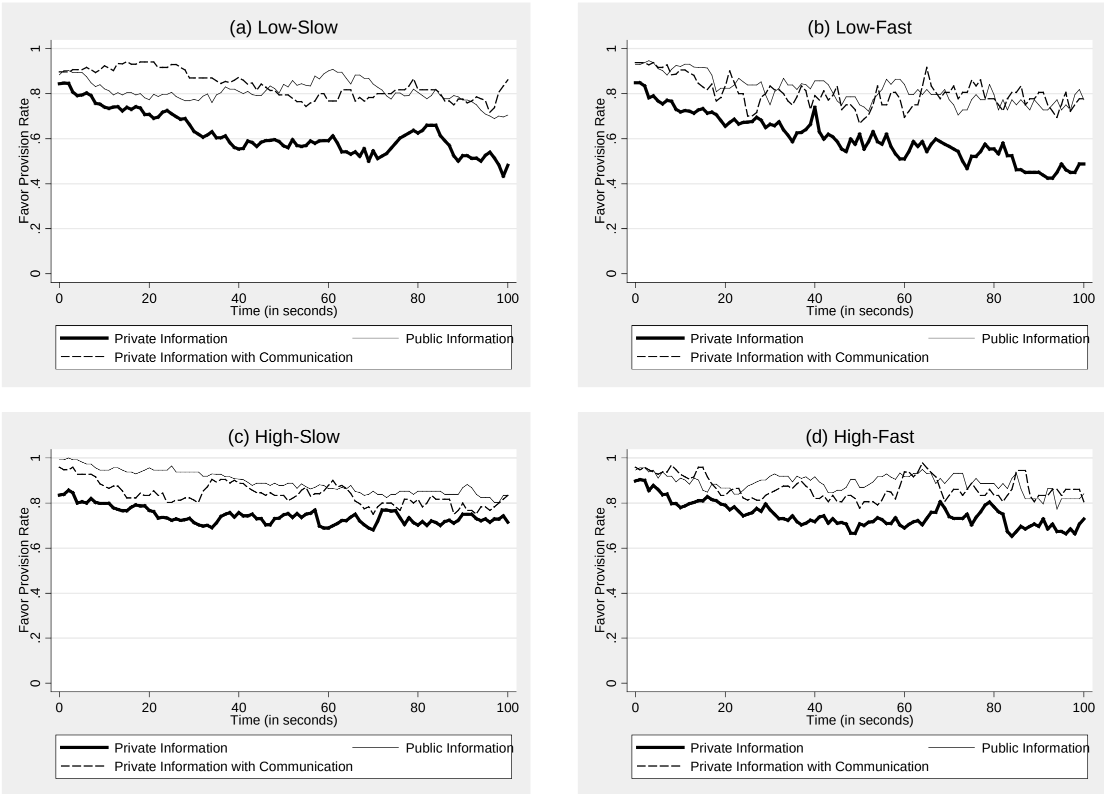

Figure 2.6: Rate of favor provision over time: comparison with the "Private Information with Communication" and comparison with "Public Information"

tunities received by their partners than when they are privately observed (support for H3).

Treatment	Private Info.	Private Info. with Communication	Public Info.
Low-Slow	62.42 [18406]	78.08 [11232]	76.45 [13682]
Low-Fast	61.47 [16686]	79.54 [7942]	82.85 [9482]
High-Slow	71.66 [23824]	83.43 [12214]	89.18 [14532]
High-Fast	74.99 [17440]	83.46 [12000]	86.99 [13992]

Number of observations in parentheses. The unit of observation is a subject per second.

Table 2.4: Mean favor provision rates: "Private Information", "Private Information with Communication", and "Public Information"

The rates of favor provision under public information treatments can be compared to the cooperation rates in the treatments with stochastic horizon in Bigoni et al. (2011)[35]. The

[35]The comparisons made here are only an attempt to relate to their study (for cooperation in continuous

mean rate of cooperation reported in their study is 52.3% for the short (20 seconds expected duration) and 66.9% for the long treatments (60 seconds expected duration). In the current study, the mean cooperation (or favor provision) rate across all treatments for the public information case is 84.06%. Thus, the cooperation rate observed in the current study is higher. However, it is also true that the expected duration of the interaction is higher (100 seconds)[36]. They also report an initial cooperation of 65.9% in the short treatments and 75.1% in the long treatments. Thus, they find that as the expected duration increases from 20 to 60 seconds, the share of initial cooperators rises. In the current study, this share is 93.75% with an expected duration of 100 seconds. This seems to be in line with their result. Also, they report that cooperation within a match declines as time progresses which is also observed in the present study.

The final result in this section aims at comparing the favor provision rates with and without pre-play communication. Table 2.4 documents the average percentage of favor provision by treatments in the private information environment with and without communication. Clearly, the rate of favor provision is significantly higher with communication in all four treatments. The difference is 15.66 in Low-Slow, 18.07 in Low-Fast, 11.77 in High-Slow and 8.47 in High-Fast treatments. The p-values are all < 0.05, except in the Low-Slow treatment (p-value is 0.0561)[37]. The initial rate of favor provision is also higher for the treatments with communication. It is 84.37 versus 89.58 in the Low-Slow, 84.82 versus 93.75 in the Low-Fast, 83.48 versus 95.83 in the High-Slow and 89.73 versus 95.83 in

time with an indefinite time horizon). However, it should be clearly noted at the outset that the setting in a favor exchange game is different from that of the prisoner's dilemma with flow payoffs.

[36]The implied discount factor in the present study is 0.99 with 1 second as the "period". However, the implied discount factor is 0.992 (0.9973) in the short (long) treatment in the study by Bigoni et al. with 16/100 second as the "period".

[37]These p-values are reported from a two-tailed unpaired Mann-Whitney U test with each observation being the time spent by each subject in the 'Do favor' mode as a proportion of the total time spent in both modes in a treatment.

the High-Fast treatment.

Figure 2.6 displays the overall rate of favor provision over time for the private information treatments without communication as well as with communication. The time path of average rate of favor provision in the private information treatment with communication is always above the time path under private information without communication. Thus, rejecting the *Communication Hypothesis (H4)*, we have the following result.

Result 4. *When the opportunities to provide a favor are privately observed, letting individuals participate in pre-play non-binding communication results in a significantly higher level of favor provision than when they are not allowed to communicate (reject H4).*

Although the theory doesn't inform us as to what should be the effect of monitoring or allowing communication in these type of games with private information, the experiments reported in this study show clearly that efficiency could be significantly enhanced in both scenarios. While an individual can directly observe the intentions of her partner regarding favor provision when there is perfect observability, in the other situation with pre-play communication but with private information one could still focus and coordinate on specific types of behavior and gain some understanding about what to expect from one's partner. It lowers the loss in efficiency by reducing the strategic uncertainty. It remains to be seen whether allowing both perfect monitoring and pre-play communication can drive the favor provision level closer to the full efficiency level.

2.5.2 Individual Strategies

The aim of this section is to gain more insights into the strategies players use in a favor exchange game. Participants were given real time information on the total number of favors done by them, the number of opportunities they had, the number of favors received, the net number of favors done by them, the time since the last favor done by them, the time since last favor done by their partner[38], their payoff, partner's payoff and the difference in payoffs. The information on some of these variables could have been derived from the already given feedback on other variables; for example, the payoff variable can be calculated given the information on the total favors done and total favors received. In spite of that, since time is continuous, it is extremely hard for the participants to calculate the derived information on payoffs and after all, since it is the payoff that finally matters, explicit information was provided on this variable too. This way it made it possible for participants to condition their strategies on a specific variable if they wanted to do so. However, they need not pay attention to each and every piece of information on screen.

To understand how the state variables affect the probability of granting a favor, a pooled probit regression is run with standard errors clustered at the subject level. The dependent variable is 'mode of person i at time t ($mode_{it}$)' which is a binary variable taking value one if individual is in a 'Do favor' mode and the independent variables include 'the mode of person i at time $t-1(mode_{it-1})$' (binary variable with one if mode is 'Do favor'), 'net favors provided by person i at time $t-1(netfavors_{it-1})$', 'time since last favor provided by partner at time $t-1(timeelapsed_{it-1})$', 'own payoff at time $t-1(ownpayoff_{it-1})$' and time. As already discussed in Section 2.4.1 all these variables were updated in real time on each subject's screen.

The marginal effects are summarized in the Tables 2.5 and 2.6 for each of the treatments. The regression results are displayed for two different sub-cases depending on whether the participant is ahead or behind and tied. A participant is termed 'ahead' ('behind') if she has done more(less) favors than her partner till the previous second of play. If she has done

[38]The term "partner" was not used in the experiment.

exactly the same number of favors as her partner till the previous second of play (or nobody has yet done a favor from the start of the play) then she is termed as being ‘tied’.

variable	Low-Slow	Low-Fast	High-Slow	High-Fast
mode_{t-1}	0.89	0.75	0.87	0.74
	(0.008)***	(0.015)***	(0.011)***	(0.016)***
netfavors_{t-1}	-0.027	0.012	-0.016	0.007
	(0.017)	(0.006)**	(0.004)***	(0.004)*
timeelapsed_{t-1}	-0.003	-0.009	-0.002	-0.004
	(0.001)***	(0.002)***	(0.0003)***	(0.001)***
ownpayoff_{t-1}	0.0005	0.0016	0.0003	0.0004
	(0.0006)	(0.0003)***	(0.0001)***	(0.0001)***
time	-0.0005	-0.0018	-0.0005	-0.002
	(0.0001)***	(0.0002)***	(0.0001)***	(0.0003)***
no. of observations	11621	9556	14126	9617
no. of groups	56	56	56	56
pseudo-R^2	0.76	0.53	0.76	0.60

* significant at 10%; ** significant at 5%; *** significant at 1%

Pooled probit regression with clustering at the subject’s level; standard errors in parentheses. The unit of observation is a subject per second; dependent variable: mode_t

Table 2.5: Marginal effects for “Private Information”: behind and tied

As is expected, since the experiments are in continuous time and data is taken at one second intervals, the choice of mode at time t is affected immensely by the choice at $t-1$ (persistence or inertia). Next, the effect of net favors on the likelihood of granting a favor is, at best, inconclusive. The signs are positive for some of the coefficients and negative for others. It is also insignificant in most of the treatments when the individual is ahead. The negative signs in front of the coefficients of the variable $timeelapsed_{t-1}$ in all of the probit regressions show that the probability of granting a favor declines as time passes without receiving a favor in return from partner. Also note that the decline is faster (a) in the

variable	Low-Slow	Low-Fast	High-Slow	High-Fast
$mode_{t-1}$	0.90	0.78	0.91	0.80
	(0.01)***	(0.02)***	(0.01)***	(0.02)***
$netfavors_{t-1}$	0.009	0.009	0.004	0.008
	(0.01)	(0.01)	(0.01)	(0.004)**
$timeelapsed_{t-1}$	-0.005	-0.009	-0.004	-0.01
	(0.001)***	(0.002)***	(0.001)***	(0.002)***
$ownpayoff_{t-1}$	0.005	0.002	0.001	0.0004
	(0.002)***	(0.0004)***	(0.0003)***	(0.0001)***
time	-0.0001	-0.0015	-0.0005	-0.0019
	(0.0005)	(0.0004)***	(0.0004)	(0.0004)***
no. of observations	6561	6906	9474	7599
no. of groups	55	54	55	55
pseudo-R^2	0.76	0.53	0.75	0.57

* significant at 10%; ** significant at 5%; *** significant at 1%

Pooled probit regression with clustering at the subject's level; standard errors in parentheses. The unit of observation is a subject per second; dependent variable: $mode_t$

Table 2.6: Marginal effects for "Private Information": ahead

treatments with high arrival rate of opportunities than when the arrival rate is slow, and (b) when the participant is 'ahead' than when she is 'behind/tied'. The marginal effects are significant with p-values < 0.01 for each of the treatments and also in both the cases where the participant is ahead or behind and tied. So, we have the following result.

Result 5. *When the opportunities to provide a favor are privately observed, the likelihood of favor provision by an individual declines with an increase in the time since the last favor provided by her partner, but is mostly insensitive to the net favors provided by her. This rejects the Behavioral Hypothesis (H5).*

Elaborate discussion on how net favors and time since last favor by partner affects the rate of favor provision is provided in sections 2.5.2.1 and 2.5.2.2. Discussion on the classification of subjects into some simple "types" is given in subsection 2.5.2.3.

2.5.2.1 Net Favors and Chips Mechanisms

Figure 2.7 displays the distribution of the net favors provided by an individual for the switches between the two decision modes. It doesn't reveal any distinctive pattern with respect to the net favors provided by the individual for the two different switches. Tables 2.5 and 2.6 already show that on average, *favor provision is non-monotonic in the net favors provided by an individual. In fact, a closer look at the individual level decisions also reveals that none of the participants employ any of the chips mechanisms or close-to-chips mechanisms.* This is verified from checking whether $Prob(mode_{it} = 1) = 1$ for $netfavors_{it-1} < k^*$ and $Prob(mode_{it} = 1) = 0$ for $netfavors_{it-1} = k^*$ for each individual. The k^* varies from 1 to k^{max}. For close-to-chips mechanisms, it is verified whether $Prob(mode_{it} = 1) \geq 0.8$ for $netfavors_{it-1} < k^*$ and $Prob(mode_{it} = 1) \leq 0.2$ for $netfavors_{it-1} = k^*$. Thus, the *Behavioral Hypothesis* is rejected. Also, it can be readily inferred from Figure 2.7 that more than 40% of the switches were done when the individual was behind or tied.

Behavior is however, monotonic in a "weak" sense. Table 2.7 documents the mean rate of favor provision in each treatment separately for the three situations when the individual is ahead or tied or behind. It is easy to see that the overall likelihood of favor provision is highest when the individual has done fewer favors than her partner as compared to the situation when she has done the same number of favors. It is lowest when she has done more favors than her partner.

Another question is whether the behavior is monotonic or not for each individual in the four treatments separately. Analysis suggests that 19 out of 56 (34%) subjects exhibit behavior that is monotonic in net favors in all four treatments, 21 out of 56 (37.5%) in

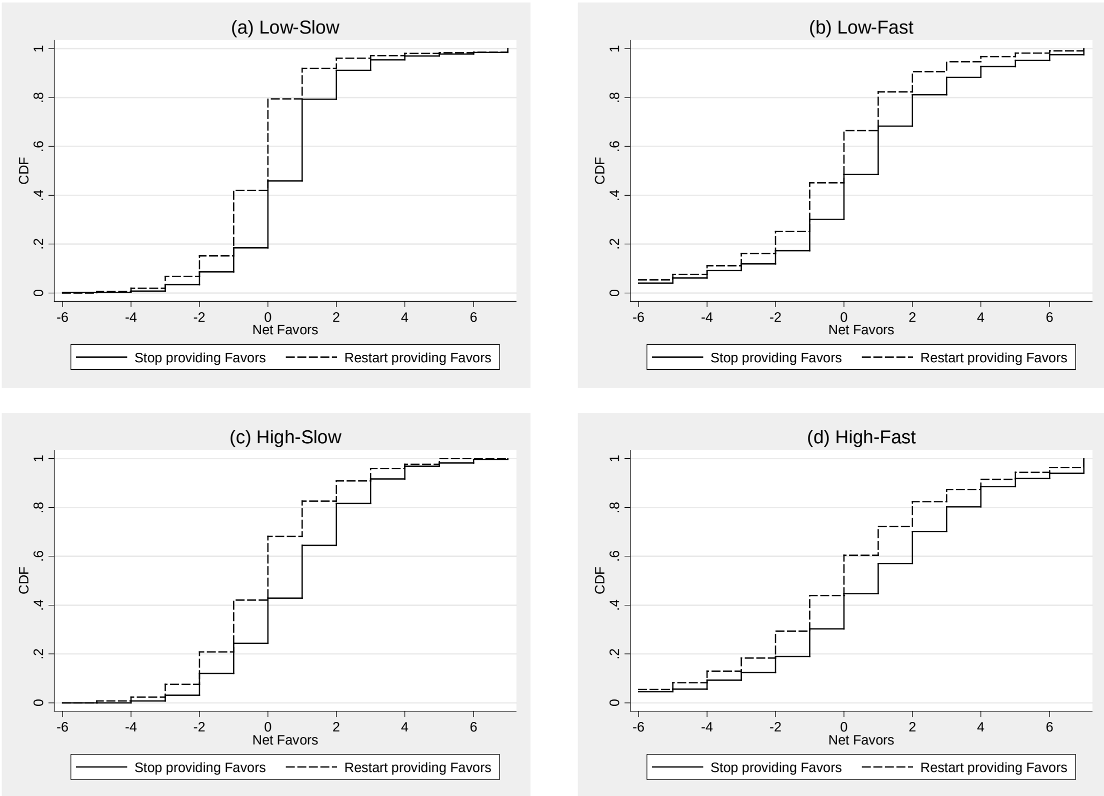

Figure 2.7: Switching pattern - CDF of net favors provided

Condition	Low-Slow	Low-Fast	High-Slow	High-Fast
ahead	52.70	55.82	62.72	70.32
tied	58.96	63.29	71.48	75.55
behind	74.70	65.87	80.72	79.31

Table 2.7: Mean favor provision by net favors

exactly three treatments, 11 out of 56 (19.6%) in exactly two treatments, and all subjects in at least one treatment. Behavior is classified as monotonic if the following is satisfied: $Prob(mode_{it} = 1|k_t \leq 0) \leq Prob(mode_{it} = 1|k_t = 1) \leq Prob(mode_{it} = 1|k_t = 2) \leq .. \leq Prob(mode_{it} = 1|k_t = k^{max})$ where k_t is net favors provided by the individual till time t. However, this obviously includes the cases where an individual always provides favors as well as situations where an individual is almost always behind or tied and only has few observations where she is ahead but still satisfies the above condition.

2.5.2.2 Time Since Last Favor Provided

Tables 2.5 and 2.6 along with Figure 2.8(a) show that the likelihood of favor provision declines monotonically in the time since last favor provided by partner. Also evident is the fact that this decline is faster in the treatments where the opportunities arrive at a faster rate. Benefit, however, has a 'level effect'. Similar behavior is documented with respect to the time since last favor provided by the individual. See Figure 2.8(b).

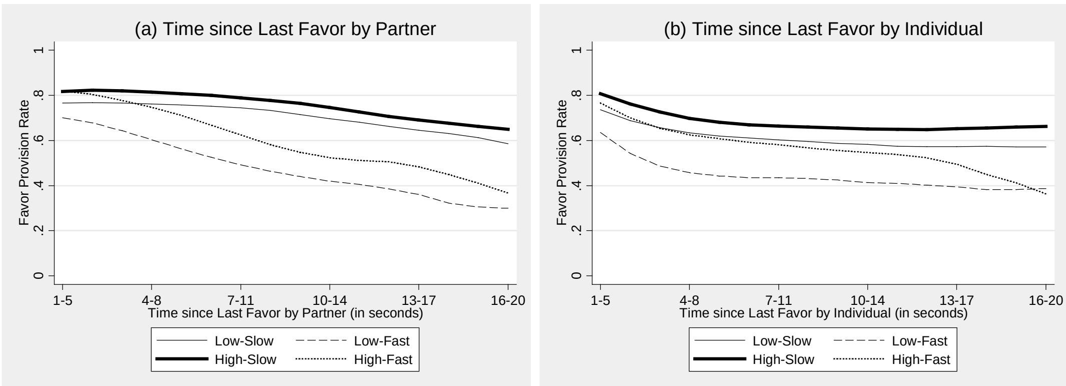

Figure 2.8: Time since last favor provided by individual and partner

Figures 2.9 and 2.10 exhibit the distribution of the variables - time since last favor provided by the individual and her partner for the switches between the two decision modes. More than 70% (90%) of the decisions to stop granting favors in the Slow (Fast) treatments take place when the time since last favor provided by the individual is less than 5 seconds. This clearly shows that *the decision to switch from a 'Do favor' to 'Do not do favor' mode is primarily governed by whether the individual has just provided a favor to her partner recently.*

Likewise, more than 50% (70%) of the decisions to restart providing favors in the Slow

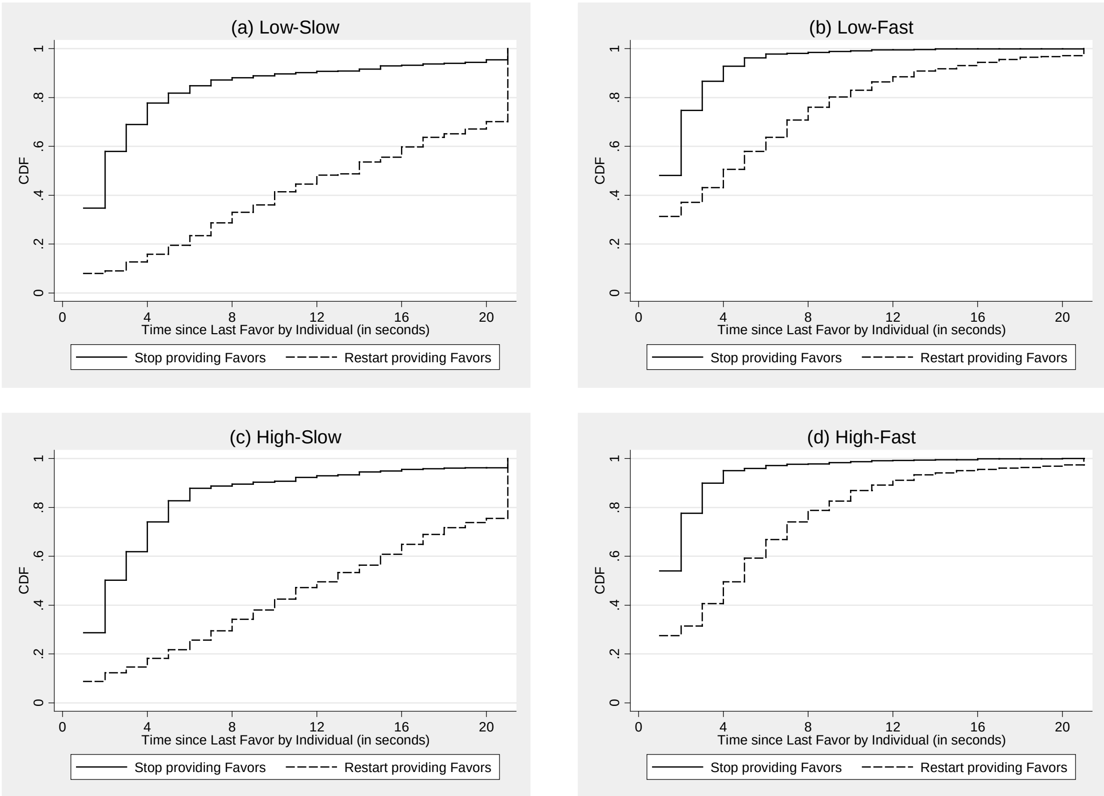

Figure 2.9: Switching pattern - CDF of time since last favor provided by the individual

(Fast) treatments take place when the time since the last favor was provided by the partner is less than 5 seconds. Thus, *whether or not an individual's partner has provided a favor in the recent past explains bulk of the switches from 'Do not do favor' to 'Do favor' mode.*

2.5.2.3 Classification of Individuals into Simple Types

Individuals can be classified into five broad types depending on the average time they spend in the 'Do favor' mode. These types are as follows along with the average time spent in 'Do Favor' mode (in percentages): (1) Very Low - [0, 20), (2) Low - [20, 40), (3) Medium - [40, 60), (4) High - [60, 80), and (5) Very High - [80, 100]. Table 2.8 shows the proportion of these types in each of the four treatments. Two observations are immediate. First, there

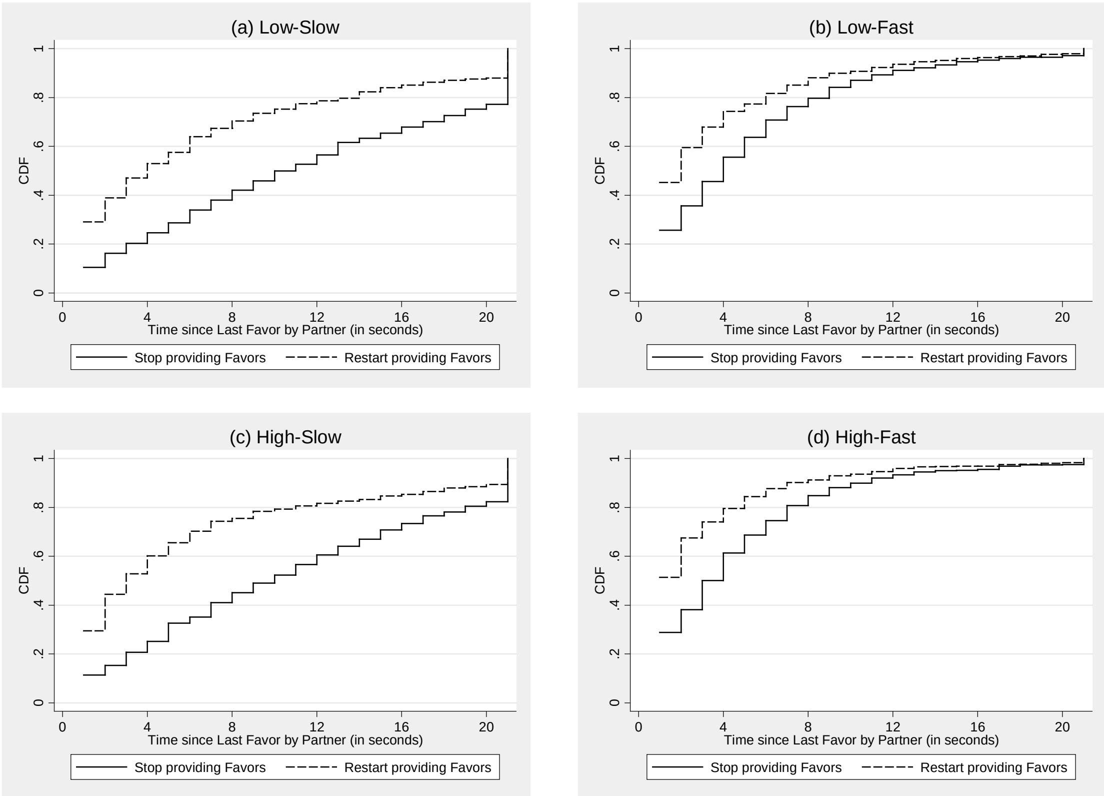

Figure 2.10: Switching pattern - CDF of time since last favor provided by partner

is no apparent difference in the distribution of the types in the slow and fast arrival rate treatments. Second, the high benefit treatments have a higher proportion of high and very high types than the low benefit treatments but smaller proportion of very low, low and medium types.

Type	Low-Slow	Low-Fast	High-Slow	High-Fast
very low	1.8	1.8	3.6	3.6
low	16.1	16.1	5.4	3.5
medium	26.8	21.4	19.6	12.5
high	21.4	28.6	33.9	39.3
very high	33.9	32.1	37.5	41.1

Table 2.8: Proportion of types in each treatment. All entries are in percentages.

Some further details of the classification are as follows. First, 2 out of 56 subjects always provide favors in all four treatments, 1 out of 56 in exactly three treatments, 4 out of 56 in exactly two treatments, and 14 out of 56 (25%) in at least one treatment provide favors all the time. Regarding the initial choice to start out in 'Do favor' mode or 'Do not do favor' mode, 29 subjects (52%) start out in the 'Do Favor' mode in all 16 matches, 20 (36%) start out in 'Do favor' mode in 10-15 matches and 7 (12%)in less than 10 matches. Second, *there is a positive correlation among the average time spent by an individual in the 'Do favor' mode and the maximum number of net favors provided by an individual.* The correlation coefficients are 0.35 in Low-Slow, 0.39 in Low-Fast, 0.52 in High-Slow and 0.54 in High-Fast treatments. All values are statistically significant with p-values < 0.01. Third, *on average, there is a positive correlation between the average time spent in 'Do favor' mode and the payoff obtained by the individual in each of the four experimental sessions.* The correlation coefficients are 0.31 in session 1, 0.63 in session 2, 0.28 in session 3 and 0.83 in session 4. Only the values for sessions 2 and 4 are significant with p-values < 0.01. Finally, the two participants who always provided favors did not earn the highest amount of payoff in the respective sessions they were in. In fact, 7 out of 11 other participants earned an amount higher than the first individual who always provided favors and for the second individual there were 2 out of 11 persons who earned more than her.

2.5.3 Inequality and Volatility

This section presents results investigating the inequality in payoffs tolerated among individuals as well as the volatility of the favor exchange relation in the situation where opportunities are privately observed and no communication is allowed.

2.5.3.1 Inequality in Payoffs

An important aspect of a favor exchange relation is that of how much inequality in payoffs is tolerated among individuals and how the average inequality level behaves over time. Obviously, the first step is to define a measure of inequality. In the current study, the inequality between a pair of individuals i and j is defined as the absolute difference of the share of either of the individual's payoffs (π_{it}) in the total payoff at time t and 0.5[39], i.e.,

$$inequality_t^{ij} = \left| \frac{\pi_{it}}{\pi_{it} + \pi_{jt}} - 0.5 \right| = \left| \frac{\pi_{jt}}{\pi_{it} + \pi_{jt}} - 0.5 \right| \quad .$$

If within a pair, no one has yet done a favor to the other person, then this measure is taken to be zero. Thus, it is perfectly possible for both the players to do no favors in the entire play and generate an inequality level equal to zero because both earn nothing.

The chips mechanisms have an upper bound on the inequality that could be tolerated as they are based on a model of even balance. Figure 2.11 shows the average level of inequality over time for each of the four treatments in the private information environment. It also displays the levels of inequality that would have been generated if everyone followed the simple CM, and the best CM and if everyone granted all possible favors (denoted as the 'social optimum'). Clearly, the simple CM is the least unequal. In fact, the level of inequality observed in the data is statistically higher than the level under the simple CM (with p-values[40] < 0.01 in all four treatments). The inequality in the data is statistically indistinguishable from the level of inequality predicted by the situation when each individual follows a best CM as well as the case when everyone grants all favors, that is, under the socially optimum situation (with p-values > 0.1). As one follows the more efficient chips

[39] Using the Euclidean distance instead of the simple absolute value generates similar qualitative results.

[40] These p-values are obtained from a two-tailed unpaired Mann-Whitney U test with each observation being the final value of inequality for each pair of participants.

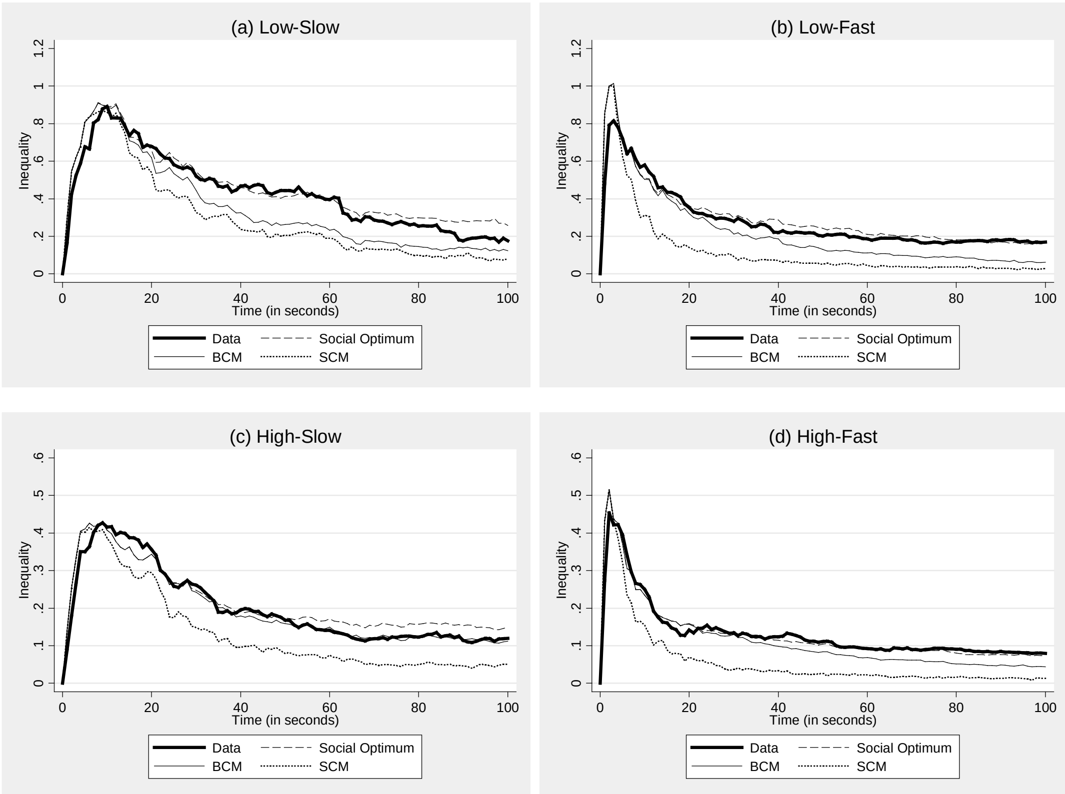

Figure 2.11: Inequality: "Private Information" treatments

mechanism, she is also willing to let the inequality levels rise. Indeed, following the strategy of always providing favors generates the highest level of efficiency but also higher inequality. Thus, there is a trade-off between efficiency and equality of payoffs. To get more of one you need to sacrifice some of the other.

Next, one could compare the inequality levels tolerated in the private information environment with that in the public information environment. Except for the High-Slow treatment, the inequality levels are statistically similar with p-values > 0.1. For the High-Slow treatment, inequality is lower (p-value $= 0.084$) in the private information environment.

Thus, summarizing the above discussion, *when opportunities are privately observed, inequality in payoffs is similar to the level seen under the BCM as well as the situation*

Treatment	Pvt. Info. Data	Simple CM	Best CM	Public Info. Data
Low-Slow	4.12	6.82	3.78	2.26
Low-Fast	10.15	19.89	5.73	4.15
High-Slow	3.71	6.67	1.83	1.29
High-Fast	7.24	19.89	3.40	3.74

Table 2.9: Average switch frequencies

when everyone grants all favors. However, inequality is significantly higher than the level seen under an SCM. Lastly, overall, the inequality in payoffs is statistically indistinguishable among private and public information environments.

2.5.3.2 Volatility of the Relationship

Volatility of play (or the favor exchange relationship) is defined as the frequency of switch between 'Do favor' and 'Do not do favor' modes. This frequency can be measured as the total number of changes in modes divided by the total number of decisions made each second by the participants. The average switch frequencies observed in the data for the private information environment under different treatments are summarized in Table 2.9. It can be readily inferred that people switch less on average when the benefit of receiving a favor is higher than when it is lower (3.71% vs. 4.12% and 7.24% vs. 10.15%). A higher benefit stabilizes the relationship by making it more valuable. Also, Table 2.9 shows that individuals switch actions more often when the opportunities arrive at a faster rate than when they arrive at a slower pace (10.15% vs. 4.12% and 7.24% vs. 3.71%). Thus, *when opportunities to provide a favor are privately observed, the volatility of play is lower if the benefit of receiving a favor is higher or if the opportunities arrive at a slower rate.*

The average switch frequencies as predicted by the simple and the best chips mechanisms are also reported in Table 2.9. Three points are worth noting from comparing these predictions with the frequencies observed in the data. First, for each treatment, the average frequency generated in the data is higher (lower) than the frequency predicted by the best (simple) CM. Thus, it lies between the predictions of the two CMs. Second, the best CM makes correct qualitative predictions about the differences in the switch frequencies. However, it under-predicts the differences quantitatively. Third, although the simple CM is able to correctly predict about the qualitative difference in changing the arrival rate (quantitatively, it over-predicts), it does not predict any difference when the benefit is changed. Thus, it even fails to make a correct qualitative prediction in this case.

Finally, comparing the switch frequencies for the different treatments under private and public information environments suggests that *monitoring the opportunities received by the partner significantly reduces the volatility of play in all four treatments.*

2.6 Conclusion

This chapter analyzed the exchange of favors in the context of a dynamic two-player game with private information. Using a novel experimental design to implement a continuous time dynamic game with a stochastic jump signal process, this study provides insights into a situation where cooperation is without immediate reciprocity. It also contributes to the growing experimental literature on infinitely repeated games with imperfect monitoring of actions. The primary finding is that the level of favor provision is considerably lower than that which individuals could have achieved under the most efficient perfect public equilibrium. This is a robust finding, even under the situation where the benefit is very high compared to the cost of providing favors and when opportunities arrive very fast.

Next, the decisions to stop providing favors and restart providing favors (that is, the switching decisions) can be justified using two competing explanations having distinctive behavioral basis. One hypothesizes that individuals follow one of the chips mechanisms that are simple threshold strategies based on net favors. The other posits that switching decisions are made based on the time since last favor provided by an individual and her partner. The results of this study show that individuals do not employ any of the chips mechanisms. In fact, the likelihood of favor provision does not decline monotonically in the net favors provided. Rather, the decision to stop granting favors seems to depend primarily on whether an individual has just provided a favor in the recent past. Likewise, the decision of an individual to restart granting favors depends mainly on whether her partner has provided a favor recently. Thus, the focus of the individual is more on the time since the last favor was provided by her and her partner. Behavior is explained more by the notion of "What have you done for me lately?" rather than individuals engaging in exact score keeping of net favors.

Efficiency, however, is enhanced by either providing information about partner's opportunities or by allowing individuals to communicate before the start of a bilateral relation. Compared to the situation where opportunities are privately observed and there is no pre-play communication, these two variations help an individual understand the intentions of her partner. While she can directly observe the intentions of her partner regarding favor provision in the situation where opportunities are publicly observed, in the other situation with pre-play communication she can at least gain some information about what to expect from her partner, that is, it helps lower the strategic uncertainty. As Farell and Rabin (1996) note that "talk is cheap (it does not directly affect payoffs), but, given that people respond to it, talk definitely affects payoffs".

There are several interesting directions to pursue in future research. First, it might be worthwhile to conduct a thorough analysis of learning with more matches under a particular parameter configuration and analyze the evolution of strategies. Do individuals learn to pay more and more attention to the net favors provided and engage in exact score-keeping?

Next, the results suggest that individuals condition their strategies on the time since the last favor provided. Thus, it would be interesting to characterize, theoretically, equilibria with the variable "time since last favor". One could also focus on the variable "difference in time since last favor by players" which is the counterpart of "net favors".

Another line of research might consider reversing the direction of private information about opportunities. The setting in this study was such that the players never had to ask for favors. The favors just became available when granted by her partner. It might be interesting to study the behavior in a setting where players have to ask for a favor whenever a need arises. When the need for a favor is private information, the "asking" itself may become an important signal and it remains to be seen whether a player asks for too many favors, even when there is no need for it. These phenomena, like signaling and reputation building, add more realistic features to a model of favor exchange but at the same time complicate the characterization of equilibria.

Fourth, a related model could be explored where the opportunities to provide a favor are publicly observable, however, the cost of providing favors is private information. For example, say cost, $c \in \{c_L, c_H\}$ and $c_L < b < c_H$, where b is the benefit to the recipient.

Fifth, in various political interactions, especially in international relations, it is perhaps more natural for one of the parties to do a disproportionate share of the giving in return for a smaller share of the taking. It would be worthwhile to investigate this situation by implementing different arrival rate of opportunities for the two paired participants. The role of partial favors could be important in these situations. It is not too difficult to allow for situations where favors are perfectly divisible so that individuals can provide fractional favors. To incorporate such a possibility one then has to add a slide bar on the user interface of each subject. The slide bar would represent the fraction of a favor that a subject would wish to grant and would range from zero (no favor) to one (full favor).

Finally, the interest in experiments in dynamic games conducted in a continuous time setting is very recent. The experimental framework developed in this chapter to study the exchange of favors in a dynamic setting with stochasticity induced by the arrival of opportunities can be used to analyze other dynamic games with jump signal processes. For example, the present design seems ideal to study a game between two firms where each of them can produce quantities ranging from zero to a maximum capacity level. However, they can produce only when an opportunity arrives. Payoffs to each firm are functions of the output produced by each firm and are updated only when an opportunity arrives. It would be interesting to track what would be the desired quantities to produce, over time, for each firm and also whether the monopoly output is produced too often or not. Similarly, one could study dynamic games where the stage game is not repeated every period, rather, it arrives according to a Poisson process.

Chapter 3

Cooperation in Revision Games: Evidence from Cournot Duopoly Experiments

3.1 Introduction

Can cooperation be sustained when individuals interact only once? This chapter addresses this question by experimentally implementing Cournot duopoly games with two players where they play this game only once at a pre-specified time but they are involved in a pre-play communication phase. During this phase, they announce the tentative quantities that they desire to produce at the end of the phase and these announcements are observed by both the players. Two types of pre-play communication are implemented. In the first type, the payoffs for the players are determined only by the choices selected at the end of the communication phase. This is known as the real time revision game. The second type of communication is known as the Poisson revision game and works as follows. Although individuals post their desired quantities in real time, opportunities to revise choices arrive according to a synchronous Poisson process and players' payoffs are given by the tentative quantities selected at the last revision opportunity.

This research is motivated in part by a recent paper by Kamada and Kandori (2011) where they analyze Poisson revision games and characterize a class of symmetric equilibria that can sustain partial cooperation in these games. Their interesting finding is that even though players interact only once in the sense that payoffs are determined only by the one-time play of the game, significant efficiency gains could be realized under this new class of equilibria compared to the static Nash equilibrium.

The act of revising actions before the play of the underlying "one-shot" game is prevalent in many areas. It is very common to announce plans for the production of cars few months ahead of actual production in the US motor vehicle and aircraft industries. The trade journal *Ward's Automotive Reports* publishes the firm's announcements of their plans for monthly U.S. production of cars as early as six months before actual production and these plans could be continuously revised until the end of the target month. As emphasized in Calcagno and Lovo (2010), revision of actions is a phenomenon that is practiced in some of the financial markets such as Nasdaq and Euronext. Prior to opening of the market, participants are allowed to submit orders which can be continuously withdrawn and changed until the opening time. During the entire pre-opening phase these orders and resulting prices are publicly posted but only the orders that are still posted at the opening time are binding and are executed. Traders do not always manage to withdraw and submit new orders simultaneously due to technological and other reasons. Other examples would include situations where communication and implementation occur at different times, possibly due to delays. The Poisson revision game implemented in this study is a stylized representation of these types of situations. It models these inefficiencies or imperfections in implementing the intended choices. At the other extreme, real time revision games implement intended actions without any imperfection.

The real time revision phase is in fact related to a cheap-talk period. However, it is different in at least one respect: no external language of communication is used during the revision phase. Players are allowed to signal only through their intended or desired choice of actions, they are not allowed to send any message or participate in any other explicit form of communication. For Poisson revision games, although the revisions are costless and there is no exogenous cost of revising quantities, players face exogenous uncertainty about the effectiveness of the revision. Choices are now binding with some exogenous probability.

There are at least two reasons to use the Cournot game for the present study. First, the unique equilibrium outcome (Cournot-Nash) is inefficient as there exist other action profiles that have higher joint profits than under the choice of Nash equilibrium actions. One can think of producing outputs that are closer to the individual collusive output (the total industry output where joint profits are maximized and divided by two) and lower than the Cournot-Nash as being "more cooperative". Thus, there is a tension between what is efficient for the industry (for both firms) and what is individually best given that the other firm chooses the cooperative action. So, similar to the prisoner's dilemma, cooperation can be studied using a Cournot duopoly game between two individuals. Second, announcing plans and revision of these subsequent plans is often a realistic description for the production industry where firms are competing in quantity and announce plans about future production. Although production is made at the end, there are many factors that influence the final quantity of output produced. Demand and cost uncertainty, information about competitors' planned production, capacity constraints and other reasons might make the final choice of production totally different from the initial plans.

This chapter uses a laboratory experiment to analyze the impact of different forms or technology of revisions on the quantities implemented and hence, on the incidence of collusion. In the laboratory, one can not only observe and control key variables, including the informational feedback available to the players, but also replicate a given scenario multiple times and make causal inferences. The primary contribution of the present chapter is to show that partial cooperation can be achieved in situations where individuals interact only once. This cooperation is achieved through a revision process of players' actions before the "one-shot" game is played at a certain pre-determined time. However, the technology of revisions is vital in determining whether or not individuals are able to take actions that are more "cooperative" than others. While real time revision does not help in achieving cooperation, situations where players' revisions have an exogenous positive probability of being the final implemented actions (the Poisson revision games) are more conducive to cooperation.

The basic results of this study are the following. When individuals play the Cournot duopoly game with real time revision but without observing the rival firm's revisions then play converges to the Cournot-Nash outcome. Quantity choices are even more competitive than Cournot-Nash in the presence of a real time revision phase with perfect observability of a rival's revisions. In contrast, when revision plans are mutually observable and revision opportunities arrive according to a Poisson process, significantly lower quantities are selected than in the choices implemented under the real time revision games. The primary contribution of this study is to show that even though individuals interact only once in the sense that the payoffs are determined only by the one-time play of the game, a significant amount of cooperation could be sustained. The other contribution is the experimental design used to analyze the Poisson revision games. This study implements a dynamic game in the laboratory with a jump signal process which in this case is the Poisson process. This type of design is not standard in the experimental economics literature and could be used to study similar games.

The remainder of the chapter is organized as follows. Section 3.2 provides an overview of the related literature. The model, treatments, and hypotheses are presented in section 3.3. Section 3.4 lays out the laboratory methods and procedures. The results are discussed in Section 3.5. The last section concludes. Section C.1 of Appendix C reproduces the instructions to subjects and Section C.2 shows the profit sheets that were handed out to the participants.

3.2 Related Literature

There have been numerous experiments conducted on Cournot duopolies and the basic conclusion seems to be that while a random matching scheme results in play converging to the Cournot-Nash equilibrium, some degree of collusion often arises in repeated Cournot settings with a fixed pair of participants. See Fouraker and Siegel (1963), Holt (1985) and Huck, Muller and Normann (2001). These studies however generally do not allow communication between the players.

Balliet (2010) provides a meta-analysis of communication and cooperation in social dilemmas and concludes that there is a large positive effect of communication on cooperation in these situations. Duffy and Feltovich (2002) report that non-binding pre-play communication enhances efficiency in one-shot prisoner's dilemma games with random matching of participants. Non-binding pre-play communication has been shown to facilitate collusive play in spatial competition (Brown-Kruse et al (1993)), price competition in differentiated products framework (Friedman (1967)), and quantity competition in a Cournot duopoly framework (Daughety and Forsythe (1987a, 1987b), Waichman, Requate and Siang (2011)).

There are only a few studies that focus on the effects of real time revision and most of them are related to public goods provision. Dorsey (1992) studies the effects of allowing real time revision of voluntary contributions for the provision of a public good. He analyzes two rules, one where individuals can only increase their contribution over time (commitment) and another where they can both revise upwards and downwards. Cooperation increases significantly only for the case in which revisions are limited to increases and a provision point exists[1]. Similar results are reported in Goren, Kurzban and Rapoport (2003, 2004), Duffy, Ochs and Vesterlund (2007). Kurzban, McCabe, Smith and Wilson (2001) also implement a real time version of the voluntary contribution mechanism (VCM) and show that the commitment mechanism is effective in sustaining cooperation over time when players have access to complete information about others' contributions, but fails to increase public good provision in the case where individuals only observe the highest contribution. Again, an important result coming out of all these studies is that efficiency goes up only when a commitment mechanism is used, that is, pledges to contribution are irreversible.

Deck and Nikiforakis (2012) study the effect of real time revision on coordination and find that real time revision coupled with perfect observability of other players' actions at each moment in time almost always lead a group of individuals to coordinate at the payoff-dominant equilibrium in a minimum-effort game. However, this is no longer true when there is imperfect monitoring of others' actions. So, real time revision is effective only when monitoring is perfect.

While the real time revision games are implemented in the laboratory in a similar way as done in the previous studies, the current work uses a design innovation to study the Poisson revision games. This design is well suited for analyzing dynamic games with a jump signal process. A similar design has been implemented in Chapter 2 of this thesis in the context

[1]The linear voluntary contribution mechanism is analogous to a Prisoners' Dilemma with n players and m strategies. The provision-point mechanism is similar to the voluntary contribution mechanism with an added attribute that the public good is not provided unless a certain level of cooperation is reached.

of dynamic favor exchange games where participants could switch between action choices asynchronously in continuous time, but payoffs can be updated only at specific discrete times determined by a Poisson process. Similar to that design, in this current research, individuals can choose to update their desired choices in continuous time. However, the relevant choice for the game is governed by the arrival of a Poisson process.

3.3 Model, Treatments and Hypotheses

The experiment uses the Cournot duopoly model of quantity competition. There are two firms, firm 1 and firm 2, producing and selling a homogeneous product in a market. Each firm's decision is to choose an output level, $q_i \in [0, 50]$ in increments of 0.1. They face a linear inverse demand:

$$P(Q) = max\{50 - Q, 0\}, \quad Q = q_1 + q_2, \tag{3.1}$$

while the cost function for each firm is given by

$$C_i(q_i) = 2q_i, \quad i = 1, 2. \tag{3.2}$$

In this static model where firms decide simultaneously, the Nash equilibrium play implies $q_i^{CNE} = 16, i = 1, 2$ and the joint profit maximization implies an aggregate market quantity of $Q^{JPM} = 24$. On a symmetric Cournot market, the symmetric joint profit maximizing output is $q_i^{JPM} = 12, i = 1, 2$. The perfectly competitive Walrasian output is $q_i^{PCW} = 24, i = 1, 2$ where price equals the marginal cost. An overview of the relevant benchmarks concerning quantities, market price, market profits, consumers' surplus and total surplus is provided in Table 3.1. The parameters were chosen so as to give rise to a unique Cournot-Nash output where the profit for each of the firm is 256 points. However, if both choose 12, which gives rise to the joint profit maximizing output, the individual profits are 288. Given that one firm is producing 12, the best-response of the other firm is 18 which gives rise to a profit of 324 for the one producing 18 and only 216 for the one choosing 12. Thus, the parameters give a large enough increase in profits from "defection" as a percentage of total profits.

	Cournot-Nash	Joint Profit Maximization	Perfect Competition
Individual Quantity	16	12	24
Market Quantity	32	24	48
Market Price	18	26	2
Market Profits	512	576	0
Consumers' Surplus	512	288	1152
Total Surplus	1024	864	1152

Table 3.1: Static benchmarks

Behavior in this basic Cournot setting is investigated in three treatments that differ in the technology used to implement the revisions of quantity choices by the firms. These different revision technologies are listed below.

3.3.1 The Cournot Game with Real Time Revision and No Monitoring (Baseline)

In this baseline treatment, the Cournot game is played at time T. However, firms could adjust their quantity choices in real time from $t = 0$ to $t = T$. The output choice that is actually implemented is the one that corresponds to time $t = T$, but, the firms are unable to monitor their competitor's choices and adjustments over time. Only at the very end does a firm get to observe the implemented choice of the other firm. Thus, this scenario is strategically equivalent to the static simultaneous move Cournot game. The static game is implemented in this manner instead of the usual one-shot version to make every aspect comparable to the other treatments, except for the observability of competitor's choices. In this game, similar to the one-shot game, there is an unique Nash equilibrium of $q_i^* = 16, i = 1, 2$ at time T. This gives rise to the first hypothesis as follows.

Hypothesis 1 (Baseline Hypothesis). *Individual quantities are equal to the Cournot-Nash output of 16 in the baseline treatment.*

3.3.2 The Cournot Game with Real Time Revision and Perfect Monitoring (Real Time)

In this treatment, firms can not only adjust their quantity choices in real time until time reaches T but also observe the output revisions of their competitor in real time. Again the output choices corresponding to time $t = T$ are implemented. Since the choices at any time $t < T$ are not binding, the entire phase $[0, T)$ constitutes a "cheap talk" period where the language of communication is the intended quantity that a firm wishes to produce at time T. There is still a unique Nash equilibrium prediction of $q_i^* = 16, i = 1, 2$ at time T regardless of the dynamics of revisions during $t \in [0, T)$. This is because at time T there is no future and the unique mutual best response is at the quantity level of 16. This leads to the following hypothesis.

Hypothesis 2 (Real Time Revision Hypothesis). *Individual quantities are the same under the real time revision game with monitoring and without monitoring (baseline).*

3.3.3 The Cournot Game with Poisson Revision and Perfect Monitoring (Poisson)

Firms submit their intended quantity choices "continuously" in this treatment and a synchronous stochastic process determines which quantity choices are implemented at time T. Specifically, there is a Poisson process with an arrival rate $\lambda > 0$ defined over the time interval $[0, T)$ and the intended choices corresponding to the last time the process takes place are implemented at time T. Both the firms observe all the past events including the "continuous" revision plans.

The revision phase $[0, T)$ is now different from a "cheap talk" period because revisions are now binding with some exogenous probability and the dynamics of revisions play a role now in this treatment. Kamada and Kandori (2011) analyze these games and call them *Poisson revision games* and identify symmetric revision game equilibria that use trigger strategies. Each of these equilibria prescribes an action path $q(t)$ which implies that when a revision opportunity arrives at time t, players are supposed to choose the quantity $q(t)$, given that there have been no deviations in the past. If any player deviates and does not choose $q(t)$, then in the future players choose the Nash equilibrium quantity of the Cournot game, whenever a revision opportunity arrives. They identify the optimal symmetric trigger strategy equilibrium that achieves (ex ante) the highest payoffs in this class of equilibrium. Certain assumptions need to be fulfilled in order to apply this equilibrium concept:

- First, a pure symmetric Nash equilibrium must exist and be different from the best symmetric action profile.

- Second, the payoff function for each player is assumed to be twice continuously differentiable.

- Third, there is an unique best response for each possible action.
- Fourth, at the best reply, the first and second order conditions are satisfied.
- Fifth, the payoff function is strictly increasing (decreasing) if an action is lower (higher) than the best symmetric action.
- Sixth, the gains from deviation are strictly decreasing (increasing) if an action is lower (higher) than the unique Nash action.

The Cournot duopoly game we study satisfies all these assumptions[2] and hence, an optimal trigger strategy equilibrium exists and is essentially unique.

Kamada and Kandori characterize paths $q_{T-\tau}(t)$, $\tau \in (-\infty, T]$ that are continuous in t and lie in $[12, 16]$, i.e., it never prescribes a quantity that is lower than the individual collusive outcome of $q^m = \frac{24}{2} = 12$ and higher than the Nash equilibrium output of 16. For $\tau \leq 0$, the path always prescribes the play of the Cournot-Nash output of 16 throughout the revision phase. Provided $\tau > 0$, $q_{T-\tau}(t)$ starts at q^m and remains there till $t_\lambda(q^m) - (T - \tau)$, where $t_\lambda(q^m)$ is the first time the prescribed path deviates from the collusive output of 12 when $\tau = T$. It follows a differential equation for $t \in (t_\lambda(q^m) - (T - \tau), \tau)$ and hits the Cournot-Nash output of 16 at time $t = \tau$. Among these paths $q_{T-\tau}(t)$, the one that reaches the Cournot Nash output level only at the end, i.e., which prescribes the play of 16 at time $t = T$ and $q(t) \neq 16, t \neq T$ has the highest payoff as it stays at q^m for a longer period of time. This is achieved for $\tau = T$ and is the *optimal trigger strategy equilibrium.*

The optimal trigger strategy equilibrium path can be calculated as follows. First, the expected payoff from cooperation[3] at time t, supposing that the other player follows the

[2]Kamada and Kandori assume continuous action spaces. In the experiments, subjects could choose any quantity in $[0, 50]$ with increments of 0.1, which is "almost" continuous.

[3]Cooperation is defined as following the action path prescribed by the trigger strategy equilibrium.

trigger strategy equilibrium, is given by:

$$V_C(t) = (48 - 2q(t))q(t)e^{-\lambda(T-t)} + \int_t^T (48 - 2q(s))q(s)\lambda e^{-\lambda(T-s)}\, ds \quad .$$

The probability that the current quantity $q(t)$ will be implemented at the end of the game is $e^{-\lambda(T-t)}$. The first term represents the payoff when there is no opportunity of revision in the future. The second term can be explained as follows. With probability density λ a revision opportunity arrives at time s, and with probability $e^{-\lambda(T-s)}$ this is the last revision opportunity. If that happens then $q(s)$ will be implemented and the realized payoff equals $(48 - 2q(s))q(s)$.

The expected payoff from deviation is maximized by deviating to the best response to the competitor's current quantity, which is $24 - \frac{q(t)}{2}$. Also, once there is a deviation, then both players choose the Nash equilibrium output. Thus, the expected payoff from deviation is given by:

$$\begin{aligned} V_D(t) &= (48 - (24 + \frac{q(t)}{2})(24 - \frac{q(t)}{2})e^{-\lambda(T-t)} + (48 - 2\frac{48}{3})(16)(1 - e^{-\lambda(T-t)}) \\ &= (24 - \frac{q(t)}{2})^2 e^{-\lambda(T-t)} + 256(1 - e^{-\lambda(T-t)}) \quad . \end{aligned}$$

Note that $V_C(T) = V_D(T) \quad \text{if} \quad q(T) = 16$, and the sufficient condition for sustaining cooperation is:

$$V_C'(t) \geq V_D'(t) \quad \forall \quad t \quad .$$

Solving the above inequality and using the fact that at the optimal equilibrium equality holds in the above expression, the following differential equation is obtained:

$$\frac{dq}{dt} = \frac{\lambda}{18}(q - 80) \quad .$$

Finally, solving the above differential equation, the optimal trigger strategy equilibrium is characterized by the following[4]:

$$\begin{aligned} q^m = 12 \quad &\text{if} \quad t \in [0, \frac{18}{\lambda} ln(\frac{17}{16})] \\ 16(5 - 4e^{\frac{\lambda}{18}(T-t)}) \quad &\text{if} \quad t \in (\frac{18}{\lambda} ln(\frac{17}{16}), T] \quad . \end{aligned}$$

Given a long enough horizon for the revision phase, players act as follows in the optimal symmetric trigger strategy equilibrium. They start with the best symmetric action which is the individual collusive output of 12. They do not revise their quantity plans (even if revision opportunities arrive) till time $t_\lambda(q^m)$. After this time, they choose the optimal path given by $q(t)$ whenever a revision opportunity arrives. The closer the revision opportunity is to the end of the revision phase, the closer the revised quantity is to the Nash equilibrium output of 16. When the revision phase is over and the game ends, the quantities chosen at the last revision opportunity are implemented.

At worst, a trigger strategy equilibrium may specify the play of Cournot-Nash throughout the revision phase. For values of $\tau \in (0, T]$, the path stays in the collusive output level for a finite period of time and also takes values in $(12, 16)$. Thus, these equilibria predict lower levels of output to be chosen (in expectation) than the Cournot-Nash equilibrium and hence, the following hypothesis is obtained.

Hypothesis 3 (Poisson Revision Hypothesis). *Individual quantities are lower and*

[4]For further details see Kamada and Kandori (2011).

the extent of collusion is higher under the game with a Poisson revision phase than with a real time revision phase without monitoring (baseline).

Similar to the case of repeated or dynamic games, the multiplicity of equilibria here is severe, even when restricted to the class of trigger strategy equilibrium. While it would be very restrictive to test whether individuals follow the optimal trigger strategy equilibrium, a more weaker but reasonable question to ask is the following: Do players begin with $q(0) < q^*$ and then revise $q(t)$ upwards (weakly) over the revision phase $[0, T]$? Individuals start out at lower quantities and move towards the static Nash output level over the course of the revision phase under these symmetric trigger strategy equilibria. This would lead to the following hypothesis about $q(t)$.

Hypothesis 4 (Dynamics Hypothesis). *(i)* $q(0) < q^*$, *(ii)* $t < t^{'} \implies q(t) < q(t^{'})$. In other words, initial tentative quantities are lower than the Cournot-Nash quantity and players revise their quantities upwards over time during a revision phase.

The optimal symmetric trigger strategy equilibrium induces a probability distribution of quantities over the individual collusive output of 12, the Nash output of 16 and quantities in $(12, 16)$. An interesting feature of this equilibrium is that the probability distribution of the quantities implemented at time T is independent of the Poisson arrival rate λ provided the time horizon is long enough. This follows from the fact that any revision game with an arrival rate of λ and a time horizon T can be rewritten as a new model by changing the time scale in such a way that one unit of time in that game would correspond to λ units in the new model. Under the new time scale, the model is identical to the revision game with arrival rate one and time horizon λT, and, the first time the optimal path starts deviating from the collusive output, $t_1(q^m)$ equals $\lambda t_\lambda(q^m) \leq \lambda T$ given that $t_\lambda(q^m) \leq T$ in the new model. Thus, the probability distribution of quantities implemented at the end of the revision phase is unchanged if the game starts at $t_1(q^m)$ and ends at λT. Hence, the probability distribution of quantities at time $t = T$ under any arrival rate λ such that $t_\lambda(q^m) \leq T$ is equal to the distribution under arrival rate one and time horizon $(\lambda T - t_1(q^m))$.

Two treatments are generated by varying the arrival rate of the Poisson process. The optimal trigger strategy equilibrium achieves the same (ex ante) payoffs under high and low arrival rates. Thus, the best individuals can do under these classes of equilibria is the same and this leads to the last hypothesis, as follows.

Hypothesis 5 (Arrival Rate Neutrality Hypothesis). *Individual quantities are identical and extent of collusion is same under the games with Poisson revision technology with low and high arrival rates.*

3.4 Methods and Procedures

The experiments reported here were all conducted at the Social Science Experimental Laboratory (SSEL), California Institute of Technology (Caltech) using the Multistage software package[5]. Subjects were recruited from a pool of volunteer subjects, maintained by the SSEL. A total of eight sessions were run, using a total of 96 subjects. No subject participated in more than one session. A total of four treatments were generated, each with a revision phase of 120 seconds. As the previous section listed, there were two treatments with real time revision of quantities, one without monitoring (baseline) and the other with monitoring. There were two treatments with the Poisson revision game, one with a low arrival rate (0.02 per second) and the other with a high arrival rate of revision opportunities (0.04 per second). Table 3.2 summarizes the characteristics of each session.

[5]Please visit http://software.ssel.caltech.edu/ for more details on the multistage software project.

Session	Subjects	Treatment	Matches	Games
1	12	Real Time with Monitoring	11	66
2	12	Real Time with Monitoring	11	66
3	12	Baseline	11	66
4	12	Baseline	11	66
5	12	Poisson-High	11	66
6	12	Poisson-Low	11	66
7	12	Poisson-High	11	66
8	12	Poisson-Low	11	66

Exchange Rate : US$ 1 was worth 150 points earned in the experiment.
Subjects (n) : Number of subjects in a session.
Matches (m) : Number of dynamic games played per subject.
Games : Total number of dynamic games played (summing across all subjects). As a game is played among two persons, there are $\frac{m \times n}{2}$ dyadic games in total.

Table 3.2: Treatments and sessions

On arrival, instructions[6] were read aloud. Subjects interacted anonymously with each other through computer terminals. There was no possibility of any kind of communication between the subjects, except as explained in the instructions. All sessions consisted of eleven matches and lasted between 45 and 50 minutes. Average earnings were US$19. Each subject also received an additional US$5 show-up fee.

In the instructions subjects were told that they were to act as a seller which, together with another seller, produces the same product in a market and that, in each match, both have to decide what quantity to produce. They were also informed that every match they would be matched to a new participant from the room and thus, they would not be matched with the same person ever again.

Participants were provided with a profit sheet and a profit calculator. The profit sheet showed the profits a seller would earn for every possible combination of integer choices (from 0 to 30) by her and the other seller. For integer numbers above 30 or non-integer numbers,

[6]Verbatim copy of the instructions is given in Appendix C.

they had to use the profit calculator provided on their screens. Each seller could choose an output level from the set $\{0, 0.1, 0.2, ..., 49.8, 49.9, 50\}$ and the payoffs were generated according to the demand and cost functions given in equations (3.1) and (3.2). The payoffs were measured in a fictitious currency unit called points and subjects were told that at the end of the experiment they would be paid the US Dollars equivalent of the sum of points earned by them across all matches. This USD equivalent was determined by using an exchange rate from 150:1.

Subjects were informed that every match they would have 120 seconds to decide the quantity to be produced. At the beginning of this "120 seconds" period they were asked to enter an initial tentative quantity from $\{0, 0.1, 0.2, ..., 49.8, 49.9, 50\}$. Once all the participants in the room made an initial choice, the match started in real time. In the baseline treatment, a subject only observed her own tentative quantity, as shown in Figure 3.1. Notice that whenever there was a change in the tentative quantity made by the subject, the screen was updated along with the time at which the change took place. At the end of the revision phase each subject observed not only her tentative quantity but also the quantity of the other seller as well as the profits earned for that match.

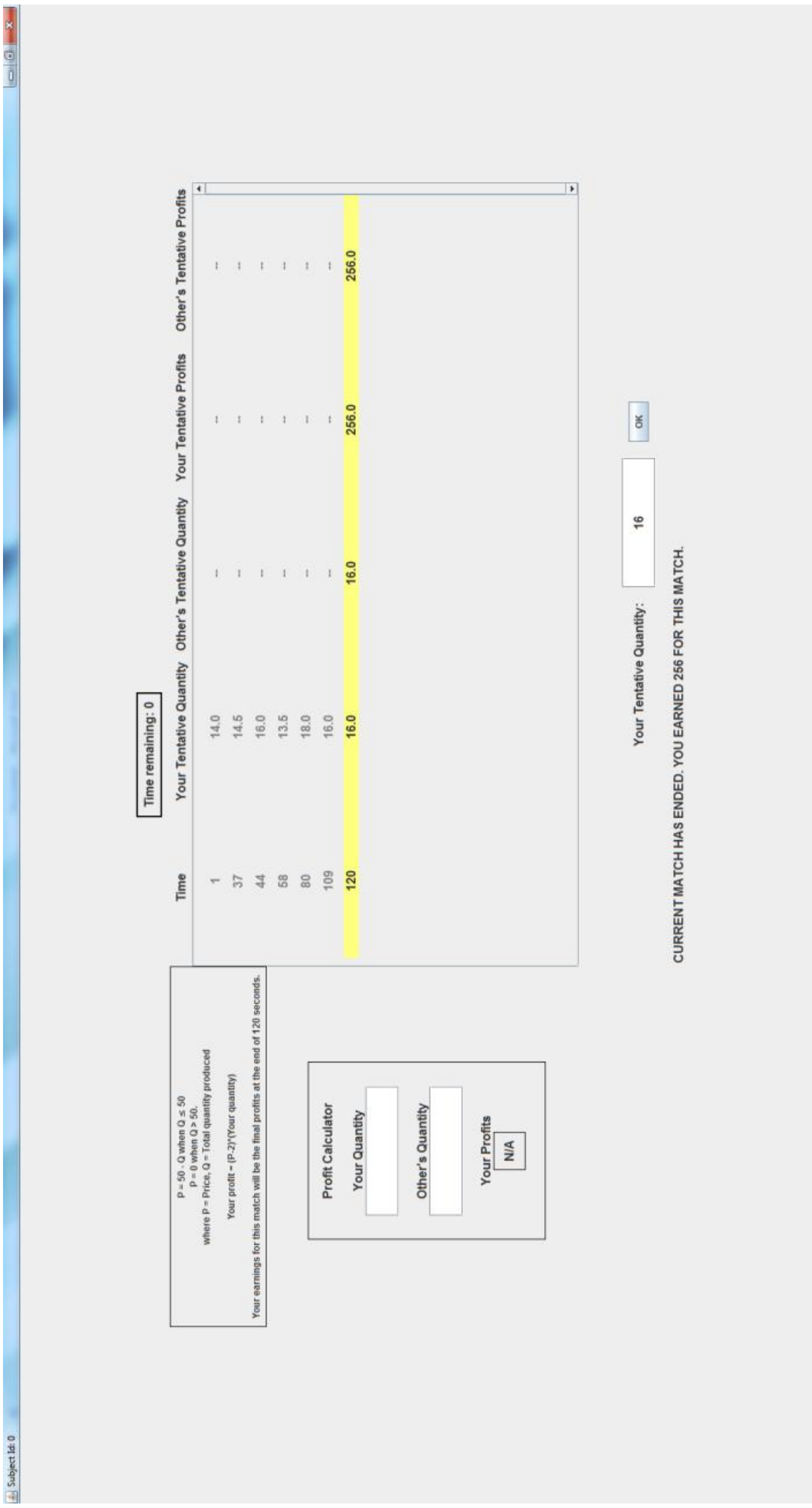

Figure 3.1: Subject screen for “Baseline”

Figure 3.2 displays the interface for the treatments with real time revision with monitoring. Here a participant gets to see her own tentative quantity, the tentative quantity chosen by her competitor and the tentative profits. The screen was updated whenever either of the two matched participants decided to change their respective quantities. Subjects were also told that the word "tentative" reflects the fact that if there were no more changes then the last entry shows the final quantity choice and the amount of profits earned by them for that match.

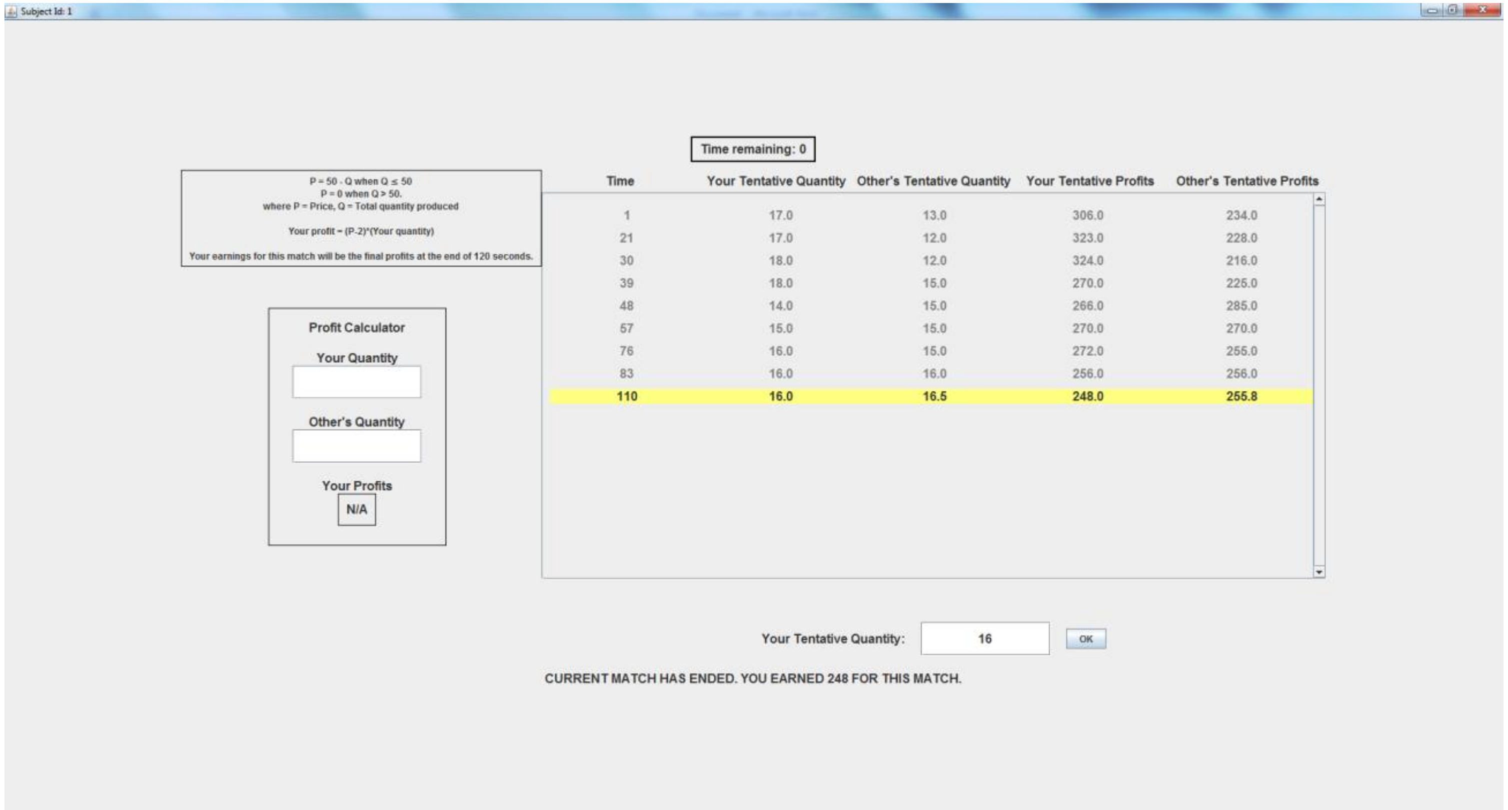

Figure 3.2: Subject screen for “Real Time Revision with Monitoring”

The screen display for the treatments with a Poisson Revision phase is given in Figure 3.3. This time the screen was updated for two separate reasons. First, there was an update if any of the two matched participants changed their respective tentative quantities. This was shown as an "unstarred" entry on the screen. The second type of update was shown with a "starred" entry and it marked the occurrence of an "event" that was responsible for the final quantity choices in a match. Specifically, the final choices and profits for a match were given by the tentative quantities and profits corresponding to the last "starred" entry on the screen. For each pair of participants, the occurrence of this event was decided using a random draw of an integer from an uniform distribution over $[1, 50]$[7]. If the draw was one, then an event occurred. Thus, the inter-arrival time between the occurrence of this event follows a geometric distribution which is an approximation to the exponential distribution in continuous time.

[7]This was for the Poisson-Low treatment. For the Poisson-High treatment, an integer was drawn from an uniform distribution over $[1, 25]$.

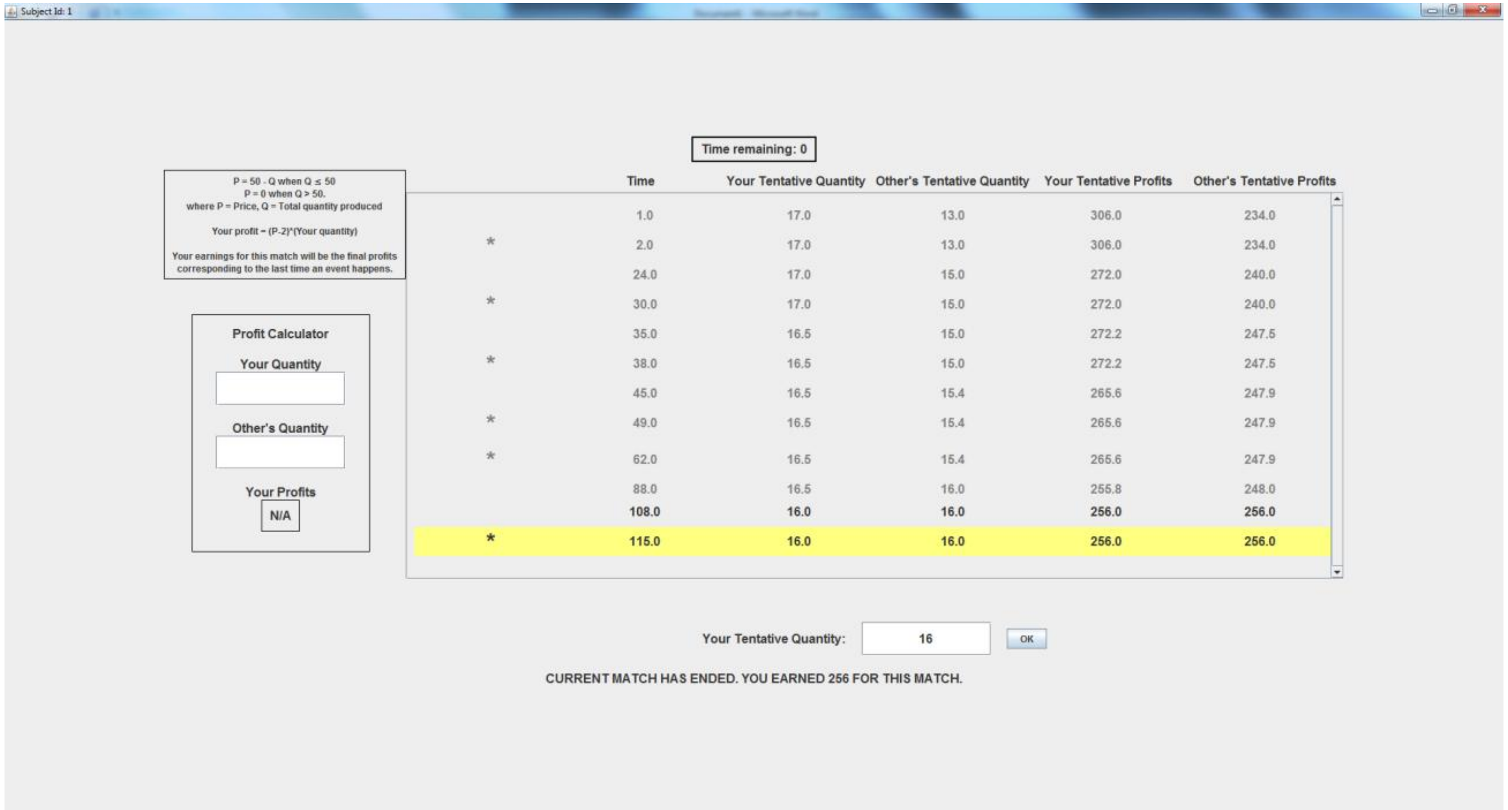

Figure 3.3: Subject screen for "Poisson Revision"

3.5 Results

The results are reported in this section. Subsection 3.5.1 collects the results from comparison of the selected quantities as well as the extent of collusion in the different treatments. The aggregate (average) dynamics and market level behavior over time are discussed in the next subsection 3.5.2. The final subsection 3.5.3 discusses in greater detail the characteristics of quantity adjustment over time.

3.5.1 Comparison of Quantities and Extent of Collusion

The essential summary statistics at an aggregate level are provided in Table 3.3 for all the four treatments. More detailed information is given in Table 3.4. There is an interesting pattern in the quantities selected over matches in the four treatments, as shown in Figure 3.4. There is an upward trend in the quantities selected over matches in the real time revision treatments. While average individual quantity is below 15 at the start of the experiment, it reaches 16 during the later stages in the baseline treatment. In the real time revision treatments with monitoring, the average quantity is close to 15 in match 1 but higher than 17 in matches 9-11. In contrast, higher quantities are chosen during the initial matches whereas there is a considerable decline over the course of the experiment in the Poisson revision treatments.

	Cournot-Nash	Baseline	Real Time	Poisson High	Poisson Low
Individual Quantity	16	15.7 (15.9)	16.4 (17.0)	15.4 (14.6)	15.5 (15.1)
Market Quantity	32	31.4 (31.9)	32.9 (33.9)	30.7 (29.0)	30.9 (30.3)
Market Price	18	18.6 (18.1)	17.2 (16.1)	19.4 (21.0)	19.4 (20.1)
Market Profits	512	514.2 (510.6)	476.0 (458.4)	485.1 (525.7)	498.6 (512.9)
Consumer Surplus	512	495.8 (509.6)	550.4 (584.9)	491.2 (432.3)	482.5 (456.8)
Total Surplus	1024	1010.1 (1020.2)	1026.4 (1043.2)	976.4 (957.9)	981.1 (969.7)

Table 3.3: Aggregate data (averages). Data from matches 7-10 in parentheses.

It is usual in Cournot duopoly experiments (Fouraker and Siegel (1963), Holt (1985) and Huck, Muller and Normann (2001)) to work with the data from the latter half of the experiment when subjects have acquired experience. Also, the last match is not considered in the data analysis due to possible end game effects. Unless otherwise mentioned, only the data for matches 7 through 10 is used in the analysis in this study.

The average quantities selected in the baseline treatment converge to 16, which is the Cournot-Nash equilibrium. This is significant with p-values > 0.1 from a two-sided test of quantities being equal to 16. The median and mode are exactly equal to 16 in matches 7-10. This observation is familiar and is consistent with the other experimental studies which report that under random matching, play converges to the Cournot Nash equilibrium. This supports the *Baseline Hypothesis* gives us the first result.

Result 1. *Average quantities are not significantly different from the Cournot-Nash output of 16 in the baseline treatment (p-value > 0.1, support for H1).*

In contrast, markets are significantly more competitive under real time revision with monitoring. In this treatment, the mean quantity chosen is close to 17 in the matches 7-10 and the modal quantity choice is 18 across all matches. A comparison of the average market quantities[8] shows that there is a highly significant difference between the two treatments (real time with and without monitoring), with p-value < 0.01 from a two-sided Mann-Whitney-U test. Thus, rejecting the *Real Time Revision Hypothesis* we have the following result.

Result 2. *The real time revision treatment with monitoring results in significantly (p-value < 0.01) higher individual and market quantities as well as lower profits than the baseline treatment (reject H2).*

[8]Using the average quantities of each market as the unit of observation is preferred to the other option of using the average of individual quantities as the former gives independent observations on each unit.

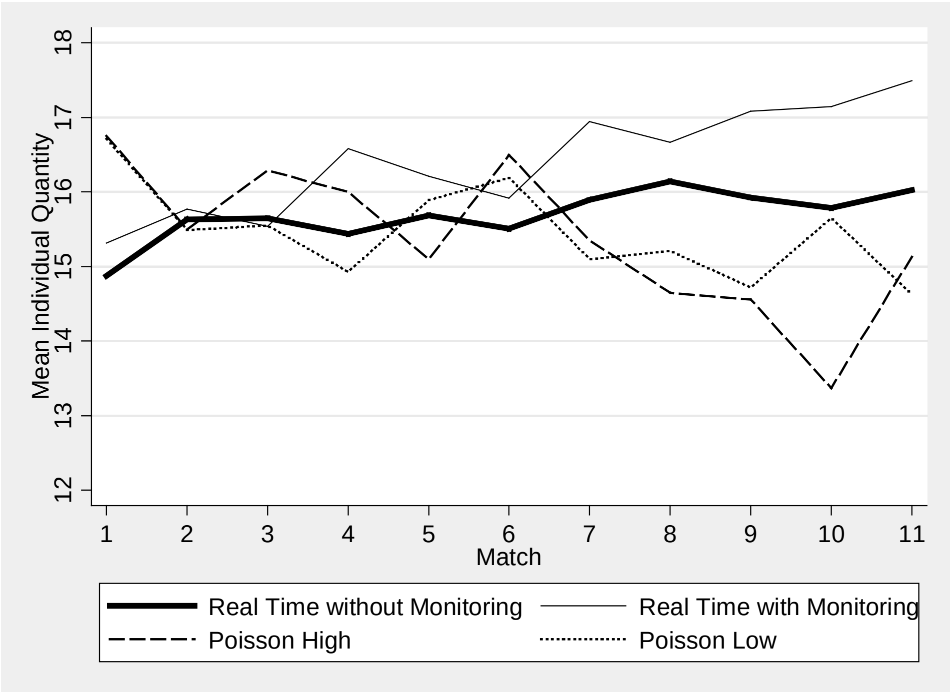

Figure 3.4: Mean individual quantities implemented over matches

	Baseline		Real Time		Poisson High		Poisson Low	
Match	Mean (SE)	Med (Mo)	Mean (SE)	Med (Mo)	Mean (SE)	Med (Mo)	Mean (SE)	Med (Mo)
1	14.9 (0.6)	15 (12)	15.3 (0.8)	15 (12)	16.8 (0.9)	16 (15)	16.7 (0.9)	15 (15)
2	15.6 (0.5)	16 (15)	15.8 (1.0)	16 (18)	15.5 (0.9)	15 (16)	15.5 (0.4)	16 (16)
3	15.6 (0.5)	16 (16)	15.5 (0.7)	15 (12)	16.3 (1.0)	15.5 (12)	15.6 (0.7)	15.5 (12)
4	15.4 (0.3)	16 (16)	16.6 (0.7)	18 (18)	16.0 (1.0)	15 (16)	14.9 (0.6)	15.5 (16)
5	15.7 (0.3)	16 (16)	16.2 (0.6)	16.5 (18)	15.1 (0.8)	14.8 (12)	15.9 (0.7)	16 (12)
6	15.5 (0.4)	15.8 (15)	15.9 (0.6)	16 (18)	16.5 (1.2)	15.3 (12)	16.2 (1.5)	15 (12)
7	15.9 (0.3)	16 (16)	16.9 (0.6)	17.8 (18)	15.4 (0.5)	16 (12)	15.1 (0.5)	15 (12)
8	16.1 (0.2)	16 (16)	16.7 (0.7)	17 (18)	14.7 (0.5)	13.5 (16)	15.2 (0.5)	16 (16)
9	15.9 (0.3)	16 (16)	17.1 (0.5)	17.5 (18)	14.6 (0.6)	13.5 (12)	14.7 (0.7)	13.3 (12)
10	15.8 (0.3)	16 (16)	17.1 (0.6)	18 (18)	13.4 (0.4)	12 (12)	15.6 (1.6)	13.5 (12)
11	16.0 (0.2)	16 (15)	17.5 (0.5)	18 (18)	15.1 (1.0)	12.8 (12)	14.6 (1.2)	12.5 (12)
1-6	15.5 (0.4)	16 (16)	15.9 (0.6)	16 (18)	16.0 (0.7)	15 (16)	15.8 (0.5)	15.5 (16)
7-10	15.9 (0.2)	16 (16)	17.0 (0.5)	17.8 (18)	14.5 (0.4)	13 (12)	15.2 (0.5)	15 (12)

Table 3.4: Mean (standard errors clustered at the individual level) and median (mode) of individual quantity implemented by matches in each treatment. Each match has 24 observations.

Now comparing the market quantities between the Poisson-High and the baseline as well as between the Poisson-Low and the baseline treatment reveal that the Poisson revision treatments result in significantly lower aggregate quantities (p-value < 0.01, Mann-Whitney-U tests). The modal quantity chosen in the majority of the matches under the Poisson revision treatment is 12. So, in contrast to the real time revision phase with monitoring, the introduction of a Poisson revision phase results in a higher incidence of collusion. However, it should also be noted that the quantities implemented in both Poisson-High and Poisson-Low treatments are higher than the expected output choice under the optimal symmetric trigger strategy equilibrium, which is equal to 14.2. Now, in matches 7-10, the collusive market outcome of 24 is chosen in 31% of the markets in Poisson revision treatments and in 2 out of 48 markets under the real time revision with monitoring treatments. On the other hand, none of the pairs achieve the fully collusive outcome in the baseline treatment. Hence, supporting the *Poisson Revision Hypothesis* the following result is obtained.

Result 3. *Individual and market quantities are significantly lower (p-values $<$ 0.01) under the Poisson revision games than the baseline treatment (support for H3).*

Finally, the market quantities in the two Poisson treatments are not statistically different from each other (p-value > 0.1, Mann-Whitney-U test). So, on an average, the arrival rate does not make a difference in the quantities chosen. The modal quantity implemented in both Poisson-High and Poisson-Low treatments is 12 for matches 7-10. Thus, supporting the *Arrival Rate Neutrality Hypothesis* we have the following result.

Result 4. *Quantities chosen under the Poisson-High and Poisson-Low treatments are not significantly different from each other (p-value $>$ 0.1, support for H5).*

Table 3.3 also shows the differences in other relevant aspects of a market, including the market price, market profits, consumer surplus and the total surplus. These follow directly from the above results. While prices are higher in the Poisson revision treatments compared to the real time revision ones, market profits in the Poisson revision treatments are only marginally higher than the baseline treatment. The average market profits, however, are considerably lower in the real time revision treatment compared to the other three treatments. The markets under the baseline treatment attain almost 89% of the collusive profits in matches 7-10, similar to those which the Cournot-Nash quantities attain. However, only 79% of the collusive profits are realized in the real time revision treatments with monitoring. On the other hand, the markets in the Poisson revision treatments attain 89-91% of the collusive profits. Consumers are significantly better off under the real time revision with monitoring treatment, followed by the baseline one. The consumer surplus is lower in the Poisson revision treatments. Finally, total surplus is highest in the real time revision treatment. Poisson revision treatments have the lowest total surplus.

3.5.2 Dynamics

3.5.2.1 Aggregate Dynamics

The distribution for the initial tentative quantity selected and the final quantity implemented for a match is given in Tables 3.5 and 3.6. Again we focus on matches 7-10. One can immediately infer that the average quantity in the baseline treatment is fairly stable over the entire revision phase. Almost 70% of the individual quantities chosen are in the range [16, 18) right from the beginning. This is also displayed in panel (a) of Figure 3.5. The p-value from a pairwise Wilcoxon rank sum test is > 0.1 showing no difference between $q(0)$ and $q(T)$. In contrast, the p-value for the same test in the real time revision treatment with monitoring is less than 0.001. In fact, around 64% of the observations for $q(0)$ are near the individual collusive quantity of 12, whereas, only 12% of the final quantities are

	Baseline		Real Time	
Qty.	Initial	Final	Initial	Final
$[0, 12)$	1	0	1	1
$[12, 13)$	5	4	64	11
$[13, 14)$	0	2	0	1
$[14, 15)$	5	5	1	2
$[15, 16)$	16	18	5	8
$[16, 17)$	43	43	5	21
$[17, 18)$	26	26	4	5
$[18, 19)$	4	2	2	38
$[19, 20)$	0	0	0	1
$[20, 30)$	0	0	18	12
$[30, 40)$	0	0	0	0
$[40, 50]$	0	0	0	0

Table 3.5: Initial and final (implemented) quantity distribution (in percentages): "Baseline" and "Real Time Revision with Monitoring". Data from matches 7-10.

close to 12. More than 60% of the final observations are in the interval [16, 19). This happens primarily because of the 'last second' increase in the tentative quantities, as shown in Figure 3.5(b). While in one session the average output choice throughout the revision phase is consistently above the Cournot-Nash quantity, it is below 16 for most of the revision phases in the other session. However, towards the end of the revision phase, there is a sharp increase in the desired output choice in both sessions making the final average choice higher than the Cournot-Nash quantity.

In the Poisson revision games the typical pattern of $q(t)$ is for $q(0)$ to start high and then decline over time. In the final seconds of the revision phase, $q(t)$ again increases towards the Cournot-Nash output. This is documented in panels (c) and (d) of Figure 3.5. In fact, in three out of four sessions, $q(0)$ reaches exactly 16. The optimal trigger strategy equilibrium also prescribes that at the end when there is no future interaction left, the tentative choice should equal the Cournot-Nash quantity. If one discards the initial dynamics of desired choices, that is, the initial downward quantity adjustments, then the

	Poisson-High			Poisson-Low		
Qty.	Initial	Final	Expected Final	Initial	Final	Expected Final
[0, 12)	3	0	0	0	0	0
[12, 13)	52	46	38	18	35	46
[13, 14)	1	5	10	3	2	15
[14, 15)	1	3	29	13	11	10
[15, 16)	5	6	21	10	14	23
[16, 17)	1	19	2	9	15	6
[17, 18)	0	3	0	7	6	0
[18, 19)	7	12	0	20	12	0
[19, 20)	6	1	0	9	0	0
[20, 30)	19	5	0	8	4	0
[30, 40)	4	0	0	0	0	0
[40, 50]	1	0	0	3	1	0

Table 3.6: Initial and implemented quantity distribution (in percentages): "Poisson-High" and "Poisson-Low". Also listed are the theoretical (expected) distribution of final quantities implemented from the play of the optimal trigger strategy. Data from matches 7-10.

rest of the dynamics have a stark resemblance, at least qualitatively, to the trigger strategy equilibrium. However, the entire dynamics follows a flat bowl shaped curve.

Comparing the initial tentative quantities and the final implemented quantities in Poisson revision games, Table 3.6 shows that around 40% of the initial observations are above 18 whereas only 17% of the implemented quantities are above 18. And, around 46% (35%) of the observations are near the collusive output of 12 for the final implemented quantities in the Poisson-High (Poisson-Low) treatment. The table also shows the distribution of quantities to be implemented if the players follow the symmetric optimal trigger strategy, using the realized draws from the experiment. A stark similarity between the theoretical and the observed distribution is the implementation of quantities near to the Collusive output of 12. The quantity range of [12, 13) has the highest percentage of observations in both data and theory. However, the optimal trigger strategy equilibrium over-predicts the implementation

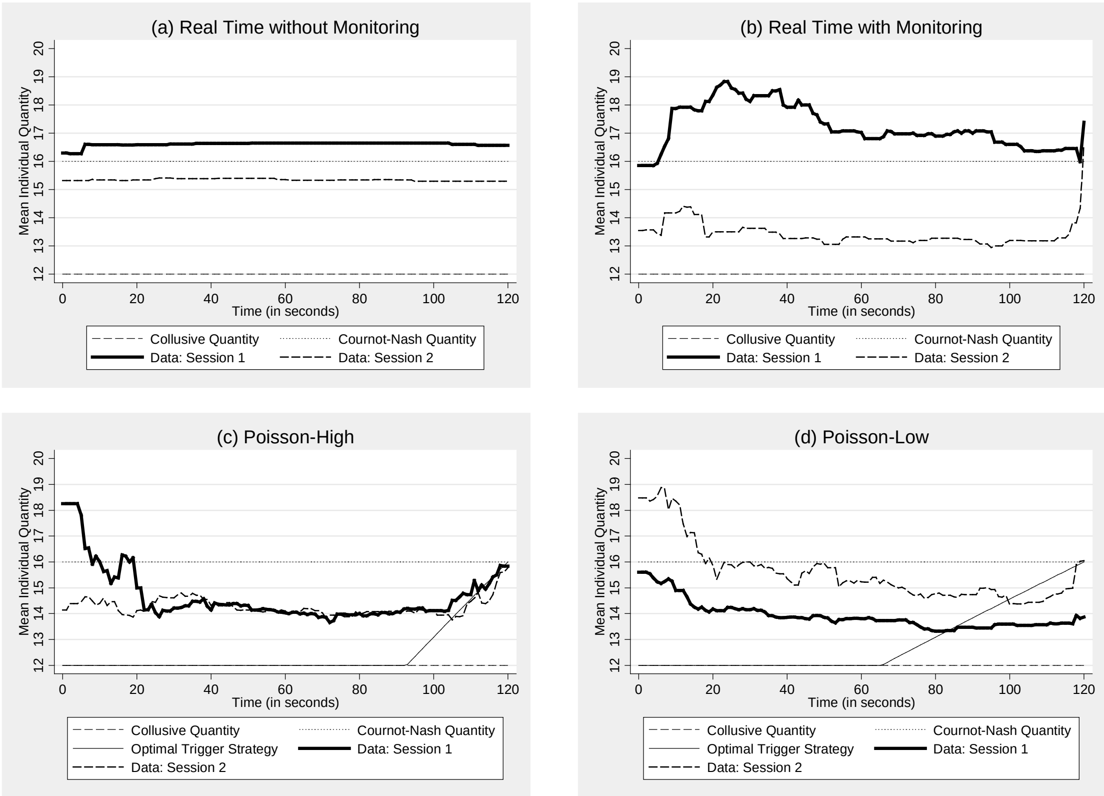

Figure 3.5: Mean individual quantity over time

of quantities in the range [13, 16) while it under-predicts implementation of quantities near the Cournot-Nash quantity.

Thus, summarizing the overall dynamics in the Poisson revision treatments, we reject the *Dynamics Hypothesis* and have the following result.

Result 5. *There is no significant upward trend in the quantity revisions over time in the Poisson revision treatments (reject H4).*

3.5.2.2 Market Level Dynamics

$Q(t)$ selected by a pair of participants is very stable over time in the baseline treatment. This is quite obvious as there is no information feedback and hence there is no opportunity to react to the competitor's plans. However, interesting dynamic patterns are exhibited in the real time revision treatments with monitoring. Figure 3.6 collects some of these patterns. The most common behavior is displayed in panels (a) and (b) which show that both the firms select the collusive quantity of 12 for a long time but as the revision phase comes to an end either one or both of them switch to a higher quantity. This 'last second' defection towards the best response of 18 explains the sharp increase in average tentative quantity towards the end of the revision phase, as in Figure 3.6(b)[9]. Another type of behavior is documented in panel (c) of Figure 3.6. Both firms start at the collusive quantity then one of them switches to a higher quantity followed by the other firm. Revisions in quantities are positively correlated and the 'last second' defection is also present.

Situations where one firm starts out with a high quantity and the other with a low quantity are shown in panels (d) and (e). While panel (d) shows that over time firms' quantities are revised downwards and are positively correlated, panel (e) shows that only one of the firms (which started out low) is trying to signal the other firm that it is willing to lower quantities provided its partner does so. Ultimately, they end up producing somewhere in between the quantities they initially intended to produce.

Finally, Stackelberg leadership emerges endogenously in a few cases. The quantity for the leader is 24 and the follower is 12. As the panels (h) and (i) show, one of the firms always selects a tentative quantity of 24 throughout the revision phase and its competitor finally 'gives in' and selects a quantity of 12 or close to 12. However, when none of the firms 'give in', as in panel (g), then both firms produce very high quantities and earn very reduced profits. Yet another situation is depicted in panel (f) wherein both the firms keep on selecting the high quantity of 25 (near the Stackelberg leader quantity) but only when

[9] 18 is a best response to the competitor's choice of 12.

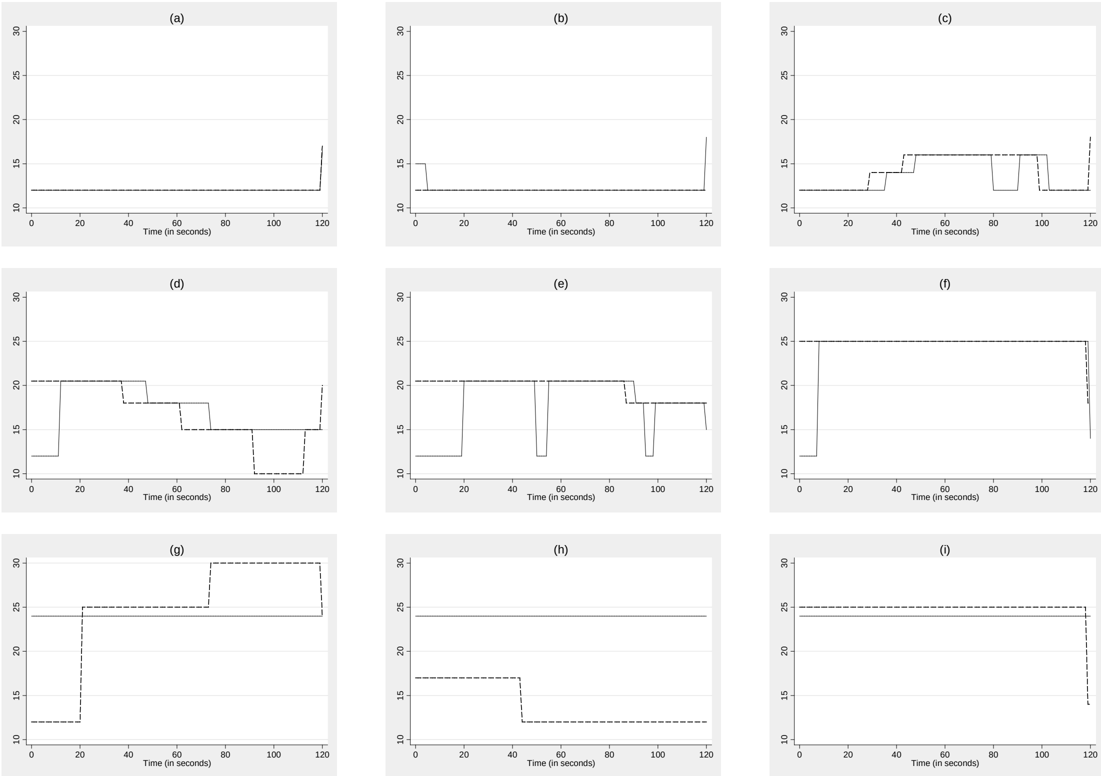

Figure 3.6: Patterns of behavior in a market in "Real Time Revision with Monitoring". The y-axis gives the tentative quantities selected by the firms in a market over a revision phase (labels are from 10 to 30 in increments of 5) and the x-axis is time in seconds (labels are from 0 to 120 in increments of 20).

just a few seconds remain for the revision phase to end do both of them switch to lower quantities.

The dynamics is remarkably different in the Poisson revision treatments. Figure 3.7 displays some typical patterns of play in these treatments. Panels (a)-(c) show situations where both the firms start out with high tentative quantities. While sometimes they revise quantities downwards in small multiple steps as in (b), in other cases they do so in fewer steps (as in (c)). Panels (d)-(i) display the situations where one firm starts out high but the other chooses a relatively lower tentative quantity. While in (d), (g) and (h) firms' downward revisions are correlated resulting in significant lower quantities than the initial plans, panel (e) shows a situation where one firm is signaling to the other one that it is interested in producing lower quantity. Panel (f) displays a pattern where both firms keep on choosing the collusive quantity for a prolonged period of time. Finally, panel (i) shows a very volatile play with both firms revising quantities upwards and downwards frequently. So, these dynamics show that there is considerable heterogeneity in these Poisson revision games. However, unlike the real time revision games, the last second defection is not the primary characteristic of the dynamics.

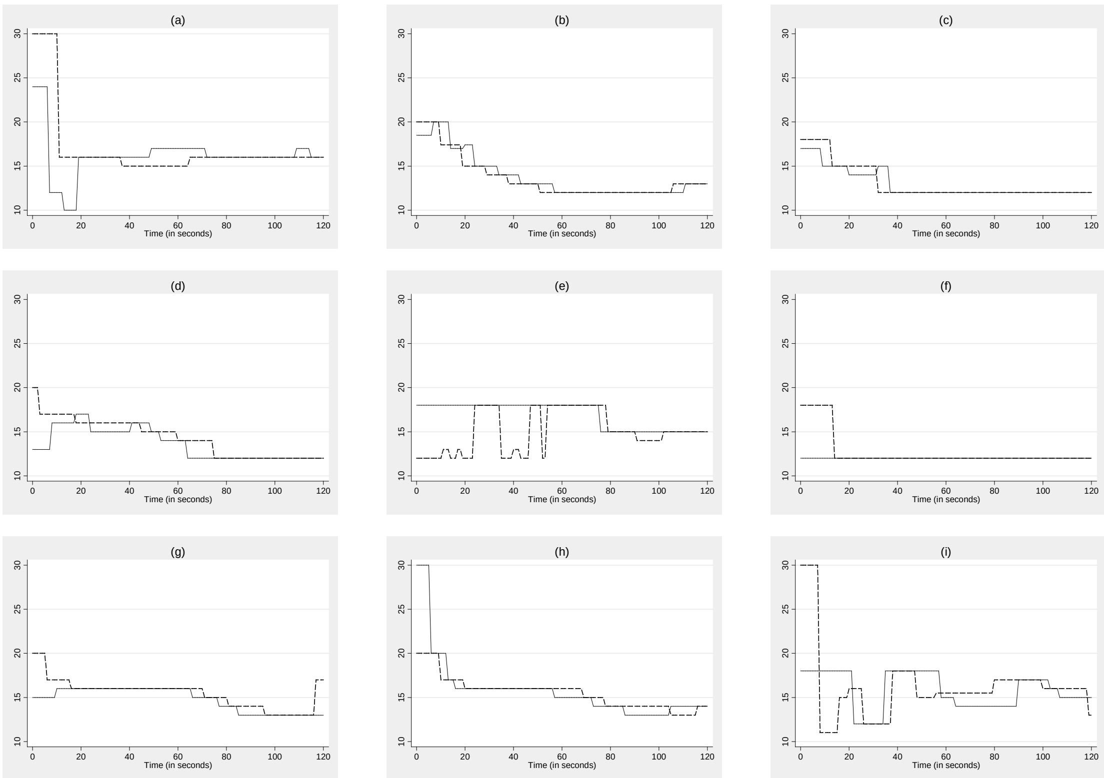

Figure 3.7: Patterns of behavior in a market in “Poisson Revision”. The y-axis gives the tentative quantities selected by the firms in a market over a revision phase (labels are from 10 to 30 in increments of 5) and the x-axis is time in seconds (labels are from 0 to 120 in increments of 20).

3.5.3 Revisions - Best Response versus Imitation

How often do individuals revise their quantities? Do they revise upwards or downwards? Do they best respond to competitor's desired output choices or try to match them (imitation)? Are these behavior different depending on the time left for the revision phase to end? The answers to these questions for all treatments are provided in this section.

To keep the analysis simple, this section uses only a part of the available data. The initial response to a competitor's tentative quantity at time $t = 0$ is taken as the first changed quantity if $t < 10$, or the tentative quantity at $t = 10$ which equals the quantity at $t = 0$ if there is no change. The subsequent data from the revision phase is taken at 10 second intervals, i.e., observations at the following times are taken for analysis in this section[10]: 10, 20, 30, ..., 110, 120. Then it is divided into three sub-phases: the first one consists of data from 10-60 (*first half*), the second from 70-110 (*second half*) and the last one is the final observation at time 120.

Table 3.7 gives the percentage of times for which there is no revision in tentative quantity and also percent of upward and downward output adjustments for the different sub-phases of the entire revision period. As there is no informational feedback in the baseline treatment, the focus of this section is only on the real time revision with monitoring and Poisson revision treatments. As is clear from Table 3.7, there are very few adjustments in tentative quantity until the final sub-phase of the revision period in the treatment with the real time revision phase. And, in the final phase, the quantities selected at time $t = 120$ are higher than those selected at $t = 110$ in more than 50% of the observations. Subjects wait till the last few seconds before revising and when they do revise, the bulk of them adjust quantities upwards.

[10]Dividing the data at 5 seconds intervals gives the same qualitative results.

	Real Time with Monitoring			Poisson High			Poisson Low		
	Downward	No	Upward	Downward	No	Upward	Downward	No	Upward
Initial	3.13	85.42	11.45	28.12	59.38	12.50	26.04	54.17	19.79
First Half	7.29	87.08	5.63	17.50	71.67	10.83	24.58	64.17	11.25
Second Half	4.58	92.08	3.34	11.46	77.08	11.46	10.00	84.58	5.42
Final	10.42	32.29	57.29	12.50	54.17	33.33	6.25	69.79	23.96

Table 3.7: Incidence of revisions subdivided by time in the revision phase (in percentages)

In the Poisson revision treatments, there is a strong downward adjustment in quantity during the initial 10 seconds of the entire revision phase. When there is a revision, 57% of the time quantity is adjusted downwards in the Poisson-Low treatment and 69% in the Poisson-High treatment. The subsequent phases are fairly stable as compared to the activity in the initial phase. As we move into the final 10 seconds of the revision phase, there is an upward trend in the adjustment of quantities. 80% of all the revisions are upward revisions in the Poisson-Low treatment whereas 73% of the time quantities are adjusted above in the Poisson-High treatment. The extent of the late upward revision in the Poisson revision treatments is, however, not as high as that in the real time revision treatment. The revision behavior in the different treatments can be summarized as follows.

Result 6. *While a real time revision game is characterized only by late upward quantity adjustments, the Poisson revision games are characterized both by initial downward adjustments and late upward revisions.*

When there is a change in quantity, this revision in the tentative output can be explained better by one of the following two "obvious behaviors". The first one is the best-response behavior where an individual adjusts her quantity towards the level that is a best response to her competitor's selected desired quantity at time $t-1$. The second type of behavior is imitation in which case a subject tries to "match" the output level of her competitor at time $t-1$. Accordingly, the following model of the adjustment process is estimated[11]:

$$q_i^t - q_i^{t-1} = \beta_0 + \beta_1(r_i^{t-1} - q_i^{t-1}) + \beta_2(q_j^{t-1} - q_i^{t-1}),$$

where q_i^t is the tentative quantity posted by subject i at time t, r_i^{t-1} is the subject i's best response to the competitor's tentative quantity at time $t-1$ and q_j^{t-1} denotes the

[11]The standard errors are clustered at the subject level.

tentative quantity of the competitor at time $t-1$. An individual who strictly plays a myopic best response (imitation) will have $\beta_1 = 1 (\beta_2 = 1)$.

The estimation results are displayed in Table 3.8. First, focus on the real time revision treatment with monitoring. While imitation significantly explains behavior in initial and subsequent phases, in the final phase it is the best-response to a competitor's choice that characterizes behavior. In the Poisson revision treatments, both imitation and best-response behavior are significant in the initial and subsequent phases. However, again similar to the real time revision treatment, best-response behavior explains output adjustment in the final phase of revisions. Thus, the one observation that is consistent across all the treatments is that towards the end of the revision phase behavior is fairly captured by best response dynamics. The above discussion can be summarized as follows.

Result 7. *While individuals imitate their competitor's choices in the beginning, they tend to best respond towards the end of the revision phase.*

3.6 Conclusion

Using a laboratory experiment, this study investigated the effect of different forms of pre-play revisions in two person Cournot games. It showed that with real time revision of quantities but without any informational feedback on competitor's revisions, quantity choices converge to the Cournot-Nash equilibrium. When perfect information is available concerning the competitor's intended quantity along with real time revisions then final output choices are significantly more competitive than the Cournot-Nash output. This shows that more information is detrimental to the firms (when they could revise their quantities in real time) as their average profits are now less than the situation without any information. In contrast, in coordination games with real time revision, information is an essential factor in

	Real Time with Monitoring			Poisson High			Poisson Low		
	β_0	β_1	β_2	β_0	β_1	β_2	β_0	β_1	β_2
Initial	0.58	0.57	1.31***	1.87	0.65***	0.49***	-0.13	0.55***	0.39**
	(4.41)	(1.56)	(0.33)	(1.31)	(0.10)	(0.11)	(0.50)	(0.13)	(0.14)
First Half	0.02	0.65***	0.50***	-0.70	0.40***	0.49***	-0.84***	0.46***	0.37***
	(0.59)	(0.10)	(0.13)	(0.47)	(0.10)	(0.10)	(0.30)	(0.04)	(0.05)
Second Half	-0.84*	0.10	0.95***	0.15	0.30***	0.67***	-0.91***	0.41***	0.42***
	(0.46)	(0.15)	(0.22)	(0.32)	(0.10)	(0.18)	(0.22)	(0.08)	(0.10)
Final	1.15**	0.73***	0.12	0.47	0.83***	0.14	0.94	0.70***	0.22
	(0.42)	(0.06)	(0.17)	(0.40)	(0.07)	(0.15)	(0.87)	(0.20)	(0.19)

Table 3.8: OLS estimation of individual revision behavior. *, **, *** denote significance at the 10%, 5% and 1% levels.

enhancing efficiency. The reason is straightforward. In coordination games, there is a profile of actions that is Pareto efficient where none of the players have any unilateral incentive to deviate, that is, there is a profile which constitutes a Nash equilibrium and is also the most efficient. The problem with coordination games is that there are other equilibria that are Pareto inferior. However, with perfect information players can completely get rid of the uncertainty about the other players' actions, whereas, in the current study, the Cournot game has an unique Nash equilibrium which is not efficient. More information does not help in achieving the efficient profile of actions, rather it serves as a tool to deceive other players.

On a more general level, this chapter showed that even though individuals interact only once in the sense that the payoffs are determined only by the one-time play of the game, a significant amount of cooperation could be sustained. This is achieved through the introduction of a Poisson revision phase where opportunities to revise quantities arrive according to a Poisson process and the payoffs are determined by the quantities chosen at the last revision opportunity before the end of the revision phase. Poisson revision games are stylized representations of a certain type of situation where there are inefficiencies and imperfections in the implementation of desired choices.

The overall dynamics of revisions in real time revision games and Poisson revision games has some interesting differences. While quantities are fairly stable in the games with real time revision phases with "last second defections" to higher quantities, there is a drastic downward adjustment of outputs in the initial phases of the Poisson revision period along with a modest upward trend in desired quantities during the last few seconds of the revision phase. However, there is one similarity in individual behavior across treatments with different forms of revisions. While the quantity adjustments during the initial period of the revision phase show that individuals imitate their opponent's desired quantity choices, behavior towards the end of the revision phase can be explained by the best response to the competitor's desired output.

There are a number of interesting directions to pursue in future research, four of which seem especially promising. First, one could explore markets with asymmetric firms where one firm is a large firm (possibly because of cost advantage) and another is a "small" firm. Do we see endogenous Stackelberg leadership emerging over the course of the revision phase? A second natural extension is to study oligopoly with more than two producers. It would be interesting to investigate whether market quantities are significantly lower under the Poisson revision games than the real time revision games when there are more than two firms in the market. Does increasing the number of competitors result in choices that are more competitive? A related area to explore would be the case of imperfect substitutes. Third, the focus of this study was on situations where revision opportunities are synchronous, yet it is very likely that different players may get these opportunities at different points in time (asynchronous). It would be worthwhile to see, both theoretically and experimentally, if quantity choices are significantly different. Fourth, an important characteristic of the revision games implemented in this chapter was perfect observability of a competitor's intended choices at every point in time. Relaxing this assumption might affect the competitiveness of the market and the equilibrium. For example, allowing participants to observe only the choices at the last received opportunity is likely to induce different dynamics of behavior depending on the initial choices.

Finally, the effect of revision of choices could also be studied in other strategic interactions, for example, in the context of equilibrium selection in coordination games and building cooperation in public goods games.

Chapter 4

Pre-play Communication, Richness of Message Space and Provision of Public Goods

with Thomas Palfrey and Howard Rosenthal

4.1 Introduction

Communication is vital to human relations and is therefore a central concern of game theory. The game-theoretic literature on pre-play communication has shown that individuals reveal "meaningful" information when the conflict is not too severe and there are greater efficiency gains in the presence of private information and coordination problems than when no communication is possible. However, an interesting question is how much communication is needed to have significant efficiency gains in a given setting. Pre-play communication can range from exchange of binary messages, to exchanging numerical messages intended to reveal certain privately observed attributes. The message space, however, doesn't have to be finite either, as is the case with text chat or with face-to-face communication. In our experimental setting, players make simultaneous binary decisions to contribute to the provision of a public good. Hence, an exchange of binary messages is akin to a pre-play version of the game where players get no payoffs in a first round. The other forms of communication are also forms of cheap talk.

In this chapter, we report a series of laboratory experiments to study the effect of three different forms of cheap talk on public good provision in games where player endowments are private information. The game involves the production of a discrete public good by soliciting binary voluntary contributions from group members. The production technology is such that one unit of input is required from at least some fraction of the group. The minimum number of inputs needed to provide the good is referred to as the threshold. The inputs may be thought of as being a fixed contribution of time, effort, or money to some common goal. The cost of the input is privately observed. No side-payments are allowed. Inputs are non-refundable; consequently, it is not possible for a player to reclaim a contribution if the threshold is not reached (under-contribution) or receive a rebate if the threshold is exceeded (over-contribution).

Threshold games like this present severe coordination problems. There is a weak incentive to free ride in the sense that each member of the group would prefer other members to contribute to the public good provision. The no-refund aspect of the game reinforces this free-riding incentive. Dawes et al. (1986) describe the game as a combination of "fear" (wasting a contribution when the good is not provided) and "greed" (the incentive not to make a contribution if contributions of others already meet the threshold). However, on the other hand, when an individual's contribution "makes or breaks" the reaching of the threshold, a player has an incentive to contribute provided her opportunity cost is lower than the benefit from the provision of the good.

There is another aspect of the coordination problem which arises because of the presence of incomplete information and heterogeneous preferences. The ex-ante efficient solution is for just enough members of a group to contribute their unit of input as is needed, and for the contributors to be the ones with relatively low opportunity costs of the inputs while the free riders are the ones who have relatively high valuations for the input. Without communication, it is not possible for players to know each others' relative valuations. However, with sufficient communication, it becomes feasible to coordinate decisions in a way that produces the desired outcome. The problem is, however, that the members of a group may communicate strategically to induce other members to bear the cost of producing the good and this strategically can diminish the potential gains from communication.

Prior to the decision move of the one-shot game, we implement different forms of pre-play communication to investigate how provision improves with communication, depending on the richness of the message space. We find that neither the one-time simultaneous exchange of binary messages to reveal one's intention to contribute, nor the one-time display of a numerical report of private costs is enough to create a significant efficiency gain relative to the situation without any form of communication. Only when participants are provided with the opportunity to engage in natural language communication in an unrestricted fashion, is public good provision (and coordination) significantly improved. Furthermore, the gains relative to a situation with no communication are higher when it is common knowledge that everyone has costs less than the 'public' benefit, than where it is possible for individuals to have higher costs. In the latter situation, individual rationality implies that some individuals have dominant strategies of not contributing. In this setting coordination is much more difficult than in the setting where it is common knowledge that every individual is a possible contributor. In both settings, we find that efficiency is weakly increasing in the richness of the message space.

The remainder of the chapter is organized as follows. The related literature is reviewed in Section 4.2. Section 4.3 specifies the experimental design and lays down the central hypotheses. Section 4.4 presents the experimental results and analysis. The last section concludes.

4.2 Related Literature

That selfish players may choose to reveal private information through costless and non-binding communication or cheap talk, and that such revelation can lead to efficiency gains, has been shown by Crawford and Sobel (1982). Palfrey and Rosenthal (1991) were the first to investigate the effects of cheap talk in a model where players have private information about costs to contribute towards the provision of a public good. They considered a 'threshold public good game' where provision requires contributions from at least a minimum number of people. Using a binary communication setting, they showed that perfect coordination is not Bayesian-incentive compatible and that players have weak incentives to free-ride in these kind of situations but there are communication equilibria that lead to higher efficiency. Using a model of continuous contributions with two privately informed players, Agastya et al. (2007) show that individuals do not have an incentive to contribute to the public good without communication. However, binary communication gives them incentives to provide the good with positive probability.

Kawamura (2011) uses an n-player setting and shows that there always exists an equilibrium where binary messages are credible and that, when n goes to infinity, the equilibrium with binary communication is the most efficient one. Costa and Moreira (2012) find that a truthful equilibrium of the communication game with a binary message space cannot be Pareto dominated by any truthful equilibrium with any finite message space. They argue that this is because players use a threshold rule to make contribution decisions and in any situation with more than one contribution threshold, they have incentives to understate their types and thereby free-ride on their partners' investment.

Other interesting theoretical works in the area of strategic information transmission include political bargaining (Matthews (1989), Matthews and Postlewaite (1989)), externalities associated with technological adoption and standardization (Farrell (1993), Farrell and Saloner (1988)), agendas and straw votes (Ordeshook and Palfrey (1988)), bargaining (Farrell and Gibbons (1989)), and applications to job markets (Forges (1990)) and arms races (Baliga and Sjostrom (2004)).

Some laboratory research has examined the effect of non-binding pre-play communication in threshold public goods games. Van de Kragt et al. (1983) show that free communication via general, unstructured discussion produces better outcomes than the same games conducted without communication in the context of public goods games similar to the one implemented by us. The message space was the entire English language and speaking order was entirely endogenous, occurring in continuous time with face-to-face communication. An important distinction from our game is that the contribution costs were equal for all players and common knowledge. Thus, Van de Kragt et al. (1983) eliminated two important impediments to coordination - private information and heterogeneity in costs.

Palfrey and Rosenthal (1991) report results from games similar to the three-person threshold games considered in the current study and conclude that in the absence of communication, behavior is closely approximated by the Bayesian equilibrium predictions, except that subjects contribute slightly more often than predicted. They implement a binary message stage prior to the decision making stage and find that this type of communication fails to provide more efficient outcomes compared to the 'no communication' outcome. They also find that subjects use a cutoff decision rule when it is optimal to do so in the 'no communication' treatment, while players' behavior is less systematic with communication.

Several experiments have implemented different forms of communication structures to evaluate the effect of cheap talk in other games. Cooper et al. (1992) studied one-way communication and two-way communication in coordination games and concluded that allowing pre-play communication does not uniformly lead to the play of the Pareto-dominant Nash equilibrium. There are situations where one-way communication performs significantly better than two-way communication. In a standard repeated VCM[1] model, Bochet et al. (2006) vary the message space used in the cheap talk games and find that the treatment with exchange of numerical messages does not affect efficiency compared to the situation with no communication. However, both types of verbal communication, face-to-face and anonymous chat, increase cooperation[2]. Costa and Moreira (2012) implement a two-person contribution game and find evidence that larger message spaces do not provide efficiency gain relative to the binary communication structure.

4.3 Experimental Design, Benchmarks and Hypotheses

The experimental design, procedures and treatments are discussed in Section 4.3.1. Section 4.3.2 provide the theoretical benchmarks. The hypotheses that are formally tested are listed in Section 4.3.3.

[1]Under a voluntary contribution mechanism (VCM), individuals voluntarily allocate their initial holdings of resources into the production of public and private goods. Certain assumptions on the payoffs and utilities are made such that there exists a dominant strategy to not contribute anything to the production of the public good and "free-ride" on the contributions of others.

[2]The result that face-to-face communication increases contributions in repeated VCM experiments has been shown by various researchers, including Isaac and Walker (1988), Cason and Khan (1999), Brosig et al. (2003), Belianin and Novarese (2005). Ostrom and Walker (1991) also find the same result in repeated common property resource situations.

4.3.1 Design, Procedures and Treatments

The experiments were conducted at the California Social Science Experimental Laboratory (CASSEL), University of California, Los Angeles (UCLA) using the Multistage software package. Participants were recruited from a pool of volunteer subjects, maintained by CASSEL. A total of sixteen sessions were run, using 183 subjects. Each session consisted of 9-15 participants and no subject participated in more than one session[3]. Upon arrival, instructions were read aloud. Subjects interacted anonymously with each other through computer terminals. Sessions lasted from 30 to 50 minutes and participants earned on average US$18.47 in addition to a show-up fee of US$5[4].

We consider a model where a group consisting of three persons is undertaking a project. Each group member is endowed with one indivisible unit of input, which may be either consumed or "contributed" to the production of the group project. The project succeeds if and only if at least two units are contributed. The value of the project to any individual is normalized to equal 100. The private value of the endowed unit of input to an individual i is denoted by c_i. Each person knows his or her own c_i but only knows that the other players' c's are independent random draws from some common probability distribution with CDF $F(.)$. We assume that there exist c_{min}, c_{max}, with $0 \leq c_{min} \leq c_{max}$, $F(c_{min}) = 0$, $F(c_{max}) = 1$; $f = F'$ exists and is continuous and strictly positive on $[c_{min}, c_{max}]$. The utility for player

[3]Eleven sessions had 12 subjects, four sessions had 9 subjects and one session had 15.
[4]Payoffs ranged from US$11 to US$25.50 with a standard deviation of US$3.18.

i with cost c is given by:

$$\begin{array}{rll} 100+c & \text{if} \quad i & \text{does not contribute and two others contribute,} \\ c & \text{if} \quad i & \text{does not contribute and fewer than two others contribute,} \\ 100 & \text{if} \quad i & \text{contributes and at least one other contributes,} \\ 0 & \text{if} \quad i & \text{contributes and none of the others contribute.} \end{array}$$

We assume that the common probability distribution is uniform and we implement the above game for two separate values of c_{max}: 100 and 150. $c_{max} = 100$ means that it is common knowledge that everyone has costs lower than the 'public benefit' and hence, no one has a dominated strategy. However, in the situation with $c_{max} = 150$, individuals have a dominated strategy not to contribute whenever $c_i > 100$. Thus, the strategic aspect is now different from the previous situation, where no strategy is dominated.

The communication structure was varied in each of the two parametric configurations, $c_{max} = 100$ and $c_{max} = 150$. We used four treatments: "no communication", "binary communication", "token revelation", and "unrestricted text chat". We ran two sessions for each of these treatments, thus a total of 8 sessions per c_{max} configuration. These are briefly discussed below[5]. All sessions had 20 rounds. After a round was over, participants were randomly rematched into new three person groups and everyone was independently and randomly assigned new token values. The random rematching was done to limit the reputation and super-game effects which can occur with repeated play.

[5]For details, please see the instructions provided in Appendix D.

4.3.1.1 No Communication

In each round, subjects were each given a single indivisible "token". Token values in integer increments between 1 to c_{max} points (also referred to as costs) were independently drawn with replacement from identical uniform distributions and randomly assigned to subjects. Each subject was told the value of her token but told only the probability distribution of values of the tokens of other subjects. Subjects were then asked to enter their decisions to spend or keep their token. If at least two of the other three subjects spent, each subject received 100 points if she was a "contributor" while the payoff was 100 points plus the token value in case where the subject was not a contributor. If a subject contributed but none of her group members spent their tokens then the person who contributed earned 0 points.

4.3.1.2 Binary Communication

Each round had two stages: a communication stage and a contribution stage. In the communication stage, subjects chose one of the two messages: "I intend to spend my token"; "I intend to keep my token". They were advised that these messages were not binding, and they could make either contribution decision regardless of which message they sent in the communication round. After these binary messages were sent, each person was told how many members in their group sent each message, and was reminded which message he or she had sent. This was directly followed by the contribution stage, where individuals made binding contribution decisions.

4.3.1.3 Token Revelation

Each round again had a communication stage and a contribution stage. In the communication stage, subjects had 20 seconds to send a message to the other members of her group.

They were told that this message can only be an integer between 1 and c_{max} and each member was allowed to send only one such message. Thus, the message space in this treatment is larger than the "binary communication" sessions. Each subject could observe the messages send by her group members. The subjects were also told that in the situation where she did not send any message, the other members of her group would see a " ? " against her subject id at the end of the 20 seconds. After the communication stage was over, individuals made contribution decisions. We called this treatment "token revelation" because subjects had the opportunity to reveal their true token values. If token values were revealed truthfully, there would be the opportunity for perfect coordination in the $c_{max} = 100$ condition, with the two individuals with the lowest token values contributing. The same could occur in the $c_{max} = 100$ condition, provided two individuals had token values no greater than 100.

4.3.1.4 Unrestricted Text Chat

This treatment consisted of the same number of rounds as in the other ones and had the same two stages, with the only difference being in the structure of the communication stage. Prior to the contribution stage, every group had a discussion period which lasted 60 seconds, during which subjects could send messages to the other members of her group. Individuals were told that the messages had to conform to certain rules, including that they must be relevant to the experiment and subjects should not send messages intended to reveal their identity. Thus, the message space under this treatment is 'anonymous, unrestricted and unstructured text messages' and is much larger than the finite message space under the "token revelation" sessions. Importantly, this treatment gave subjects the opportunity to employ natural language.

4.3.2 Benchmarks

In this subsection, we discuss the different benchmarks against which we will compare the performance of our experimental data. The first benchmark is the "trivial" equilibrium that exists in our game where each of the three players keep their tokens. Thus, in this '*zero contribution equilibrium*' the total group earnings equal the sum of the members' token values.

Without communication, a natural benchmark "mechanism" is one in which individuals either contribute or do not and make their decisions independently and simultaneously. The symmetric Bayesian equilibria of our game are of a particular simple form[6]. For any beliefs that player i has about the other players' decisions to contribute, there is a unique best response strategy which is a *cutpoint rule*. That is, there is a threshold cost level (say c^*), such that contribution is optimal if $c_i < c^*$ and non-contribution is optimal if $c_i > c^*$. c^* can be determined from the solution to the following equation:

$$c^* = \binom{3-1}{2-1}[F(c^*)][1 - F(c^*)] \cdot 100 \quad .$$

The interpretation of the above equation is that a person with a private cost of c^* faces an opportunity cost of contributing equal to the expected gain from contributing. At equilibrium, everyone with a cost below c^* is better off contributing, given others are using the c^* decision rule and everyone with private costs greater than c^* is better off not contributing. Individuals with a cost of exactly c^* are indifferent between contributing and not contributing. There might be multiple solutions for c^*. In fact, $c^* = c_{min}$ is always an equilibrium corresponding to the situation where nobody contributes. However, for the parameters of

[6] This is shown in Palfrey and Rosenthal (1988). Also see Palfrey and Rosenthal (1991) for more details.

our experiments, there is a unique solution for $c^* > c_{min}$. The '*Bayes-Nash equilibrium*' strategy for the treatment with $c_{max} = 100$ is: contribute if $c_i < 50$; do not contribute otherwise. The cutpoint c^* is lower for the $c_{max} = 150$ case and the equilibrium strategy is: contribute if $c_i < 37.5$; do not contribute otherwise.

With communication, additional equilibria arise. Players can now use their joint messages as a correlating device. While there are multiple equilibria with communication[7], we will focus on the optimal equilibrium in the class of incentive-compatible mechanisms. Under a mechanism, individuals submit their costs to a "mediator" and then the planner determines whether the good is provided and decides who contributes. In such a mechanism, the good is provided if and only if there are at least two costs reported that are less than a threshold value and the contributors are randomly chosen if all three costs are less than the threshold; otherwise the two people who report costs less than the threshold value contribute. For an equilibrium, we need to find a $\hat{c}$, such that a person with $c_i = \hat{c}$ is indifferent to the choice between contributing and not contributing, that is:

$$100 \cdot [F(\hat{c})]^2 = [F(\hat{c})]^2(100 - \frac{2}{3}\hat{c}) + 2F(\hat{c})[1 - F(\hat{c})](100 - \hat{c}) \quad .$$

The left hand side of the above equation gives the expected payoff from not contributing and the right hand side is the expected payoff from contributing. For $c_{max} = 100$, we conjecture that the *best IC mechanism* would prescribe that the good is always produced and two out of the three people in a group are randomly selected to contribute. This randomization is possible to achieve without a mediation, through jointly controlled lotteries, in our communication treatments, except for the binary communication game. In the case of $c_{max} = 150$, our conjecture is that the *mechanism* would prescribe the following: (a) if everyone submits

[7]Palfrey and Rosenthal (1991) discuss some of the equilibria in greater detail.

a cost, $c_i \leq 75$, then the good is produced and two out of the three people in a group are randomly selected to contribute, (b) if two people submit $c_i \leq 75$, then those two contribute and the good is provided, and (c) if less than two submit $c_i \leq 75$, then nobody contributes.

Group earnings are maximized if the individuals with the lower two token values contribute while the individual with the highest token value keeps her token. This '*first best outcome*', however, does not constitute an equilibrium as there would be incentives to over report token values. While this is feasible in the $c_{max} = 100$ configuration, in the $c_{max} = 150$ situation individual rationality will make the '*first best outcome*' infeasible if one or both of the lower two token values are higher than 100. In this $c_{max} = 150$ configuration, one can calculate the group earnings under the '*constrained first best outcome*' which would prescribe that (a) individuals with the lower two token values only contribute if both of the lower two token values are less than or equal to 100, and (b) no one contributes if one or both of the lower two token values are higher than 100. This is the first-best situation subject to the individual rationality constraint.

4.3.3 Hypotheses

We test several hypotheses regarding (a) the effect of communication on efficiency, (b) comparison across communication treatments differing in the richness of message space, and (c) differences across $c_{max} = 100$ and $c_{max} = 150$ sessions.

The effect of communication on efficiency is broken down into four separate hypotheses with respect to the total earnings generated, the likelihood of public good provision, the number of contributors and the costs of contributors. First, as there exist equilibria under communication that have higher payoffs for players than in the Bayesian Nash equilibrium without communication, one can conjecture that the earnings would indeed be higher if communication is allowed. Also, intuitively one can expect that the public good is provided more often under communication, thereby attaining Pareto superior outcomes, than under no communication. These give rise to our first two hypotheses:

Hypothesis 1 (Earnings Hypothesis). Total earnings are higher with communication than without communication.

Hypothesis 2 (Provision Hypothesis). The likelihood of public good provision is greater in the communication sessions than in the non-communication sessions.

The ex ante efficient solution is for just enough individuals to contribute their token as needed, and for the contributors to be the ones with relatively low costs while the free rider is the one with the highest cost. Without communication, it is impossible for players to know who has relatively high valuations and who has relatively low ones. Thus the "efficient" outcome can occur only by chance. However, with sufficient communication, it is at least feasible to coordinate decisions in a way that produces this desired outcome. Thus, we can expect to have fewer wasteful contributions, that is, fewer incidences of one or three individuals contributing in the case with communication. In other words, communication leads to lower production inefficiency. Also, conditional on the public good being provided, we should have a higher percentage of the subjects with the two lowest costs contributing when communication is allowed. Thus, communication can help in two ways: one by reducing coordination failures (such that only two contribute) and other by making the two lowest costs contributing. Hence, the next two hypothesis are as follows:

Hypothesis 3 (Coordination Hypothesis). The incidence of wasteful contributions is lower in the communication sessions than in the non-communication sessions.

Hypothesis 4 (Efficiency Hypothesis). Conditional on the public good being provided, the incidence of individuals with the lower two costs contributing is higher with communication *than under no communication.*

We have three different forms of communication structures that we study. While the 'binary communication' sessions use a simultaneous exchange of binary messages, the 'token revelation' sessions have a much larger although finite message space and also are not simultaneous. The timing of the token reports is endogenously determined. The 'unrestricted text chat', while anonymous, uses the entire English language in continuous time and the text order is entirely endogenous. Hence, we have a progression in the treatments in terms of the richness of the message space implemented. One would expect that it would be easier to implement the "efficient" outcome as the message space is enlarged and hence we would have a monotonic relation between efficiency and richness of message space. This gives rise to the following hypothesis:

Hypothesis 5 (Monotone Hypothesis). Efficiency increases with the richness of the message space. It is highest in the 'unrestricted text chat' sessions followed by the 'token revelation' sessions, followed by the 'binary communication' sessions. Efficiency is lowest under no communication.

The final hypothesis compares the provision of public good across the two different c_{max} sessions. When it is common knowledge that everyone's costs are less than the benefit from the public good, 100% provision of public good is possible. However, in the $c_{max} = 150$ case, it is no longer possible to provide the public good all the time assuming subjects do not use dominated strategies. Indeed the probability that the good isn't provided equals $\frac{7}{27}$, when at least two of the subjects have costs greater than 100. This gives us the following hypothesis:

Hypothesis 6 (c_{max} Hypothesis). Keeping the communication protocol fixed, the probability of public good provision is greater in the $c_{max} = 100$ sessions than in the $c_{max} = 150$

sessions.

4.4 Results

The results from the experiments are reported in this section. We first present the results from testing the specific hypotheses that we discussed earlier. Subsection 4.4.2 compares the performance of the data in all treatments to the various theoretical benchmarks. The final subsection briefly discusses the behavior in the communication stages of the different treatments.

4.4.1 Test of Hypotheses

The first objective is to compare the average earnings across the different treatments. Tables 4.1 and 4.2 document the average group earnings net of token values in rounds 1-10, 11-20 and all rounds across the four treatments, in both c_{max} sessions. In the $c_{max} = 100$ sessions, average net earnings are higher only in the 'unrestricted chat' treatment. The other two communication treatments yield similar earnings as those under 'no communication'. In the $c_{max} = 150$ sessions, 'unrestricted chat' results in higher earnings only in the second half of the experiment. We use the net earnings as opposed to gross earnings to reduce the 'randomness' present due to the realization of the drawn token values. Figure 4.1 plots the group earnings (net of token values) as a fraction of the first-best earnings possible (net of token values) for rounds 11-20. We use this measure as the variable for statistical tests. Also, unless otherwise mentioned, throughout this section, we use the data from rounds 11-20 only (when subjects have gathered some experience) for performing the statistical analyses.

Taking each group as an independent observation, one-tailed Mann-Whitney tests show

Rounds	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
1-10	113.2	89.7	95.8	176.8
	(15.7)	(14.0)	(14.8)	(11.9)
11-20	88.0	80.2	101.5	221.1
	(15.3)	(14.7)	(14.6)	(7.1)
All rounds	100.6	84.9	98.6	198.9
	(10.9)	(10.1)	(10.4)	(7.2)

Table 4.1: Average group earnings net of token values: $c_{max} = 100$ sessions. Standard errors in parentheses.

Rounds	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
1-10	67.0	82.0	55.1	48.5
	(14.0)	(13.8)	(13.8)	(12.4)
11-20	39.6	58.7	42.1	83.4
	(12.4)	(13.4)	(11.9)	(13.9)
All rounds	53.3	70.3	48.6	66.0
	(9.4)	(9.6)	(9.1)	(9.4)

Table 4.2: Average group earnings net of token values: $c_{max} = 150$ sessions. Standard errors in parentheses.

that there is no significant difference in the net group earnings as a proportion of the maximum net earnings possible across 'binary communication' and 'no communication' treatments and between 'token revelation' and 'no communication' treatments. However, the differences are statistically significant at the 1% level between 'no communication' and 'unrestricted chat' implying that only when subjects are provided with unstructured communication in the form of text chat, they end up with higher earnings. This is true for both c_{max} sessions. While differences are significant between 'unrestricted chat' and 'no communication' for all rounds in $c_{max} = 100$, this is no longer the case if all rounds are considered in the $c_{max} = 150$ sessions. Indeed, it is clear from Tables 4.1 and 4.2 that earnings are lower or similar in rounds 1-10 and 11-20 in all treatments except 'unrestricted chat', in which case there is an increase in payoffs in the second half of the sessions as compared to the first half. The above discussion leads us to support *Earnings Hypothesis* and conclude the following:

Result 1. Total earnings are higher with communication only in the third, 'text' condition, where natural language is employed (support for H1).

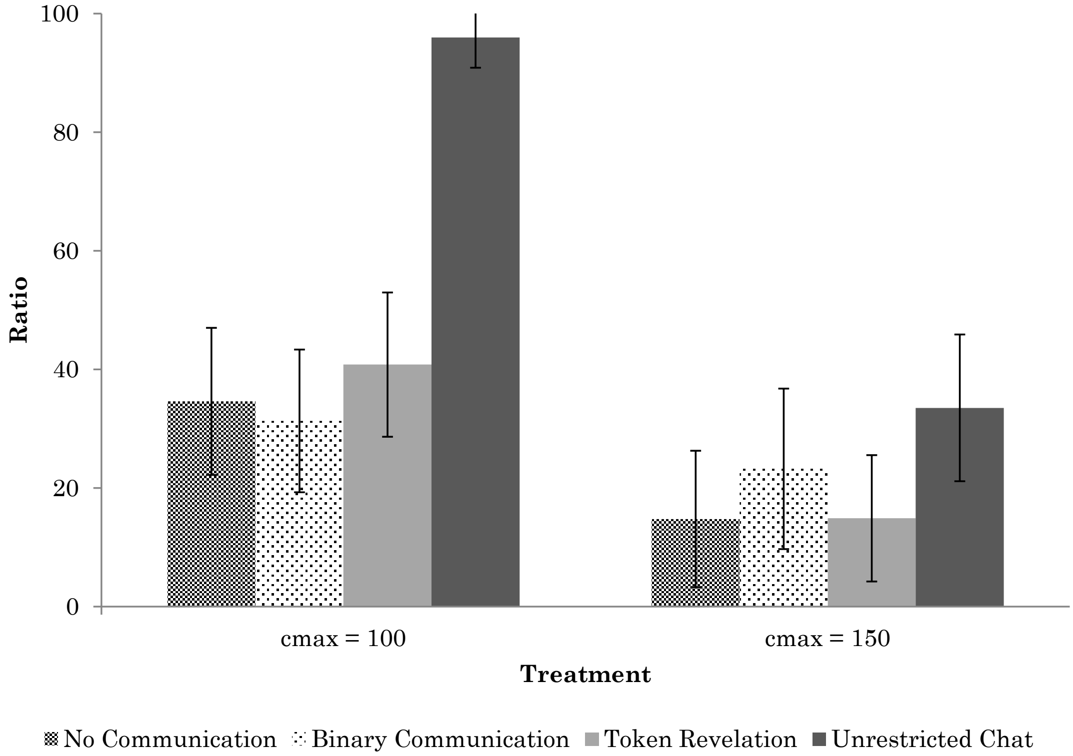

Figure 4.1: Average group earnings (net of token values) as a percentage of the first-best earnings (net of token values). Also shown are the 95% confidence intervals. Data from rounds 11-20.

Tables 4.3 and 4.4 collect the probabilities of the public good provision for all rounds, rounds 1-10, and rounds 11-20 in each of the four session. An immediate finding is that in line with the previous result, not all forms of communication increase the likelihood of the good being provided. Clearly, 'binary communication' and 'token revelation' sessions result in either similar or lower public good provision when compared to 'no communication' sessions, considering the two c_{max} sessions. It is only in the 'unrestricted chat' that the provision of good is higher in both $c_{max} = 100$ and $c_{max} = 150$ sessions. Moreover, comparing the numbers for rounds 1-10 and 11-20 reveal that provision declines over time in all other treatments[8] except the 'unrestricted chat' where the opposite is true. A difference in proportions z-test for each of the binary comparisons of the 'communication' sessions with the 'no communication' sessions shows that there is significant difference only in the 'binary communication' and 'no communication' in $c_{max} = 150$ (at 10% level), and in 'unrestricted chat' and 'no communication' in both c_{max} sessions (at 5% level)[9]. Supporting *Provision Hypothesis*, we have the following:

Result 2. The likelihood of public good provision is greater with communication only when natural language is employed (support for H2).

Rounds	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
1-10	54.3 (70)	47.5 (80)	48.8 (80)	85.0 (60)
11-20	42.9 (70)	41.3 (80)	50.0 (80)	98.3 (60)
All rounds	48.6 (140)	44.4 (160)	49.4 (160)	91.7 (120)

Table 4.3: Probability of public good provision: $c_{max} = 100$ sessions. Number of observations in parentheses.

Rounds	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
1-10	38.8 (80)	43.8 (80)	38.8 (80)	31.3 (80)
11-20	27.5 (80)	37.5 (80)	25.0 (80)	42.5 (80)
All rounds	33.1 (160)	40.6 (160)	31.9 (160)	36.9 (160)

Table 4.4: Probability of public good provision: $c_{max} = 150$ sessions. Number of observations in parentheses.

Next, we turn to the following question: Does communication help in reducing production inefficiency in the form of wasteful contributions? The answer here is yes. As is

[8]For the 'token revelation' there is a marginal increase in $c_{max} = 100$ session but a considerable decline in the other $c_{max} = 150$ session.

[9]These are results from one-tailed tests of difference in two independent proportions. For the $c_{max} = 100$ situation, the data numbers do not satisfy the criteria: $n(1-p) > 5$, where n is the number of observations and p is the proportion. But its easy to infer that the proportion is almost 1 under 'unrestricted chat' and hence there would be a significant difference when compared with the 'no communication' session.

evident from Tables 4.5 and 4.6 taken together, the percentage of either one person or three persons contributing is always lower with communication. Again, one-tailed proportions z-tests show that while these differences are significant for 'binary communication' and 'unrestricted chat' sessions[10], there is no significant difference between incidence of wasteful contributions between 'no communication' and 'token revelation', although the numbers are lower with 'token revelation'. Also, there is a decline in the proportion of wasteful contributions over time in two out of the three communication treatments, whereas there is always an increase in production inefficiency in the 'no communication' sessions over time. Overall, we support *Coordination Hypothesis*, and the following result is obtained.

Result 3. Incidence of wasteful contributions is lower in the communication sessions than in the non-communication sessions (support for H3).

Rounds	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
1-10	50.0 (70)	51.3 (80)	53.8 (80)	20.0 (60)
11-20	60.0 (70)	47.5 (80)	51.3 (80)	3.3 (60)
All rounds	55.0 (140)	49.4 (160)	52.5 (160)	11.7 (120)

Table 4.5: Percentage of wasteful contributions: $c_{max} = 100$ sessions. Number of observations in parentheses.

Rounds	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
1-10	52.5 (80)	36.3 (80)	52.5 (80)	48.8 (80)
11-20	55.0 (80)	41.3 (80)	47.5 (80)	35.0 (80)
All rounds	53.8 (160)	38.8 (160)	50.0 (160)	41.9 (160)

Table 4.6: Percentage of wasteful contributions: $c_{max} = 150$ sessions. Number of observations in parentheses.

The most efficient outcome occurs when exactly two people in a group contribute and those two happen to be the subjects with the lower two costs. Conditional on the good being

[10]The p-values are < 0.01 for the 'unrestricted chat' vs. 'no communication' test, while they are < 0.05 (0.10) in the case of 'binary communication' vs. 'no communication' with $c_{max} = 150 (c_{max} = 100)$.

provided, Tables 4.7 and 4.8 document the percentage of times the subjects with the lower two costs contribute across the four treatments and for each c_{max} session. The numbers are quite high, even without communication. With $c_{max} = 100$, full efficiency would require 100% of the time the two lowest costs contribute. Obtaining 86.4% in the last 10 rounds of unrestricted chat is impressive. The loss in value from the misallocated contributions is only 8.2 per round (See Figure 4.3). With $c_{max} = 150$, the benchmark for the actual draw of tokens is 74% given individual rationality constraints.

Again, the incidence of subjects with the two lowest costs in a group contributing is higher in the 'unrestricted chat' than the 'no communication' sessions (difference is significant at the 1% level for $c_{max} = 100$), but this is not true when the other 'communication' sessions are compared with 'no communication' ones. Thus, similar to Results 1 and 2, we conclude that only one form of communication helps and have the following result.

Result 4. Conditional on the public good being provided, the incidence of individuals with the lower two costs contributing is higher with communication than under no communication only when natural language is used (support for H4).

Rounds	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
1-10	60.5 (38)	42.1 (38)	61.5 (39)	74.5 (51)
11-20	56.7 (30)	57.6 (33)	65.0 (40)	86.4 (59)
All rounds	58.8 (68)	49.3 (71)	63.3 (79)	80.9 (110)

Table 4.7: Percentage of two lowest costs contributing: $c_{max} = 100$ sessions. Number of observations where the public good is provided is in parentheses.

Efficiency loss of one form or another occurs when the individual with the highest token value in a group contributes. So, the count of such occurrences also qualifies as a measure of inefficiency. This happens in 13 out of 70, 17 out of 80, 14 out of 80 and 8 out of 60 cases in the 'no communication', 'binary communication', 'token revelation' and 'unrestricted chat'

Rounds	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
1-10	64.5 (31)	71.4 (35)	67.7 (31)	80.0 (25)
11-20	68.2 (22)	70.0 (30)	55.0 (20)	79.4 (34)
All rounds	66.0 (53)	70.8 (65)	62.7 (51)	79.7 (59)

Table 4.8: Percentage of two lowest costs contributing: $c_{max} = 150$ sessions. Number of observations where the public good is provided is in parentheses.

cases respectively, in the $c_{max} = 100$ sessions[11]. In the $c_{max} = 150$ sessions, this count is 9 out of 80 in 'no communication' as well as 'token revelation' treatments, 10 out of 80 in the 'binary communication' and 7 out of 80 in the 'unrestricted chat' treatment.

Summarizing the above results, we find that the structure of communication is crucial in determining whether we can achieve any efficiency gain over the situation where no communication is allowed. The simultaneous exchange of binary messages aimed at disclosing intentions as well as the one-time broadcasting of a numerical message fail to improve average earnings and the probability of the public good provision. In these sessions, subjects also are not able to coordinate such that only the ones with lower two token values contribute. Only when unstructured communication with a "common language" is allowed, are there gains in efficiency. While these gains are enormous in the sessions where it is common knowledge that everyone has a cost that is less than the benefit from the public good, it takes time to capture some of the gains in the other sessions where individuals might have costs that exceed the benefit. All three communication structures, however, do help in reducing the wasteful contributions.

Given results 1-4, it is now straightforward to test Hypothesis 5. Whether it is measured with respect to the average group earnings, likelihood of public good provision, incidence of wasteful contributions, or the subjects with lower two costs contributing, there is no

[11]Again this is using data from last ten rounds. This also includes situations where two of the individuals in a group shared the highest token value and both of them contributed. There was no such instance of all three subjects having same token value draws in one group.

strict progression in efficiency as we move from no communication to binary to a larger numerical message space to finally the "infinite" communication structure. The relation between efficiency and the richness of message space is hence non-monotonic. Rejecting *Monotone Hypothesis*, we now have the following result.

Result 5. Efficiency does not increase monotonically in the richness of the message space of the communication stage (reject H5). It does increase when 'unrestricted chat' is permitted.

Now, focusing on the $c_{max} = 150$ sessions, do subjects use dominated strategies? The answer is a clear no. In 632 cases where the token values were greater than 100, there were only 11 observations of contribution. Broken down by treatments, the number of times subjects used the dominated strategy was 1 out of 160 in 'no communication', 4 out of 159 in 'binary communication', 5 out of 160 in 'token revelation' and 1 out of 153 in the 'unrestricted chat' sessions. Thus, in the $c_{max} = 150$ sessions, a subject with a cost which exceeds the benefit of 100 does not contribute.

The final result compares the likelihood of providing the public good in $c_{max} = 100$ and $c_{max} = 150$ sessions. There were 489 instances out of 1918 total observations (that is 25.5% of the cases) when at least two of the subjects in a group had costs higher than 100 in the $c_{max} = 150$ sessions. The public good was provided in 6 out of 114 observations in the 'binary communication' treatment. In the other treatments, the good was never provided when there was only one person in a group having a cost less than 100 (out of a total of 375 such instances). These numbers strongly support the intuition behind Hypothesis 6. Figure 4.2 plots the probability of public good provision in the two different c_{max} sessions for each of the treatments in the last ten rounds. As can be seen from the figure (and also Tables 4.3 and 4.4), the pubic good provision is higher in the $c_{max} = 100$

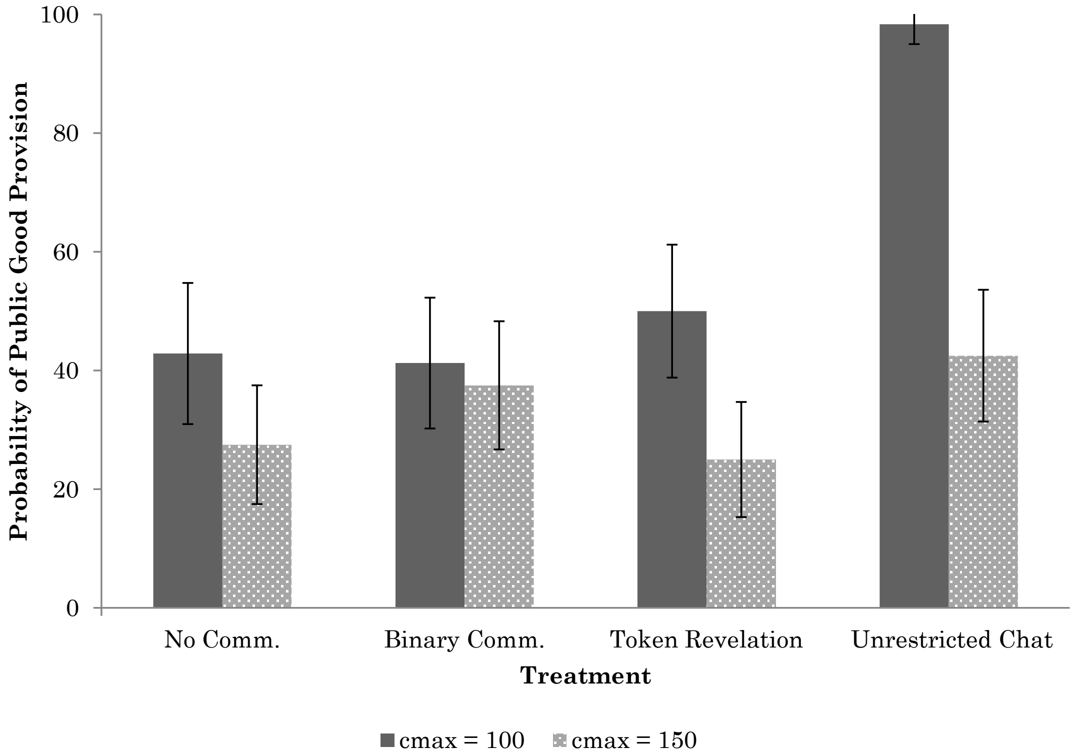

Figure 4.2: Probability of public good provision: $c_{max} = 100$ vs. $c_{max} = 150$. Also shown are the 95% confidence intervals.

sessions than in the corresponding $c_{max} = 150$ sessions. The differences are significant at the 1% level for 'token revelation' and 'unrestricted chat' sessions and at the 5% level for the 'no communication' sessions. The difference in the 'binary communication' treatment is, however, not significant, even at the 10% level. Hence, supporting c_{max} *Hypothesis*, we have:

Result 6. Keeping the communication protocol fixed, the probability of public good provision is greater in the $c_{max} = 100$ *sessions than in the* $c_{max} = 150$ *sessions (support for H6).*

4.4.2 Comparison to Benchmarks

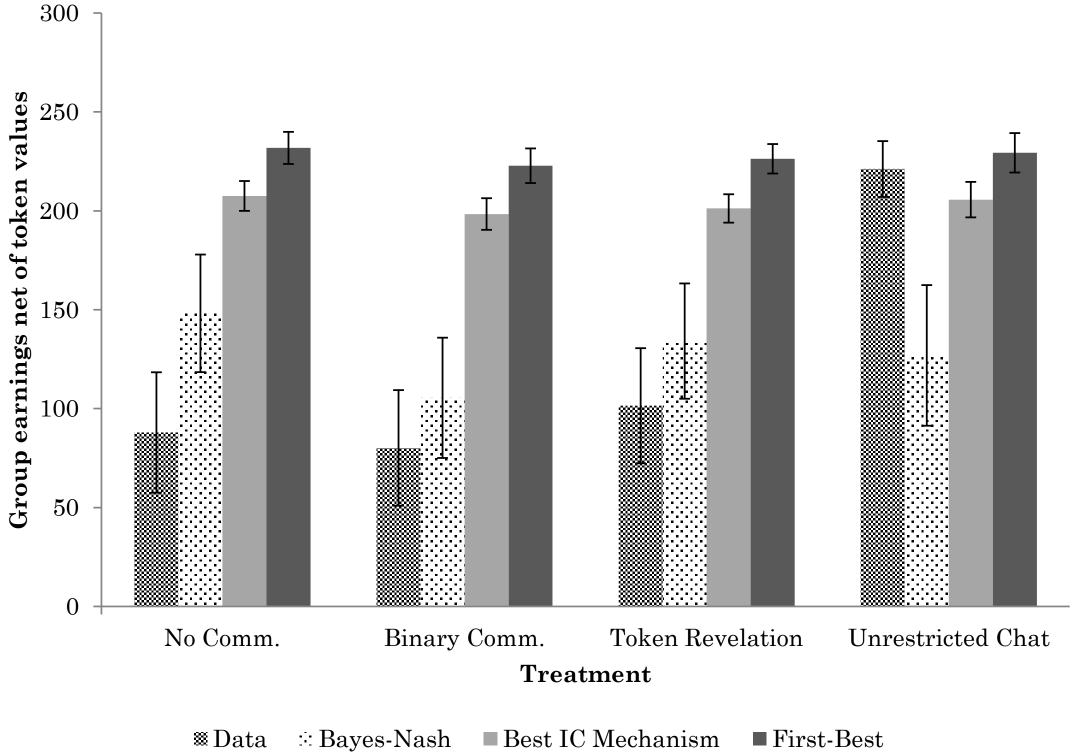

Figure 4.3: Average group earnings net of token values by treatments (rounds 11-20) in c_{max} = 100 sessions: comparison to benchmarks. Also shown are the 95% confidence intervals.

We now turn to the analysis of how the data fares when compared to the benchmarks discussed in section 4.3.2. Figure 4.3 plots the average group earnings (net of token values) for rounds 11-20 from data in the $c_{max} = 100$ sessions as well as the net earnings that one would get under the Bayes-Nash, the best incentive compatible mechanism and the first-best situation. Individuals do much better in terms of payoffs than what they would have earned if all of them decided to just keep their tokens. However, subjects are worse off (in payoffs) than in not only the best IC mechanism and subsequently the first-best, but also than in the Bayes-Nash outcome in all other treatments except 'unrestricted chat'. See Table 4.9 for the differences in net group earnings as well as the significance of these differences from a Wilcoxon signed-rank paired difference test. Interestingly, in the 'unrestricted chat', subjects do even better than the best IC mechanism and reach closer to the first-best outcome.

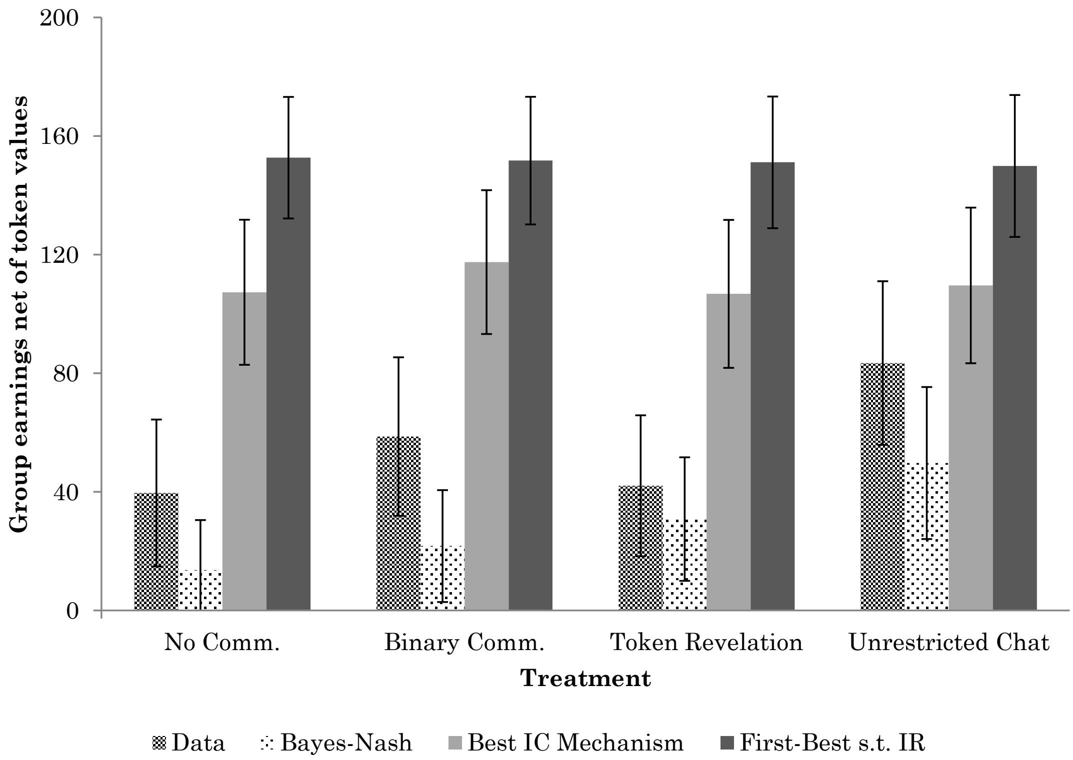

Figure 4.4: Average group earnings net of token values by treatments (rounds 11-20) in c_{max} = 150 sessions: comparison to benchmarks. Also shown are the 95% confidence intervals.

Next, Figure 4.4 graphs the net earnings for rounds 11-20 from the data and other benchmarks in the $c_{max} = 150$ sessions. Unlike the $c_{max} = 100$ sessions, payoffs are not

Comparison	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
Data minus Bayes-Nash	-60.2***	-25.3**	-32.7**	94.2***
Data minus Best IC Mechanism	-119.6***	-118.2***	-99.7***	15.5***
Data minus First-Best	-143.8***	-142.6***	-124.8***	-8.2

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 4.9: Differences in average net group earnings: $c_{max} = 100$ sessions.

always significantly higher than the zero contribution situation. Under 'no communication' and 'token revelation' treatments, subject earnings are statistically indistinguishable from the payoffs that would have been generated if they had kept their tokens. Also, in these two treatments, earnings are similar to the Bayes-Nash payoffs. Table 4.10 shows the differences in net group earnings as well as the significance of these differences from a Wilcoxon signed-rank paired difference test. Subjects do better than the Bayes-Nash payoffs in the 'binary communication' and 'unrestricted chat' sessions, at 5% and 10% levels of significance, respectively. Of course, part of the reason is that in this situation with $c_{max} = 150$, the Bayes-Nash equilibrium yields no or marginal gains in earnings compared to the equilibrium where nobody contributes. However, in all four treatments, payoffs are significantly (at 1%) lower than those which could have been attained under the best IC mechanism and hence, lower than the earnings under the first-best situation subject to the individual rationality constraint. But, in the 'unrestricted chat' sessions, the difference in earnings between what is observed in data and the best IC mechanism is much smaller than the difference in other three treatments.

Comparison	No Comm.	Binary Comm.	Token Revelation	Unrestricted Chat
Data minus Bayes-Nash	26.0	36.9**	11.2	33.7*
Data minus Best IC Mechanism	-67.7***	-58.8***	-64.7***	-26.2***
Data minus First-Best st. IR	-113.1***	-93.0***	-109.1***	-66.5***

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 4.10: Differences in average net group earnings: $c_{max} = 150$ sessions.

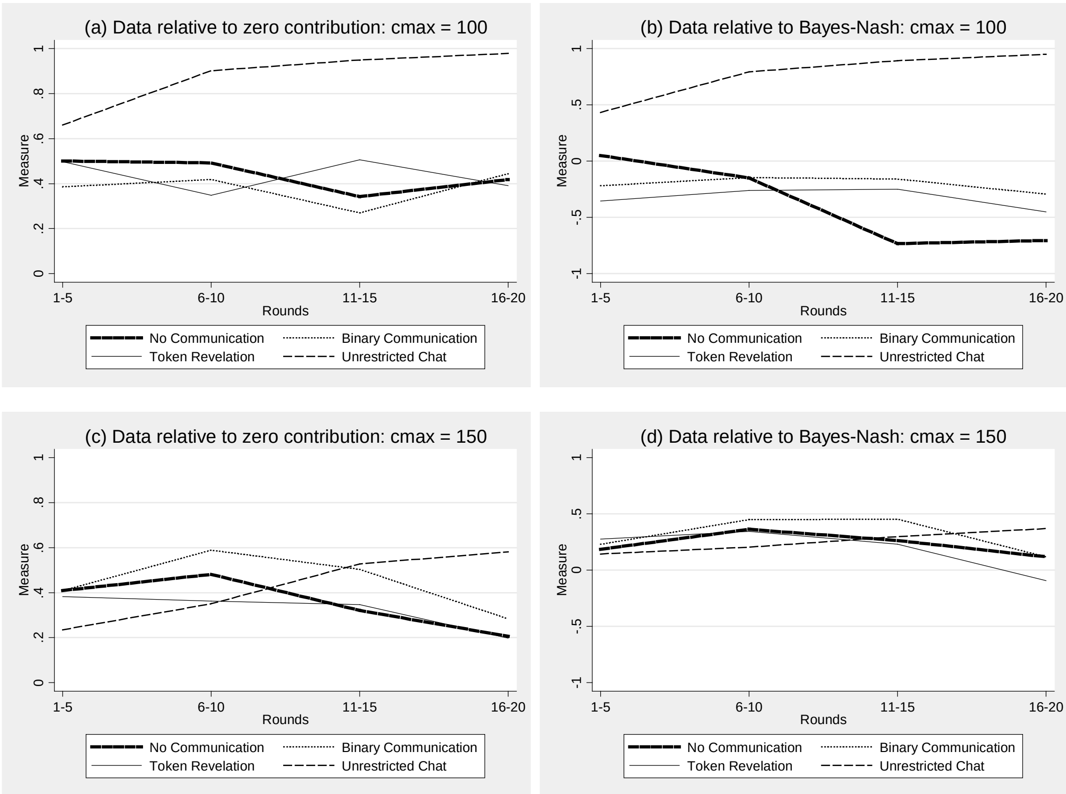

Figure 4.5: Performance of data over rounds.

As we have seen earlier, there are differences in performance during first and second halves of the sessions, so the question arises as to whether there is any interesting trend over time in the four different treatments. Figure 4.5 gives the performance of the data with respect to zero contribution as well as the Bayes-Nash equilibrium over the rounds divided into four phases. It graphs the efficiency measures: ((group earnings in data) - (zero contribution earnings)) divided by (first best earnings subject to individual rationality - (zero contribution earnings)) and ((group earnings in data) - (Bayes-Nash earnings)) divided by (first best earnings subject to individual rationality - (Bayes-Nash earnings)) for rounds 1-5, 6-10, 11-15, and 16-20. It is evident from the graphs that only in the 'unrestricted chat' sessions is there a clear upward trend in efficiency as we move from rounds 1-5 to 16-20.

In fact, subjects reach very close to the first-best outcome in the $c_{max} = 100$ sessions. In the other two communication treatments and in the 'no communication' sessions, efficiency either goes down or stays the same over rounds.

4.4.3 Strategies in Communication Stage

In this section we provide insights into the strategies used by subjects in the communication stage which drives the results we already discussed. These are presented below, separately for each of the communication structures implemented. We use data from all rounds in this section. The same observations hold if we only use the data from the last 10 rounds.

4.4.3.1 Binary Communication

When c_{max} was 100, an individual contributed 56% of the time after she said that she intended to contribute in the message stage, while this number was 55% for $c_{max} = 150$ sessions. 29% of the time a person contributed after sending the message "I intend to keep my token" in the $c_{max} = 100$ sessions. This number was 13% for $c_{max} = 150$ sessions. More interestingly, the frequency of contributions among individuals who used the message "I intend to spend my token" was higher when there were exactly two people using that message than when all of the three members of a group used that message or the individual was the only person to use the message[12]. This is evident from Tables 4.11 and 4.12. In the $c_{max} = 100$ sessions, out of 112 observations with "I intend to spend my token" with exactly two people using this message, 77 of them ended up contributing, whereas, only 20 of 48 contributed when there was only one person using that message and 74 out of 144 contributed when all of the members said that they intended to spend. For $c_{max} = 150$

[12]This result is in line with the one obtained in Palfrey and Rosenthal (1991) in their 'binary communication' experiments with $c_{max} = 150$.

Message sent by me	Number of others saying "I intend to spend" 0	1	2
"I intend to spend"	20 (48)	77 (112)	74 (144)
"I intend to keep"	5 (24)	30 (96)	15 (56)

Table 4.11: Number of times spent: $c_{max} = 100$. Total observations in parentheses.

Message sent by me	Number of others saying "I intend to spend" 0	1	2
"I intend to spend"	11 (38)	87 (103)	72 (112)
"I intend to keep"	2 (5)	9 (38)	7 (24)

Table 4.12: Number of times spent for token values < 100: $c_{max} = 150$. Total observations in parentheses.

sessions, we first consider the subsample where an individual's token value is less than 100. In comparison to the 87 of 103 individuals who said "I intend to spend my token" who contributed when there were exactly two people using that message, 72 out of 112 contributed when three of them used that message and only 11 out of 38 contributed when they were the sole person using that message.

In $c_{max} = 150$ sessions, there were 160 instances when the token values were greater than 100. An individual used the message "I intend to spend my token" in 60 of those cases but contributed only in 4. This can be thought of as 'lying' because from an individual's point of view, she is "never" going to contribute when her token value is higher than 100. So, she is certainly lying about her intention of making a contribution in the decision round. Subjects possibly use this strategy in order to induce others to contribute by signaling that she was willing to share the costs.

4.4.3.2 Token Revelation

Now, in the 'token revelation' treatment, a numerical message was sent around 85% of the time in $c_{max} = 100$ sessions and 92% of the time when c_{max} was 150. In $c_{max} = 100$

sessions, 48% of the reports were truthful while 40% of the messages were lower than the true token values and individuals over-reported their token values around 12% of the time. A similar pattern was seen in the $c_{max} = 150$ sessions, where 43% of the time subjects reported their token values truthfully, 38% of the time there was under-reporting and 19% of the time subjects over-reported. Again, from the prevalence of under-reporting of token values it seems that a subject is trying to signal their group members that they can expect a contribution from her. Probably they tend to believe somehow that this will drive others to contribute too.

In $c_{max} = 150$ sessions, individuals under-reported 53% of the time while they reported truthfully around 37% of the time when their token values were greater than 100. Out of 136 such instances when their token values were higher than 100, 71 of the time they sent a message which indicated their token values were less than 100. That is indeed one of the primary reasons of why efficiency is so low under the 'token revelation' when c_{max} is 150. In contrast, out of 273 cases when the token value was less than 100, around 13% of the time individuals sent a message greater than 100. They were truthful 47% of the time and over-reported 22% of the time.

4.4.3.3 Unrestricted Text Chat

We do a content analysis of the discussion in the communication round of the 'unrestricted text chat' sessions by coding each message sent by a subject into one of the nine mutually exclusive categories, as enlisted in Table 4.13. The table also contains some verbatim examples of sentences/messages that fall under each category that were used in our experiments. While the instructions did not indicate that the chat was to be in English, subjects communicated in a language that was closer to SMS texts than to ordinary English. Table 4.14 gives the percentage of all messages that fall in these nine categories in all four sessions of the 'unrestricted chat' treatment. While "acknowledgement"/"strategy suggestion"/"revelation of own token value" were used a lot in the $c_{max} = 100$ sessions, there was a predominance of informative discussions and messages related to revealing the intent to spend in the $c_{max} = 150$ sessions. Messages falling in categories "strategy suggestion" and "acknowledgement" were used less in sessions with $c_{max} = 150$ than in the $c_{max} = 100$ sessions. Table 4.14 also shows that individuals used conditional statements or ambiguous messages in very few number of cases in all of the sessions. Lastly, there was also no notable difference in the messages sent over time.

Code category	Examples
Acknowledgement	okay; cool; yes; alright; done; great; yep
Irrelevant/Junk	lol; hehehe
Own token value	token value is 39; I have a really high value; its so low
Others' token value or plan of action	what is your token value 3?; are you going to spend 2?
Strategy suggestion about others/own decision/group decision	I should keep my token; 1 and 2 should spend and 3 keep; everyone should spend; can I keep?
Informative/explaining something to group members but not any strategy suggestion	if 2 spend then they both get 100; token values can never be higher than 100
I plan to spend	spending; I will spend
I plan to keep	keeping; I will keep
Conditional statement or ambiguous/contradictory statement	I will spend if someone else spends; I will keep if you two spend; "I will keep" then later on says "I will spend"; I will spend or keep

Table 4.13: Content analysis: code categories.

Thus, we see that revealing one's own token value coupled with a strategy suggestion about one's own as well as the entire group's strategies helps in achieving "efficient" outcomes in the situation when it is common knowledge that everyone has costs less than the 'public' benefit. However, in the other situation, there is a lack of use of these message categories and efficiency is also not as high. Apart from this, a closer look at the transcript of the text chat reveals that in the $c_{max} = 150$ sessions there is a lot of discussion regarding lying, trust and promises whereas, none of these terms are used when $c_{max} = 100$. So, in the $c_{max} = 150$ sessions, there is an atmosphere of mistrust which is certainly a reason as to why subjects do not do as well in these sessions even with unrestricted text chat. A quote from one of the texts from a subject aptly summarizes the situation: "this test really shows how low humans have fallen".

Code category	$c_{max} = 100$ session 1	$c_{max} = 100$ session 2	$c_{max} = 150$ session 1	$c_{max} = 150$ session 2
Acknowledgement	9.3	24.1	6.3	5.5
Irrelevant/Junk	22.5	2.4	20.8	12.2
Own token value	31.4	17.3	8.7	15.2
Others' token value or plan of action	3.2	8.4	10.7	11.3
Strategy suggestion about others/own decision/group decision	12.5	19.7	5.6	6.9
Informative/explaining something to group members but not any strategy suggestion	5.7	8.0	13.9	22.5
I plan to spend	9.8	13.8	21.9	14.5
I plan to keep	4.6	4.2	7.4	9.2
Conditional statement or ambiguous/contradictory statement	0.9	2.1	4.6	2.5

Table 4.14: Content analysis: percentage of messages falling in the code categories.

4.5 Conclusion

We investigated the effect on efficiency, in a public good game where player endowments are private information, of non-binding pre-play communication structures differing in the richness of the message space. Neither the one-time simultaneous exchange of binary messages meant to reveal one's intention to contribute nor the one-time display of a numerical message is enough to create any efficiency gain relative to the situation without any form of communication. Only when participants are provided with the opportunity to engage in text chat in an unrestricted fashion (though anonymous), earnings and public good provision are significantly higher. Thus, it seems that without a 'common language' there is no 'obvious' way to interpret the binary or the numerical messages among subjects. The natural language results hold even when the groups are rotated every round, as is the case in our experiments. Unrestricted chat gives the subjects the opportunity to understand and interpret each others' intentions and messages. Also, gains relative to the situation with no communication are higher when it is common knowledge that everyone has costs less than the 'public' benefit. This is intuitive because in this situation the question is not whether the good should be provided but rather which of the two people in a group are going to provide it. In contrast, in the sessions where it is possible for individuals to have higher costs than the 'public' benefit, first one needs to figure out whether a good should be provided and hence, there are more chances of mis-coordination and mis-representation.

Appendices

Appendix A

Experimental Instructions for Chapter 2

The following is the instructions from one of the sessions in the "private information" treatment. The only difference to the other sessions under this treatment was in the sequence of the parametric configurations. The instructions in the "public information" treatment closely follows the "private information" treatment differing in the informational aspects wherever necessary. The instructions in the sessions with communication also closely follows the "private information" treatment except for additional instructions on the rules of communication provided in section A.2.

A.1 Instructions

Thank you for agreeing to participate in this decision-making experiment. During the experiment we require your complete, undistracted attention, and ask that you follow instructions carefully. You should not open other applications on your computer, chat with other students, or engage in other distracting activities, such as using your phone, reading books, etc.

You will be paid for your participation, in cash, at the end of the experiment. Different participants may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

The entire experiment will take place through computer terminals, and all interactions between you will take place through the computers. It is important that you not talk or in any way try to communicate with other participants during the experiment.

During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, raise your hand and your question will be answered out loud so everyone can hear. If you have any questions after the experiment has begun, raise your hand, and I will come and assist you.

The experiment will have two practice matches followed by the paid session. The paid session will consist of four separate parts, each consisting of four matches. In each match, you will be matched with one of the other participants in the room. In each match, both you and the participant you are matched with will make some decisions. Your earnings for that match will depend on both of your decisions, but are completely unaffected by decisions made by any of the other participants in the room. I will explain exactly how these payoffs are computed in a minute.

At the end of the session, you will be paid the sum of what you have earned in each of the 16 matches, plus the show-up fee of $ 5. Everyone will be paid in private and you are under no obligation to tell others how much you earned. Your earnings during the experiment are denominated in points. At the end of the experiment you will be paid $1 for every 160 points you have earned.

[Description of part I of the experiment]

A match begins by matching you with another person from the room. Both you and the person you are matched with will receive opportunities in real time to undertake an investment. Every time you invest, you incur a cost of investment equal to 5 points, but the person you are matched with gets a return of 10 points. You do not get any return on your own investment. Similarly, every time the person you are matched with makes an investment, he/she incurs a cost of 5 points but you earn a return of 10 points. Again, the person you are matched with does not get any return on his/her own investment.

Throughout a particular match, you will have two options: INVEST and DO NOT INVEST. At time zero, before a match begins in real time, each of you need to activate either the INVEST or the DO NOT INVEST option by clicking on the respective button. Once everyone in the room activates either of the two options, the match begins in real time. When you have the INVEST option activated, then you automatically keep investing whenever an opportunity arrives. Whereas, if you have the DO NOT INVEST option activated, then you automatically forego the opportunity to invest whenever such an opportunity arrives. You may switch between these two options as many times as you wish during a match.

The investment opportunities that you receive are observed only by you and not by the person you are matched with. The person you are matched with only gets to see when you actually invest. Similarly, the opportunities received by the person you are matched with are seen by him or her only. You only get to see the investments made by the person you are matched with.

After a match is over, you will be randomly re-matched to another person from the room and we will proceed to the next match. I will now explain in more detail the frequency of investment opportunities and how the length of a match is determined.

After each second, the computer generates a random integer for each pair of matched participants. This integer is drawn from a uniform distribution over [1,100]. If the integer is in the interval [1,10] then you get an investment opportunity; if the integer is in the interval [11,20] then the person you are matched with gets an investment opportunity; whereas if the integer is in the interval [21,100], then neither receives an investment opportunity. So on an average, every 10 seconds, you will get one investment opportunity and the person you are matched with will get one investment opportunity. Of course these are just averages. This random number generation is independent across each second.

Also, after each second and after a random integer has been generated for the arrival of opportunities, another random integer is instantly generated independent of the previous process and common to all pairs of matched participants. This decides whether a match is finished or whether we continue. This integer is independently drawn from a uniform distribution over [1,100]. If the outcome is 100, we stop. Otherwise we continue the match. This is again done independently after each second. Thus, there is a 1% chance that the match ends after every second. Therefore, regardless of how much time has already elapsed, the match is still expected to last another 100 seconds. Are there any questions about how this works?

[Explain the screenshots before the practice match begins in real time and also go through the screen display in detail]

We will now go through the first practice match. This practice match will last for a definite length of 180 seconds. This is only for illustrative purposes. In the paid matches, the length of each match will be random and there will be a 1 percent chance that the match ends after each second, as explained earlier. During the practice match, please do not hit any keys until I tell you, and when you are prompted by the computer to enter information, please wait for me to tell you exactly what to enter. You are not paid for this practice match.

[AUTHENTICATE CLIENTS]

Please double click on the icon on your desktop that says MULTISTAGE CLIENT. When the computer prompts you for your name, type your First and Last name. Then click SUBMIT and wait for further instructions.

[START GAME]

You now see the first screen of the experiment on your computer. It should look similar to the screen in front of the room. Each of you please activate either the INVEST or the DO NOT INVEST option. Remember that you can always switch back and forth between these two options during the match. Once everyone in the room activates one of the two options, the match begins in real time. Familiarize yourself with the screen display and the choice buttons (for INVEST and DO NOT INVEST). You will not be paid for this practice match.

[End of practice match]

Now we will go through the second practice match. The length of this practice match is random and there will be a 1 percent chance that the match ends after each second, as explained earlier.

[End of second practice match]

Are there any questions before we begin with the paid session?

[WAIT FOR QUESTIONS]

We will now begin with the first of the four paid matches of part one of the experiment. Please pull out your dividers for the paid session of the experiment. If there are any problems or questions from this point on, raise your hand and I will come and assist you.

[display summary screen for matches 2-5 and go over it quickly]

[Run matches 2-5]

This completes part 1 of the experiment. Part 2 will also consist of four matches. The rules are identical to part 1, except that the frequency of random investment opportunities will now be different. As before, the computer will generate a random integer from a uniform distribution over [1,100] for each pair of matched participants after each second. But now, if the integer is in the interval [1,30], then you get an investment opportunity, and if the integer is in the interval [31,60] then the person you are matched with gets an investment opportunity, whereas if the integer is in the interval [61,100], then nobody receives an investment opportunity. So on an average, every 10 seconds, you will get about three investment opportunities and the person you are matched with will get three investment opportunities. Of course, these are just averages. This random number generation is again independent across each second.

[display summary screen for matches 6-9 and go over it quickly]

[Run matches 6-9]

This completes part 2 of the experiment. Part 3 will also consist of four matches. The rules are identical to part 2, except that the return to the person you are matched with is now 25 instead of 10 when you invest.

[display summary screen for matches 10-13 and go over it quickly]

[Run matches 10-13]

This completes part 3 of the experiment. Part 4 will also consist of four matches. The rules are identical to part 3, except that the nature of random investment opportunities will now be different. As before, the computer will generate a random integer from a uniform distribution over [1,100] for each pair of matched participants after each second. But, now if the integer is in the interval [1,10] then you get an investment opportunity and if the integer is in the interval [11,20] then the person you are matched with gets an investment opportunity, whereas if the integer is in the interval [21,100], then nobody receives an investment opportunity. So on an average, every 10 seconds, you will get one investment opportunity and the person you are matched with will get one investment opportunity. Of course these are just averages. This random number generation is again independent across each second.

[display summary screen for matches 14-17 and go over it quickly]

[Run matches 14-17]

[AFTER 17th MATCH, READ THE FOLLOWING]

This completes the experiment. A popup window that gives your total earnings has appeared on your screen. Please record this amount on your record sheet, rounding up to the nearest 25 cents. You will be paid this amount plus your show-up fee of $5.00. Please complete all the information on your record sheet and wait until your ID is called to be paid privately in the next room.

Please remain in your seat while you are waiting. Do not talk or use the computers. Please take all belongings with you when you leave to receive payment. You are under no obligation to reveal your earnings to the other participants. Thank you for your participation.

A.2 Additional Instructions in the Communication Treatments

At the very beginning of every match you will have a 60 seconds discussion period, during which you are allowed to send messages to the person you are matched with. The messages you send are seen by both you and the person you are matched with. The messages must conform to the following rules. 1. Your messages must be relevant to the experiment. Do not engage in social chat or use emoticons. 2. You are not permitted to send messages intended to reveal your identity. 3. Do not use threatening or offensive language.

Appendix B

Additional Analysis for Chapter 2

When the opportunities to provide a favor are publicly observed by both players, then players can condition their behavior on whether or not their partner provided them a favor at the last received opportunity. Table B.1 gives the mean rates of favor provision depending on whether a partner provided or declined to provide a favor at the last received opportunity, separately for the situation when the individual is ahead or behind/tied. Clearly, *the likelihood of favor provision by an individual is higher if her partner provided a favor at the last opportunity received compared to the situation where her partner declined to grant a favor at the last opportunity received.* Also, this likelihood is higher if the individual is behind/tied compared to the situation where she has done more favors than her partner.

Favor Provided/Declined	Condition	Low-Slow	Low-Fast	High-Slow	High-Fast
provided	ahead	85.18	86.45	87.43	92.53
provided	behind/tied	90.45	90.55	97.22	93.08
declined	ahead	29.27	43.54	42.12	41.85
declined	behind/tied	37.70	57.62	58.61	51.77

Table B.1: Mean favor provision by whether partner provided favor at the last received opportunity

A pooled probit regression with standard errors clustered at the subject level is run for each treatment, separately depending on whether or not the participant is ahead, or behind and tied. The probit equations are similar to the ones under private information treatments, except that $nice_{it-1}$ is added to the set of exogenous variables[1], where $nice_{it-1}$ is a binary variable taking value 0 (1) if the partner provided the favor (declined to grant a favor) at the last received opportunity[2].

variable	Low-Slow	Low-Fast	High-Slow	High-Fast
mode_{t-1}	0.79	0.56	0.55	0.55
	(0.04)***	(0.04)***	(0.06)***	(0.03)***
nice_{t-1}	-0.06	-0.02	-0.04	-0.03
	(0.02)***	(0.01)*	(0.01)***	(0.01)***
netfavors_{t-1}	-0.01	-0.002	-0.0007	-0.003
	(0.005)***	(0.001)	(0.0003)**	(0.001)**
timeelapsed_{t-1}	-0.0005	-0.001	0.000003	-0.0004
	(0.0002)*	(0.0005)***	(0.00002)	(0.0004)
ownpayoff_{t-1}	-0.0002	-0.00005	-0.000004	-0.00001
	(0.0002)	(0.00005)	(0.00000)	(0.00001)
no. of observations	8187	5230	8426	7587
no. of groups	28	28	28	28
pseudo-R^2	0.72	0.46	0.75	0.41

* significant at 10%; ** significant at 5%; *** significant at 1%

Pooled probit regression with clustering at the subject's level; Standard errors in parentheses. The unit of observation is a subject per second; Dependent variable: mode_t

Table B.2: Marginal effects for "Public Information": behind and tied

Tables B.2 and B.3 collect the marginal effects of the probit regression. As before, the choice of mode at time t continues to be affected tremendously by the choice at $t-1$. The variables 'net favors', 'time since last favor by partner', and 'nice' have significant negative coefficients in most of the cases. Thus, it is clear that when the information is available on whether or not an individual's partner has been 'nice' or 'mean', individuals pay attention to this new piece of information in addition to the other 'obvious' state variables. Rewards

[1] Including 'time' as a covariate does not affect the analysis.

[2] This variable takes a value 0 by default if the partner has not yet received her first opportunity. So, until first refusal of favor, everyone is 'nice'.

and punishments are conditioned on this new and simple piece of information. However, importantly, it can be noted from the magnitudes of the marginal effects that the likelihood of favor provision is less responsive to 'net favors' and 'time since last favor by partner' under public information than under private information for each treatment. Finally, the effect of 'own payoff' is still inconclusive.

variable	Low-Slow	Low-Fast	High-Slow	High-Fast
$mode_{t-1}$	0.91	0.80	0.81	0.66
	(0.01)***	(0.04)***	(0.04)***	(0.05)***
$nice_{t-1}$	-0.13	-0.03	0.004	-0.09
	(0.03)***	(0.02)*	(0.005)	(0.02)***
$netfavors_{t-1}$	-0.007	-0.007	-0.007	-0.004
	(0.007)	(0.002)***	(0.002)***	(0.001)***
$timeelapsed_{t-1}$	-0.0009	-0.006	-0.0005	-0.0004
	(0.0005)*	(0.002)***	(0.0002)**	(0.0006)
$ownpayoff_{t-1}$	0.002	0.0002	0.000009	0.00003
	(0.0005)***	(0.00009)***	(0.00002)	(0.00001)**
no. of observations	5383	4140	5994	6293
no. of groups	27	27	28	28
pseudo-R^2	0.73	0.64	0.76	0.53

* significant at 10%; ** significant at 5%; *** significant at 1%

Pooled probit regression with clustering at the subject's level; Standard errors in parentheses. The unit of observation is a subject per second; Dependent variable: $mode_t$

Table B.3: Marginal effects for "Public Information": ahead

Appendix C

Experimental Instructions for Chapter 3

C.1 Instructions

Thank you for agreeing to participate in this decision-making experiment. During the experiment we require your complete, undistracted attention, and ask that you follow instructions carefully. You should not open other applications on your computer, chat with other students, or engage in other distracting activities, such as using your phone, reading books, etc.

You will be paid for your participation, in cash, at the end of the experiment. Different participants may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

The entire experiment will take place through computer terminals, and all interactions between you will take place through the computers. It is important that you not talk or in any way try to communicate with other participants during the experiment.

During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, raise your hand and your question will be answered out loud so everyone can hear.

If you have any questions after the experiment has begun, raise your hand, and I will come and assist you.

The experiment will consist of 11 matches. In each match, you will be matched with one of the other participants in the room. In each match, both you and the participant you are matched with will make some decisions. Your earnings for that match will depend on both of your decisions, but are completely unaffected by decisions made by any of the other participants in the room.

At the end of the experiment, you will be paid the sum of what you have earned in each of the 11 matches, plus the show-up fee of $ 5. Everyone will be paid in private and you are under no obligation to tell others how much you earned. Your earnings during the experiment are denominated in points. At the end of the experiment you will be paid $1 for every 150 points you have earned.

The Structure of a Match

Every match proceeds as follows. You are matched with another person from the room at the beginning of a match. Both you and the other person are sellers in a market producing and selling the same product. **Your decision is to choose the quantity of the product to be produced and sold in the market. You can choose any number (up to one decimal place) between (and including) 0 and 50.**

It costs 2 points to produce an unit of the product. Thus, if you produce q units, your total cost is $2q$ points.

The total market quantity (Q) is the sum of the quantity chosen by you and the quantity chosen by the other person. The price of the product depends on the total market quantity according to the following rule. Whenever the total market quantity is less than or equal to 50, then the price of the product is 50 minus the total market quantity. However, if the total market quantity exceeds 50, then the price of the product is zero. That is,

$$P = \begin{cases} 50 - Q & Q \leq 50 \\ 0 & Q > 50 \end{cases} .$$

Your revenue is given by price times the quantity you produce, costs are two times the quantity you produce, and, your profit equals your revenue minus your cost:

Profit = Revenue - Cost = $(P - 2)*$(quantity produced by you)

In order to make an informed decision, each of you has been provided with a profit sheet and your screen contains a profit calculator. The profit sheet shows the profit that you will earn for every possible combination of integer choices (for 0 to 30). Each row in the profit sheet refers to your quantity, and each column refers to the quantity produced by the person you are matched with. For any combination of (integer) quantities chosen by you and the person you are matched with, you find the corresponding cell to find the resulting profits. If you wish to find out the resulting profits when you enter integer numbers above 30 or non-integer numbers for your quantity and other person's quantity, then use the profit calculator. Here you can enter any number up to one decimal place between 0 and 50.

[The following paragraph only in the 'Baseline' treatments.] You and the person you are matched with will have 120 seconds to decide the quantities to be produced. At the beginning of a match, each of you should enter an initial tentative quantity. After you enter an initial tentative quantity, you will get to see your tentative quantity on your screen. However, **you will not be able to observe** the tentative quantity chosen by the person you are matched with, your tentative profits, or the other person's tentative profits. Once everyone in the room has entered a tentative quantity, the "120 seconds" timer starts. Remaining time will be displayed on top of your screen. At any point, you can change your tentative quantity (during the 120 seconds). Every time you make a change it is updated on your screen. The word *tentative* reflects the fact that if there are no more changes then the last entry shows your final quantity choice for the match. Your final quantity choice for the match is whatever your tentative quantity is at the end of 120 seconds. At the end of the 120 seconds you will get to observe not only your own final quantity choice but also the final quantity choice made by the person you are matched with. Depending on these quantity choices, you will then get to know the profits earned by you and the profits earned by the person you are matched with for the match.

[The following paragraph only in the 'Real Time Revision with Monitoring' treatments.] You and the person you are matched with will have 120 seconds to decide the quantities to be produced. At the beginning of a match, each of you should enter an initial tentative quantity. After both of you enter an initial tentative quantity, you will get to see your tentative quantity, the tentative quantity chosen by the person you are matched with, your tentative profits and the other person's tentative profits on both of your screens. Once everyone in the room has entered a tentative quantity, the "120 seconds" timer starts. Remaining time will be displayed on top of your screen. At any point, you can change your tentative quantity (during the 120 seconds). Every time you or the person you are matched with makes a change it is updated on the screen for both of you to see along with the tentative profits. The word *tentative* reflects the fact that if there are no more changes then the last entry shows the final quantity choice and the amount of profits you earn for this match. Your final quantity choice for the match is whatever your tentative quantity is at the end of 120 seconds. Your earnings for this match will be the tentative profits at the end of 120 seconds.

[The following two paragraphs only in the 'Poisson Revision' treatments.] You and the person you are matched with will have 120 seconds to decide the quantities to be produced. At the beginning of a match, each of you should enter an initial tentative quantity. After both of you enter an initial tentative quantity, you will get to see your tentative quantity, the tentative quantity chosen by the person you are matched with, your tentative profits and the other person's tentative profits on both of your screens. Once everyone in the room has entered a tentative quantity, the "120 seconds" timer starts. Remaining time will be displayed on top of your screen. At any point, you can change your tentative quantity (during the 120 seconds). Every time you or the person you are matched with makes a change it is updated on the screen for both of you to see along with the tentative profits.

The word *tentative* suggests that the quantity choice is not the final choice for the match. Your final quantity choice for the match is determined according to the following procedure: After each second, the computer generates a random integer for each pair of matched participants. This integer is drawn from a uniform distribution over [1,50]. If the integer drawn is 1, then an "event" occurs. This means that every second the probability that an "event" happens is 0.02 (as every integer is equally likely to be drawn). So on an average, an "event" takes place every 50 seconds[1]. Of course this is just an average; this random number generation is independent across each second. **Your final quantity choice for the match is given by the tentative quantity corresponding to the last time an "event" happened before the "120 seconds time" ends. Your earnings**

[1] For the Poisson High treatment, the following was presented: This integer is drawn from a uniform distribution over [1,25]. If the integer drawn is 1, then an "event" occurs. This means that every second the probability that an "event" happens is 0.04 (as every integer is equally likely to be drawn). So on an average, an "event" takes place every 25 seconds.

for this match will be the tentative profits corresponding to the last time an "event" happened before the "120 seconds time" ends. Also note that if an "event" does not happen in the entire "120 seconds time", then the tentative quantity choices at time 1 (initial choices) will be implemented as the final quantity choices. So, to summarize, whether or not you change your tentative quantity is always in your "hands" but whether its implemented for the match is dependent on the random occurrence of the "event".

After a match is over we will go to the next match. **You will be matched to a new participant from the room. It is very important to note that you will not be matched with the same person ever again.** Every match proceeds according to exactly the same rules as described above. When the 11^{th} match is over, you will be paid the sum of what you have earned in all matches plus the show up fee.

Before we begin with the experiment, it is important that you understand the range of profits that you can earn depending on your quantity choice and the other person's quantity choice. For this reason, you will have two minutes before the start of the experiment to explore different quantity combinations with the profit sheet and the profit calculator. During the actual experiment, you are free to consult both the profit sheet and the profit calculator whenever you wish.

C.2 Profit Sheet

The profit sheets that were handed out to the participants are given in Figures C.1 and C.2.

Other →

You ↓

You ↓ \ Other →	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33
2	92	90	88	86	84	82	80	78	76	74	72	70	68	66	64
3	135	132	129	126	123	120	117	114	111	108	105	102	99	96	93
4	176	172	168	164	160	156	152	148	144	140	136	132	128	124	120
5	215	210	205	200	195	190	185	180	175	170	165	160	155	150	145
6	252	246	240	234	228	222	216	210	204	198	192	186	180	174	168
7	287	280	273	266	259	252	245	238	231	224	217	210	203	196	189
8	320	312	304	296	288	280	272	264	256	248	240	232	224	216	208
9	351	342	333	324	315	306	297	288	279	270	261	252	243	234	225
10	380	370	360	350	340	330	320	310	300	290	280	270	260	250	240
11	407	396	385	374	363	352	341	330	319	308	297	286	275	264	253
12	432	420	408	396	384	372	360	348	336	324	312	300	288	276	264
13	455	442	429	416	403	390	377	364	351	338	325	312	299	286	273
14	476	462	448	434	420	406	392	378	364	350	336	322	308	294	280
15	495	480	465	450	435	420	405	390	375	360	345	330	315	300	285
16	512	496	480	464	448	432	416	400	384	368	352	336	320	304	288
17	527	510	493	476	459	442	425	408	391	374	357	340	323	306	289
18	540	522	504	486	468	450	432	414	396	378	360	342	324	306	288
19	551	532	513	494	475	456	437	418	399	380	361	342	323	304	285
20	560	540	520	500	480	460	440	420	400	380	360	340	320	300	280
21	567	546	525	504	483	462	441	420	399	378	357	336	315	294	273
22	572	550	528	506	484	462	440	418	396	374	352	330	308	286	264
23	575	552	529	506	483	460	437	414	391	368	345	322	299	276	253
24	576	552	528	504	480	456	432	408	384	360	336	312	288	264	240
25	575	550	525	500	475	450	425	400	375	350	325	300	275	250	225
26	572	546	520	494	468	442	416	390	364	338	312	286	260	234	208
27	567	540	513	486	459	432	405	378	351	324	297	270	243	216	189
28	560	532	504	476	448	420	392	364	336	308	280	252	224	196	168
29	551	522	493	464	435	406	377	348	319	290	261	232	203	174	145
30	540	510	480	450	420	390	360	330	300	270	240	210	180	150	120

Figure C.1: Profit sheet: page 1

	Other →															
	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17
2	62	60	58	56	54	52	50	48	46	44	42	40	38	36	34	32
3	90	87	84	81	78	75	72	69	66	63	60	57	54	51	48	45
4	116	112	108	104	100	96	92	88	84	80	76	72	68	64	60	56
5	140	135	130	125	120	115	110	105	100	95	90	85	80	75	70	65
6	162	156	150	144	138	132	126	120	114	108	102	96	90	84	78	72
7	182	175	168	161	154	147	140	133	126	119	112	105	98	91	84	77
8	200	192	184	176	168	160	152	144	136	128	120	112	104	96	88	80
9	216	207	198	189	180	171	162	153	144	135	126	117	108	99	90	81
10	230	220	210	200	190	180	170	160	150	140	130	120	110	100	90	80
11	242	231	220	209	198	187	176	165	154	143	132	121	110	99	88	77
12	252	240	228	216	204	192	180	168	156	144	132	120	108	96	84	72
13	260	247	234	221	208	195	182	169	156	143	130	117	104	91	78	65
14	266	252	238	224	210	196	182	168	154	140	126	112	98	84	70	56
15	270	255	240	225	210	195	180	165	150	135	120	105	90	75	60	45
16	272	256	240	224	208	192	176	160	144	128	112	96	80	64	48	32
17	272	255	238	221	204	187	170	153	136	119	102	85	68	51	34	17
18	270	252	234	216	198	180	162	144	126	108	90	72	54	36	18	0
19	266	247	228	209	190	171	152	133	114	95	76	57	38	19	0	-19
20	260	240	220	200	180	160	140	120	100	80	60	40	20	0	-20	-40
21	252	231	210	189	168	147	126	105	84	63	42	21	0	-21	-42	-42
22	242	220	198	176	154	132	110	88	66	44	22	0	-22	-44	-44	-44
23	230	207	184	161	138	115	92	69	46	23	0	-23	-46	-46	-46	-46
24	216	192	168	144	120	96	72	48	24	0	-24	-48	-48	-48	-48	-48
25	200	175	150	125	100	75	50	25	0	-25	-50	-50	-50	-50	-50	-50
26	182	156	130	104	78	52	26	0	-26	-52	-52	-52	-52	-52	-52	-52
27	162	135	108	81	54	27	0	-27	-54	-54	-54	-54	-54	-54	-54	-54
28	140	112	84	56	28	0	-28	-56	-56	-56	-56	-56	-56	-56	-56	-56
29	116	87	58	29	0	-29	-58	-58	-58	-58	-58	-58	-58	-58	-58	-58
30	90	60	30	0	-30	-60	-60	-60	-60	-60	-60	-60	-60	-60	-60	-60

Figure C.2: Profit sheet: page 2

Appendix D

Experimental Instructions and User Interface for Chapter 4

D.1 Instructions

Thank you for agreeing to participate in this group decision making experiment. During the experiment we require your complete, undistracted attention, and ask that you follow instructions carefully. You should not open other applications on your computer, chat with other students, or engage in other distracting activities, such as using your phone, reading books, etc.

You will be paid for your participation in cash, at the end of the experiment. Different participants may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

The entire experiment will take place through computer terminals, and all interaction between you will take place through the computers. It is important that you not talk or in any way try to communicate with other participants during the experiment, except as instructed.

We will start with a brief instruction period. During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, raise your hand and your question will be answered out loud so everyone can hear. If you have any questions after the experiment has begun, raise your hand, and an experimenter will come and assist you.

You will make choices over a sequence of 20 matches. In each match, you will be assigned to a group with two other participants in the room. In every match, you and the two other participants you are matched with each make a single decision. Your earnings for that match will depend on all three group members' decisions, but are completely unaffected by the decisions made by participants assigned to other groups. We will explain exactly how these payoffs are computed in a minute.

At the end of the session, you will be paid the sum of what you have earned in all matches, plus the show-up fee of \$ 5. Everyone will be paid in private and you are under no obligation to tell others how much you earned. Your earnings during the experiment are denominated in points. At the end of the experiment you will be paid \$ 1 for every 100 points you have earned.

Every match proceeds as follows. At the beginning of each match, we randomly divide you into three-member groups called committees. The committees are completely independent of each other and payoffs and decisions in one committee have no effect on payoffs and decisions in other committees. Each member has a single token and can either spend or keep that token. Each member is also assigned a private token value, which is equally likely to be any amount of points between 1 and 100. Token value assignments are completely independent across members, across committees, and across matches. Thus, your own token value tells you absolutely nothing about the token value of the other members, and has no effect on any future token values that will be assigned to you or anyone else.

Payoffs are computed as follows. If you keep your token you earn your token value in that match plus you earn 100 points if both other members of your committee decide to spend their tokens. If you choose to spend your token, then you earn 100 points if at least one other member of your group spends their token, and you earn 0 points if no other member of your group spends their token. This is summarized in the following table.

Your Decision	#Others Spending	Your Earnings
SPEND	0	0
	1	100
	2	100
KEEP	0	YOUR TOKEN VALUE
	1	YOUR TOKEN VALUE
	2	YOUR TOKEN VALUE + 100

[The following paragraph only in the 'No Communication' treatments.] Every match you are prompted to make your choice to either keep or spend your token. When everyone has made a choice, the outcome and the choices of the other members of your committee are revealed, and this determines your earnings for the match.

[The following two paragraphs only in the 'Binary Communication' treatments.] Before anyone makes a spending decision, each of you will have an opportunity to give the other members of your committee some indication of what your spending decision might be. There are exactly two messages which you can send. They are:

MESSAGE A: "I INTEND TO SPEND MY TOKEN"

MESSAGE B: "I INTEND TO KEEP MY TOKEN"

Please remember that these are only messages and are not binding in any way. When the message stage ends, you are told the intent messages of the others in your committee and the decision stage begins wherein you are prompted to make your choice to either keep or spend your token. This decision is binding. When everyone has made a choice, the outcome and the choices of the other members of your committee are revealed, and this determines your earnings for the match.

[The following two paragraphs only in the 'Token Revelation' treatments.] Before anyone makes a spending decision, your committee has a 20 second message stage, during which you are allowed to send a message to the other members of your committee. This message can only be an integer between 1 and 100 and you are allowed to send only one such message. The integer value you send are seen by both other members of your committee. In the situation where you do not send any message, it will be shown as a "question mark" to other members of your committee at the end of the message stage.

When the message stage ends, the decision stage begins wherein you are prompted to make your choice to either keep or spend your token. When everyone has made a choice, the outcome and the choices of the other members of your committee are revealed, and this determines your earnings for the match.

[The following two paragraphs only in the 'Unrestricted Text Chat' treatments.] Before anyone makes a spending decision, your committee has a 60 second discussion period, during which you are allowed to send messages to the other members of your committee. The messages you send are seen by both other members of your committee. The messages must conform to the following rules: (1) Your messages must be relevant to the experiment. Do not engage in social chat or use emoticons. (2) You are not permitted to send messages intended to reveal your identity. (3) The use of threatening or offensive language, including profanity, is not permitted. (4) Do not send blank or nonsense messages.

When the discussion period ends, the decision stage begins wherein you are prompted to make your choice to either keep or spend your token. When everyone has made a choice, the outcome and the choices of the other members of your committee are revealed, and this determines your earnings for the match.

When all committees have finished the first match, we then go to the next match. You will be randomly re-matched into new three-person committees and everyone is independently and randomly assigned a new token value between 1 and 100. Every match proceeds according to exactly the same rules as described above.

D.2 User Interface

Figures D.1-D.7 show the screen-shots for the user interface in the different treatments.

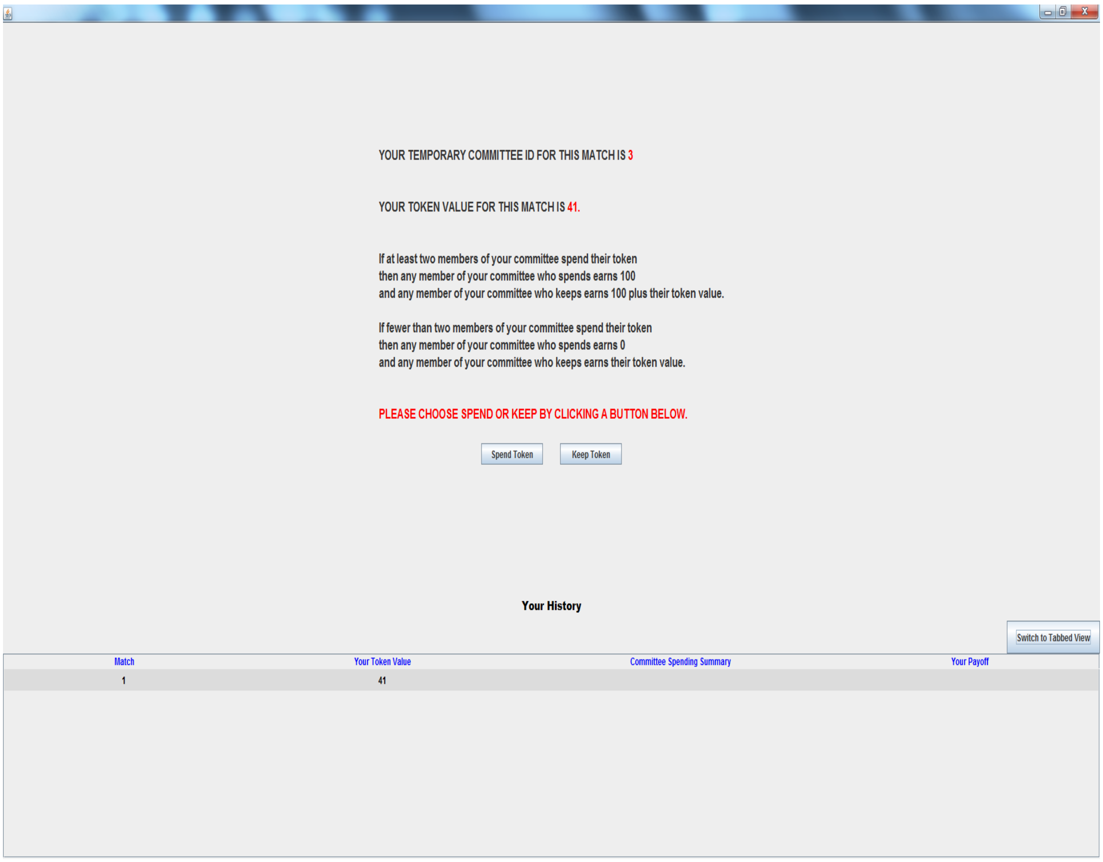

Figure D.1: User interface: “No Communication”

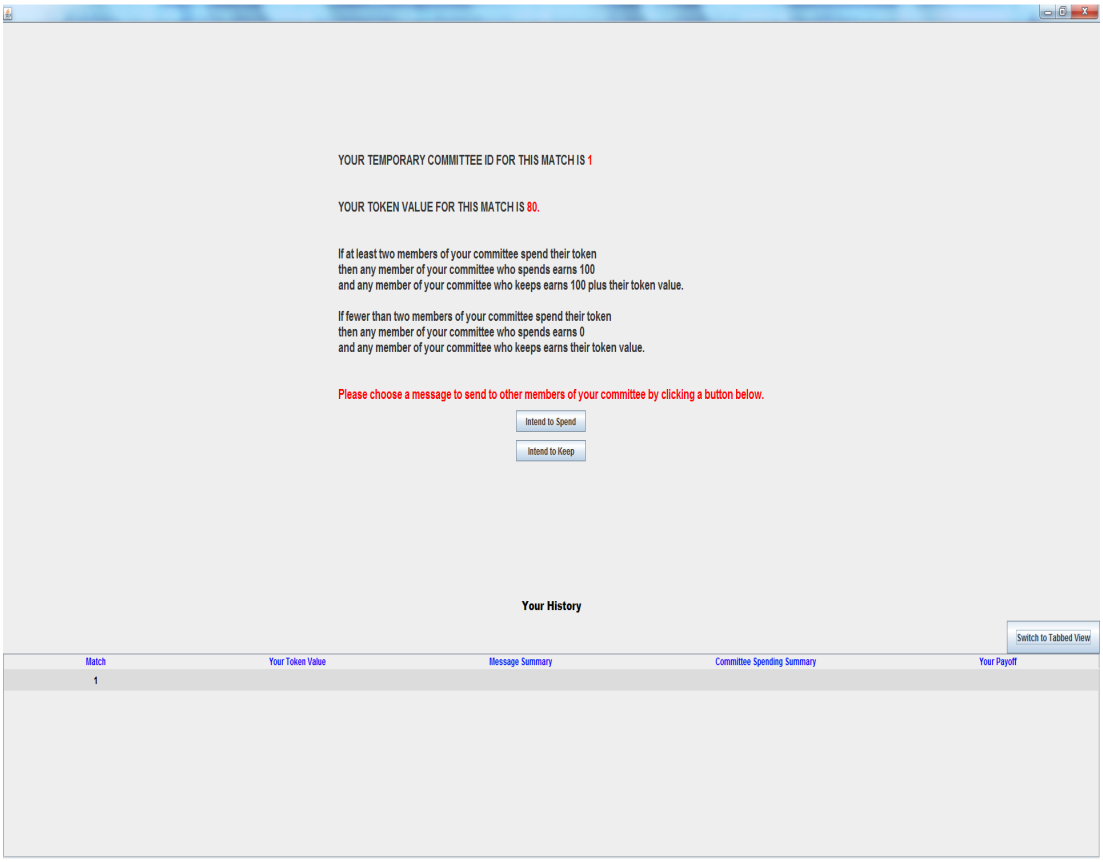

Figure D.2: User interface: “Binary Communication” I

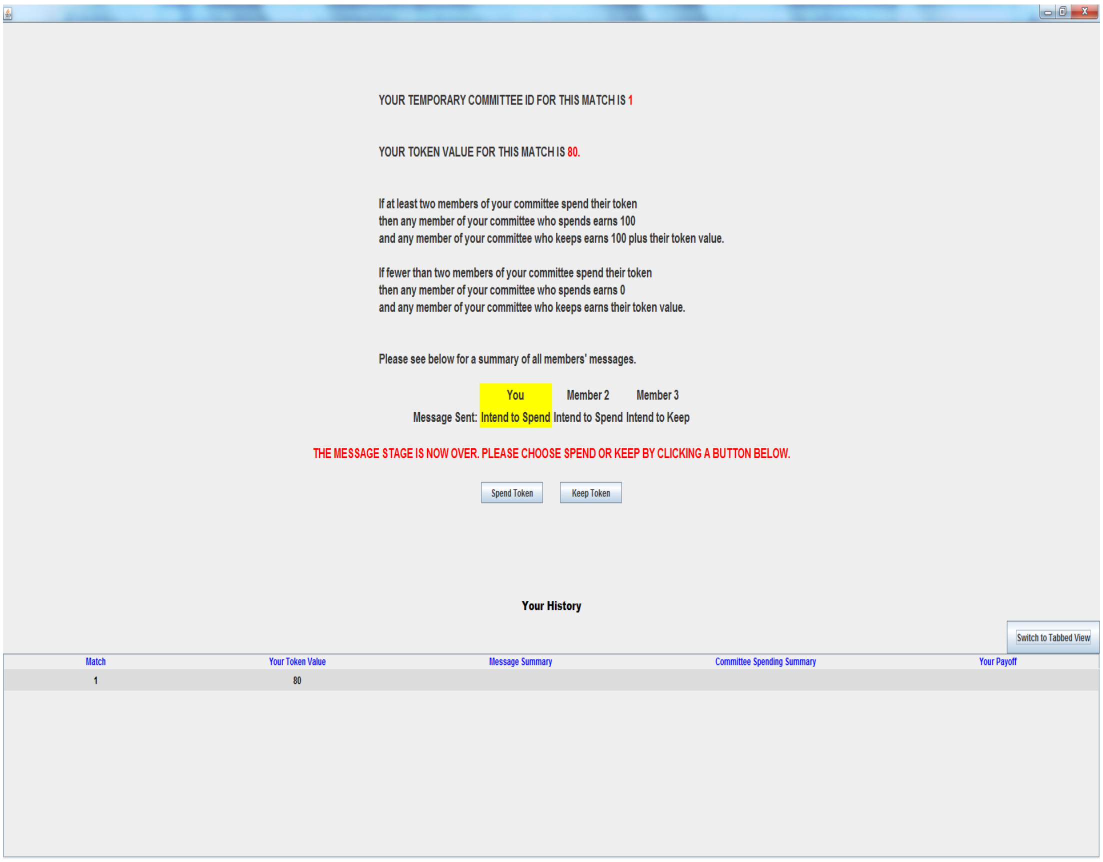

Figure D.3: User interface: “Binary Communication” II

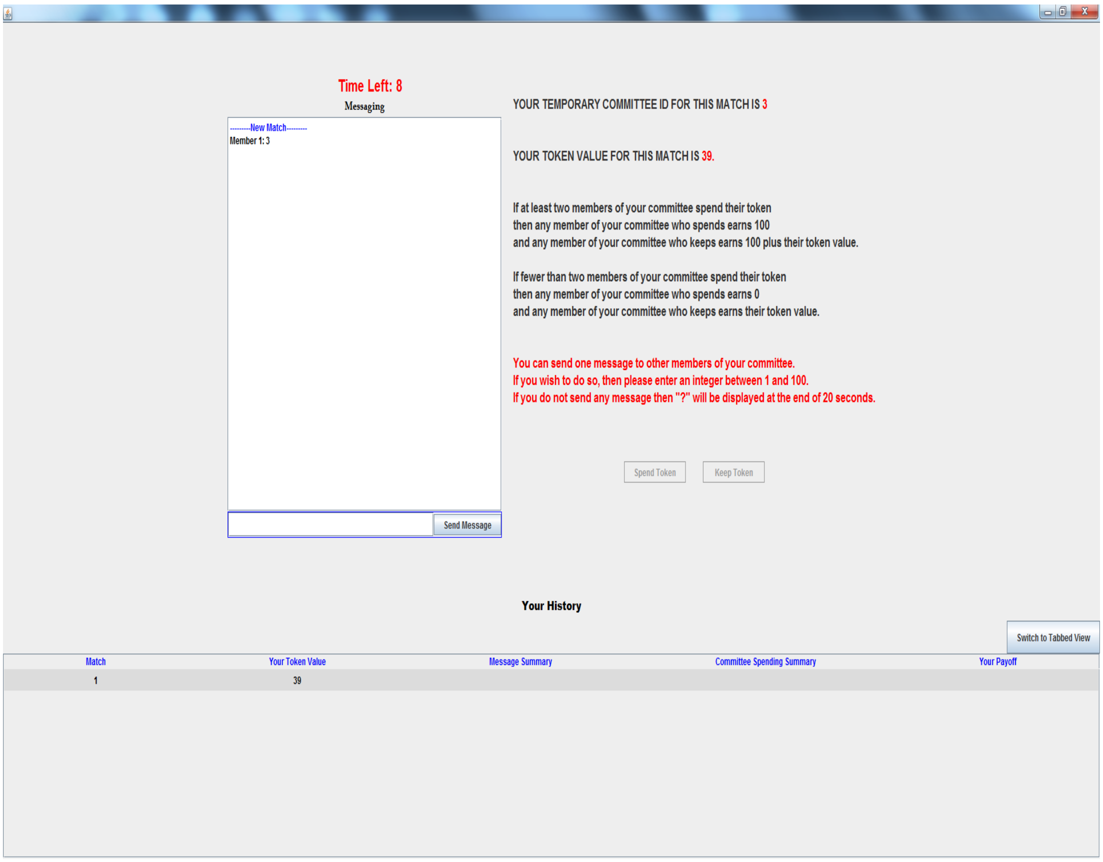

Figure D.4: User interface: “Token Revelation” I

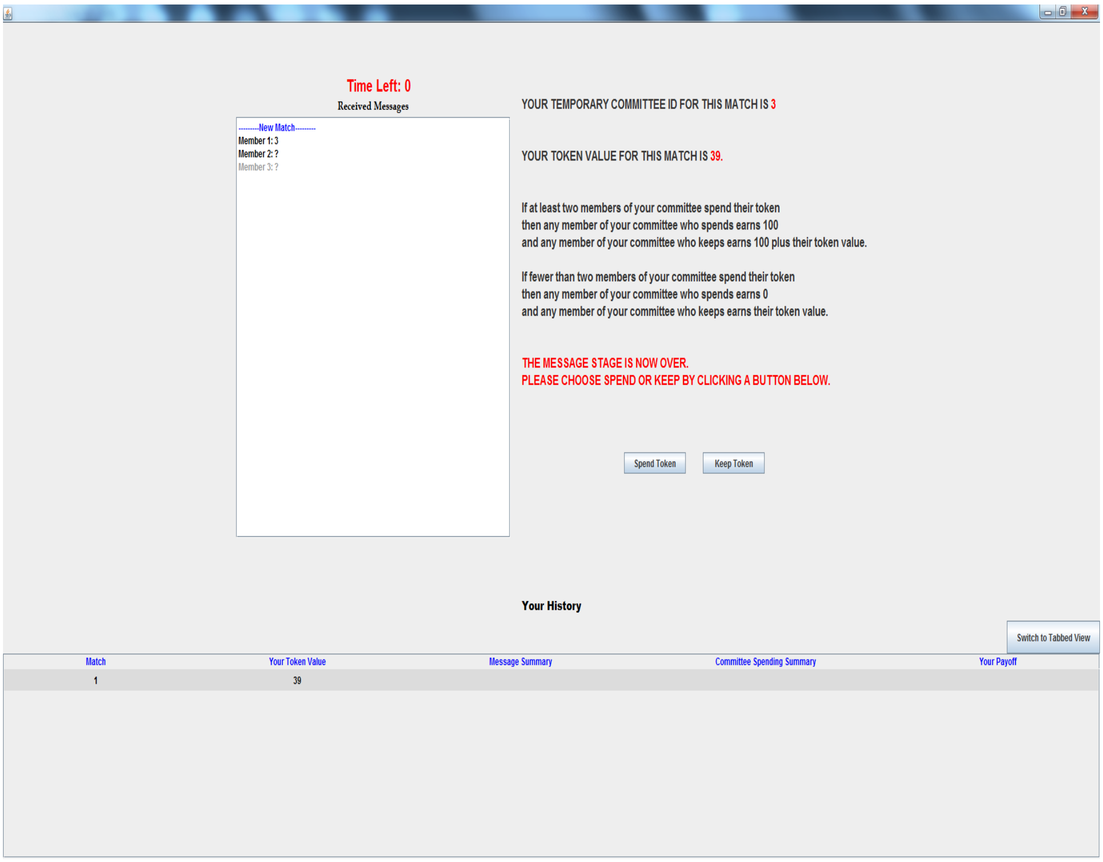

Figure D.5: User interface: "Token Revelation" II

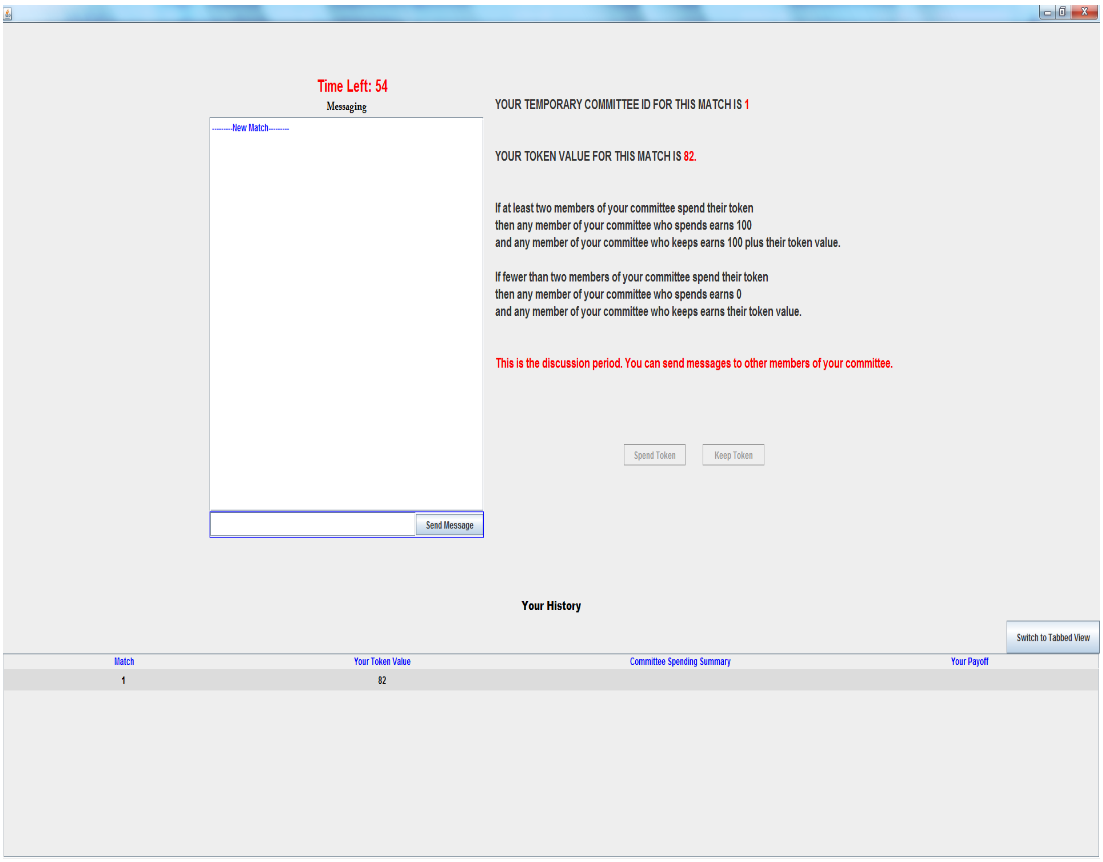

Figure D.6: User interface: “Unrestricted Text Chat” I

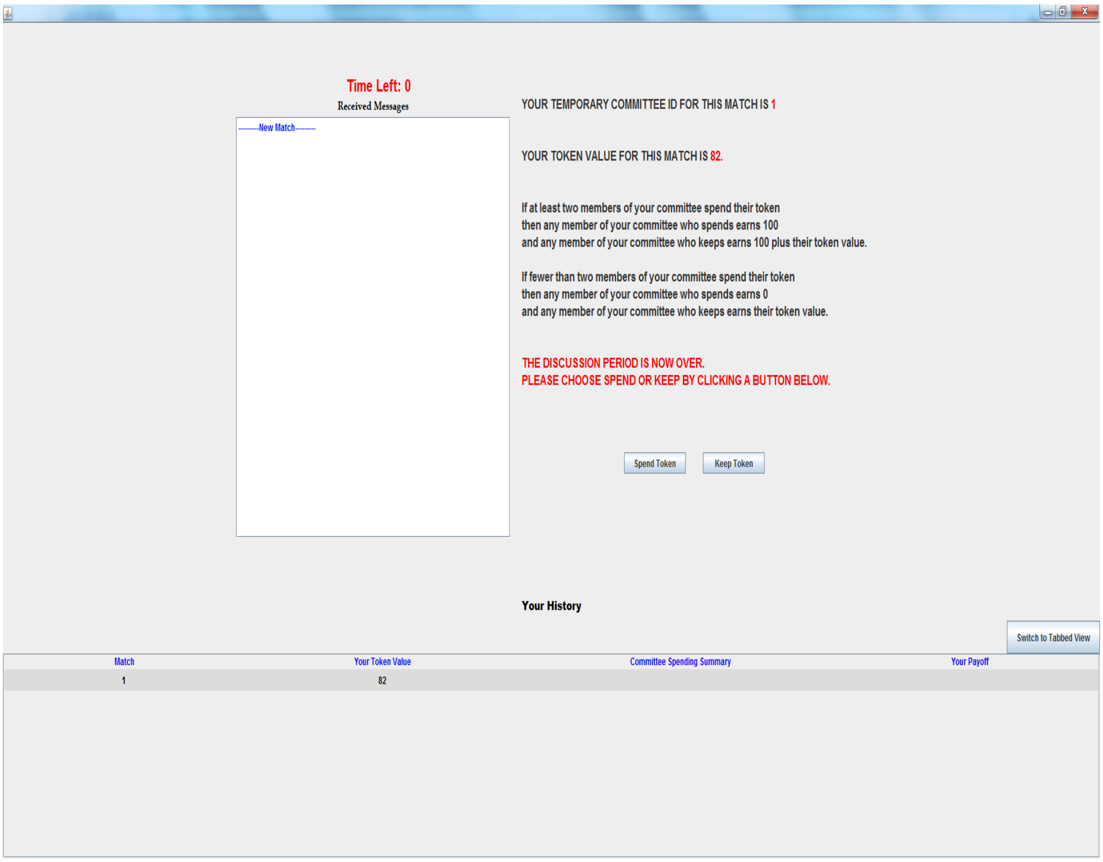

Figure D.7: User interface: “Unrestricted Text Chat” II

Bibliography

[1] Abreu, D., Pearce, D. and Stacchetti, G. (1990), "Towards a theory of discounted repeated games with imperfect monitoring," *Econometrica*, 58(5), 1041-1063.

[2] Abdulkadiroglu, A. and Bagwell, K. (2012a), "The optimal chips mechanism in a model of favors," *Mimeo.*

[3] Abdulkadiroglu, A. and Bagwell, K. (2012b), "Trust, reciprocity and favors in cooperative relationships," *Mimeo.*

[4] Agastya, M., Menezes, F. and Sengupta, K. (2007), "Cheap talk, efficiency and egalitarian cost sharing in joint projects," *Games and Economic Behavior*, 60, 1-19.

[5] Aoyagi, M. and Frechette, G. (2009), "Collusion as public monitoring becomes noisy: experimental evidence," *Journal of Economic Theory*, 144(3), 1135-1165.

[6] Athey, S. and Bagwell, K. (2001), "Optimal collusion with private information," *The RAND Journal of Economics*, 32:3, 428-465.

[7] Athey, S., Bagwell, K. and Sanchirico, C. (2004), "Collusion and price rigidity," *The Review of Economic Studies*, 71, 317-349.

[8] Atkeson, A. and Lucas Jr., R.E. (1992), "On efficient distribution with private information," *The Review of Economic Studies*, 59(3), 427-453.

[9] Azevedo, E. and Moreira, H. (2007), "On the impossibility of an exact imperfect monitoring folk theorem," *Mimeo.*

[10] Baliga, S. and Sjostrom, T. (2004), "Arms races and negotiations," *Review of Economic Studies*, 71, 351-369.

[11] Balliet, D. (2010), "Communication and cooperation in social dilemmas: a meta-analytic review," *Journal of Conflict Resolution*, 54, 39-57.

[12] Belianin, A. and Novarese, M. (2005), "Trust, communication and equilibrium behavior in public goods," *Working paper.*

[13] Berg, J., Dickhaut, J. and McCabe, K. (1995), "Trust, reciprocity, and social history," *Games and Economic Behavior*, 10(1), 122-142.

[14] Bigoni, M., Casari, M., Skrzypacz, A. and Spagnolo, G. (2011), "Time horizon and cooperation in continuous time," *Working paper.*

[15] Bochet, O., Page, T. and Putterman, L. (2006), "Communication and punishment in voluntary contribution experiments," *Journal of Economic Behavior and Organization*, 60, 11-26.

[16] Brosig, J., Ockenfels, A. and Weimann, J. (2003), "The effect of communication media on cooperation," *German Economic Review*, 4, 217-242.

[17] Brown-Kruse, J., Cronshaw, M. B., and Schenk, D. J. (1993), "Theory and experiments on spatial competition," *Economic Inquiry*, 31, 139-165.

[18] Brunnermeier, M.K. and Morgan, J. (2010), "Clock games: theory and experiments," *Games and Economic Behavior*, 68(2), 532-550.

[19] Calcagno, R. and Lovo, S. (2010), "Preopening and equilibrium selection," *Mimeo.*

[20] Calvert, R. (1989), "Reciprocity among self-interested actors: uncertainty, asymmetry and distribution," *Peter C. Ordeshook, ed., Models of Strategic Choice in Politics*, University of Michigan Press.

[21] Camerer, C. (1988), "Gifts as economic signals and social symbols," *American Journal of Sociology*, 94, S180-S214.

[22] Camerer, C., Kang, M.J. and Ray, D. (2012), "Measured anxiety and choices in experimental timing games," *Working paper.*

[23] Carmichael, H.L. and MacLeod, W.B. (1997), "Gift giving and evolution of cooperation," *International Economic Review*, 39, 485-509.

[24] Cason, T.N. and Khan, F.U. (1999), "A laboratory study of voluntary public goods provision with imperfect monitoring and communication," *Journal of Development Economics*, 58, 533-552.

[25] Cason, T.N., Lau, S.P. and Mui, V.L. (2012), "Learning, teaching and turn-taking in the repeated assignment game," *Working paper.*

[26] Charness, G. and Haruvy, E. (2002), "Altruism, equity, and reciprocity in a gift-exchange experiment: an encompassing approach," *Games and Economic Behavior*, 40(2), 203-231.

[27] Cheung, Y-W. and Friedman, D. (2009), "Speculative attacks: a laboratory study in continuous time," *Journal of International Money and Finance*, 28(6), 1064-1082.

[28] Cooper, R., Dejong, D.V., Forsythe, R. and Ross, T.W. (1989), "Communication in the battle of the sexes game: some experimental results," *Rand Journal of Economics*, 20, 568-587.

[29] Cooper, R., Dejong, D.V., Forsythe, R. and Ross, T.W. (1992), "Communication in coordination games," *Quarterly Journal of Economics*, 107, 739-771.

[30] Costa, F. J. and Moreira, H. A. (2012), "On the limits of cheap talk for public good provision," *Working paper*.

[31] Crawford, V. P. and Sobel, J. (1982), "Strategic information transmission," *Econometrica*, 50(6), 1431-1451.

[32] Dal Bo, P. (2005), "Cooperation under the shadow of the future: experimental evidence from repeated games," *American Economic Review*, 95(5), 1591-1604.

[33] Dal Bo, P. and Frechette, G. (2011), "The evolution of cooperation in infinitely repeated games: experimental evidence," *American Economic Review*, 101(1), 411-429.

[34] Daughety, A. F. and Forsythe, R. (1987a), "The effects of industry-wide price regulation on industrial organization," *Journal of Law, Economics, and Organization*, 3, 397-434.

[35] Daughety, A. F. and Forsythe, R. (1987b), "Industry-wide regulation and the formation of reputations: a laboratory analysis," In Bailey, E. E., editor, *Public Regulation: New Perspectives on Institutions and Policies*, pages 347-398. MIT Press, Cambridge, Massachusetts.

[36] Dawes, R. M., Orbell, J. M., Simmons R. T. and Van de Kragt, A. (1986), "Organizing groups for collective action," *American Political Science Review*, 80(4), 1171-1185.

[37] Deck, C. and Nikiforakis, N. (2012), “Perfect and imperfect real-time monitoring in a minimum effort game,” *Experimental Economics*, 15(1), 71-88.

[38] Dorsey, R. E. (1992), “The voluntary contributions mechanism with real time revisions,” *Public Choice*, 73, 261-282.

[39] Duffy, J. and Feltovich, N. (2002), “Do actions speak louder than words? an experimental comparison of observation and cheap talk,” *Games and Economic Behavior*, 39(1), 1-27.

[40] Duffy, J., Ochs, J., and Vesterlund, L. (2007), “Giving little by little: dynamic voluntary contribution games,” *Journal of Public Economics*, 91, 1708-1730.

[41] Engle-Warnick, J. and Slonim, R.L. (2006), “Learning to trust in indefinitely repeated games,” *Games and Economic Behavior*, 54(1), 95-114.

[42] Farrell, J. (1993), “Meaning and credibility in cheap-talk games,” *Games and Economic Behavior*, 5, 514-531.

[43] Farrell, J. and Gibbons, R. (1989), “Cheap talk can matter in bargaining,” *Journal of Economic Theory*, 48 (1), 221-237.

[44] Farrell, J. and Rabin, M. (1996), “Cheap talk,” *Journal of Economic Perspectives*, 10(Summer), 103-118.

[45] Farrell, J. and Saloner, G. (1988), “Coordination through committees and markets,” *The Rand Journal of Economics*, 19, 235-252.

[46] Feeley, T.H., Tutzauer, F., Rosenfeld, H.L. and Young, M.J. (1997), “Cooperation in an infinite-choice continuous-time prisoner’s dilemma,” *Simulation & Gaming*, 28(4), 442-459.

[47] Feinberg, R.M. and Snyder, C. (2002), "Collusion with secret price cuts: an experimental investigation," *Economics Bulletin*, 3(6), 1-11.

[48] Forges, F. (1990), "Equilibria with communication in a job market example," *Quarterly Journal of Economics*, 105, 375-398.

[49] Fouraker, L. and Siegel, S. (1963), *Bargaining Behavior*, McGraw-Hill, New York.

[50] Friedman, D., Henwood, K. and Oprea, R. (2011), "Separating the hawks from the doves: evidence from continuous time laboratory games," *Journal of Economic Theory, forthcoming.*

[51] Friedman, D. and Oprea, R. (2012), "A continuous dilemma," *American Economic Review*, 102(1), 337-363.

[52] Friedman, J. W. (1967), "An experimental study of cooperative duopoly," *Econometrica*, 35, 379-397.

[53] Fudenberg, D., Levine, D. and Maskin, E. (1994), "The Folk theorem with imperfect public information," *Econometrica*, 62, 997-1039.

[54] Fudenberg, D., Rand, D.G. and Dreber, A. (2012), "Slow to anger and fast to forgive: cooperation in an uncertain world," *American Economic Review*, 102(2), 720-749.

[55] Fudenberg, D., Rand, D.G. and Dreber, A. (2013), "It's the thought that counts: the role of intentions in reciprocal altruism," *Mimeo.*

[56] Garicano, L. and Santos, T. (2004), "Referrals," *American Economic Review*, 94(3), 499-525.

[57] Goren, H., Kurzban, R. and Rapoport, A. (2003), "Social loafing vs. social enhancement: public goods provisioning in real-time with irrevocable commitments," *Organizational Behavior and Human Decision Processes*, 90, 277-290.

[58] Goren, H., Kurzban, R. and Rapoport, A. (2004), "Revocable commitments to public goods provision under the real-time protocol of play," *Journal of Behavioral Decision Making*, 17, 17-37.

[59] Green, E.J. (1987), "Lending and the smoothing of uninsurable income," *Contractual Arrangements for Intertemporal Trade (edited by E.C. Prescott & N. Wallace)*, The University of Minnesota Press, 3-25.

[60] Green, E.J. and Porter, R.H. (1984), "Non-cooperative collusion under imperfect price information," *Econometrica*, 52, 87-100.

[61] Hauser, C. and Hopenhayn, H. (2011), "Trading favors: optimal exchange and forgiveness," *Working paper.*

[62] Hertel, J. (2004), "Efficient and sustainable risk sharing with adverse selection," *Mimeo, Princeton University.*

[63] Holcomb, J.H. and Nelson, P.S. (1997), "The role of monitoring in duopoly market outcomes," *Journal of Socio-Economics*, 26(1), 79-93.

[64] Holt, C. H. (1985), "An experimental test of the consistent-conjectures hypothesis," *American Economic Review*, 75, 314-325.

[65] Horisch, H. and Kirchkamp, O. (2010), "Less fighting than expected- experiments with wars of attrition and all pay auctions," *Public Choice*, 144(1), 347-367.

[66] Huck, S., Muller, W. and Normann, H-T. (2001), "Stackelberg beats Cournot: on collusion and efficiency in experimental markets," *The Economic Journal*, 111, 749-765.

[67] Isaac, R. M. and Walker, J. M. (1988), "Communication and free-riding behavior: the voluntary contribution mechanism," *Economic Inquiry*, 264, 585-608.

[68] Jackson, M.O., Rodriguez-Barraquer, T. and Tan, X. (2011), "Slow to anger and fast to forgive: cooperation in an uncertain world," *American Economic Review, forthcoming.*

[69] Kahan, J. and Rapoport, A. (1974), "Decisions of timing in bipolarized conflict situations with complete information," *Acta Psychologica*, 38, 183-203.

[70] Kahan, J., Rapoport, A. and Stein, W. (1976), "Decisions of timing in experimental probabilistic duels," *Journal of Mathematical Psychology*, 13(2), 163-191.

[71] Kalla, S.J. (2010), "Essays in favor trading," *PhD dissertation, University of Pennsylvania.*

[72] Kamada, Y. and Kandori, M. (2011), "Revision games," *Mimeo.*

[73] Kaplan, T.R. and Ruffle, B.J. (2011), "Which way to cooperate," *The Economic Journal, forthcoming.*

[74] Kawamura, K. (2011), "A model of public consultation: why is binary communication so common?," *The Economic Journal*, 121, 819-842.

[75] Kurzban, R., McCabe, K., Smith, V. L. and Wilson, B. J. (2001), "Incremental commitment and reciprocity in a real-time public goods game," *Personality and Social Psychology Bulletin*, 27(12), 1662-1673.

[76] Kuzmics, C., Palfrey, T.R. and Rogers, B.W. (2012), "Symmetric play in repeated allocation games," *Working paper.*

[77] Landa, J.T. (1994), "Trust, ethnicity and identity," *University of Michigan Press, Ann Arbor.*

[78] Lau, S.P. and Mui, V.L. (2008), "Using turn-taking to mitigate coordination and conflict problems in the repeated battle of the sexes game," *Theory and Decision*, 65, 153-183.

[79] Lau, S.P. and Mui, V.L. (2012), "Using turn-taking to achieve intertemporal cooperation and symmetry in infinitely repeated 2x2 games," *Theory and Decision*, 72, 167-188.

[80] Luce, R. and Raiffa, H. (1957), *Games and Decisions: Introduction and Critical Survey*, Wiley, New York, NY.

[81] Matthews, S. (1989), "Veto threats: rhetoric in a bargaining game," *Quarterly Journal of Economics*, 104, 347-369.

[82] Matthews, S. and Postlewaite, A. (1989), "Preplay communication in two-person sealed-bid double auction," *Journal of Economic Theory*, 48, 238-263.

[83] Mobius, M. (2001), "Trading favors," *Working paper, Harvard University.*

[84] Nayyar, S. (2009), "Essays on repeated games," *PhD dissertation, Princeton University.*

[85] Neill, D.B. (2003), "Cooperation and coordination in the turn-taking dilemma," *In 14th International Conference on Game Theory at Stony Brook.*

[86] Neilson, W.S. (1999), "The economics of favors," *Journal of Economic Behavior and Organization*, 39, 387-397.

[87] Oprea, R., Wilson, B. and Zillante, A. (2011), “War of attrition: evidence from a laboratory experiment on market exit,” *Working paper.*

[88] Ordeshook, P. and Palfrey, T. R. (1988), “Agendas, strategic voting, and signaling with incomplete information,” *American Journal of Political Science*, 32, 441-466.

[89] Ostrom, E. and Walker, J. M. (1991), “Communication in a commmons: cooperation without external enforcement,” in T. R. Palfrey, ed. *Laboratory Research in Political Economy*, Ann Arbor: University of Michigan Press.

[90] Palfrey, T. R. and Rosenthal, H. (1988), “Private incentives and social dilemmas: the effects of incomplete information and altruism,” *Journal of Public Economics*, 35, 309-332.

[91] Palfrey, T. R. and Rosenthal, H. (1991), “Testing for effects of cheap talk in a public goods game with private information,” *Games and Economic Behavior*, 3, 183-220.

[92] Palfrey, T.R. and Rosenthal, H. (1994), “Repeated play, cooperation and coordination: an experimental study,” *Review of Economic Studies*, 61(3), 545-565.

[93] Thomas, J. and Worrall, T. (1990), “Income fluctuations and asymmetric information: as example of the repeated principal-agent problem,” *Journal of Economic Theory*, 51, 367-390.

[94] Van de Kragt, A., Orbell, J. and Dawes, R. (1983), “The minimal contributing set as a solution to public goods problems,” *American Political Science Review*, 77, 112-121.

[95] Waichman, I., Requate, T., and Siang, C. K. (2011), “Pre-play communication in Cournot competition: an experiment with students and managers,” *Mimeo.*