DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

COOPERATION WITHOUT IMMEDIATE RECIPROCITY: AN EXPERIMENT IN FAVOR EXCHANGE

Nilanjan Roy

SOCIAL SCIENCE WORKING PAPER 1358

January 2012

Cooperation Without Immediate Reciprocity: An Experiment in Favor Exchange

Nilanjan Roy

Abstract

This paper presents experimental evidence concerned with behavior in indefinite horizon two-person dynamic favor exchange games. Subjects interact in pairs in continuous time and occasionally one of them receives opportunity to provide a favor to her partner. The effects of changing the benefit of receiving a favor and the arrival rate of opportunities to do a favor are studied when the opportunities are privately observed. Also considered are the impacts of informational access to partner's opportunities on efficiency and the overall behavior of individuals with respect to "obvious" state variables.

JEL classification numbers: C72, C73, C91, C92, D82

Key words: favor exchange, continuous time games, private information, infinite horizon, chips mechanisms

Cooperation Without Immediate Reciprocity: An Experiment in Favor Exchange

Nilanjan Roy*

A favor is a costly action that an individual can take which benefits some other person but yields no direct benefit to the individual performing the favor. In theory, an individual may perform a favor in the expectation of return of favor from her partner in the future. Exchange of favors[1] is the economic basis of much of the informal cooperation that goes on among individuals in a society. An enormous part of economic activities that remain fairly uncaptured by the 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, when trying to pass their own legislations. These behaviors generate mutual gains for self-interested individuals.

Using controlled experiments, this study wishes to explore the characteristics of relationships in which people exchange favors using the following basic setting, based on models developed by Mobius (2001) and Hauser and Hopenhayn (2011). Two players interact in continuous time for an indefinite length and occasionally, one of the players receive opportunity to provide a favor to the other player. Only one of the players 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, thus, it is socially optimal to always provide a favor. Different scenarios are considered depending on whether or not the ability to do a favor is private information. Both the complete and private information settings capture certain aspects of everyday interactions among self-interested individuals. For example, the complete information environment seems to be a reasonable assumption in the context of village economies in developing countries. People living there 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

*Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena CA 91125, USA ⟨*royn@hss.caltech.edu*⟩; I am grateful to Thomas Palfrey for excellent guidance and generous support for running the experiments. The funding from Gordon and Betty Moore Foundation is gratefully acknowledged. I am also thankful to Chris Crabbe for developing the experimental software. Federico Echenique, Ben Gillen, Julian Romero and Matthew Shum provided helpful suggestions.

[1]Favor exchange is also known as backscratching. There is a famous idiom that says: You scratch my back and I'll scratch yours.

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[2]. On the other hand, to appreciate the need for a model with a private information component, consider the scenario where two firms are involved in a partnership. At random times one of them finds a discovery which if disclosed can make total payoffs higher. But, disclosure comes at the cost of lowering the discovering firm's payoffs. It is perfectly possible for these random discoveries to be privately observed and this is also a situation without immediate reciprocity.

The favor exchange game is however, different from the widely studied prisoner's dilemma in that the players act sequentially. The stochasticity of the opportunities and the introduction of private information makes the environment completely different from the classic prisoner's dilemma. However, providing a favor can be thought of as cooperating, thus, favor provision is synonymous to cooperation. Also, generally, an individual performing a favor decides neither when to perform the favor nor how large a favor to perform. These parameters are exogenously given to an individual because the other individual needs a specific favor at a specific time. The only remaining decision is a binary one concerning whether or not to perform the favor.

As it is impossible to gather private information on the arrival of opportunities, as well as on the costs and benefits of favors in the field, a laboratory seems perfect to examine the behavior of individuals in a favor-exchange setting. 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. This paper implements a two-person favor exchange model in the laboratory in which opportunities to perform a favor arises stochastically over time. Time is continuous and there are no discrete periods. A continuous time setting seems ideal to study the behavior of individuals, at least for the current dynamic game. The advantages are two-fold. First, it is much easier for the subjects to make decisions in a continous time game. They only need to switch their actions whenever they want to. In contrast, in a discrete time game, each period they need to make a choice and the experiment gets stuck if some subjects are slow to respond. Also, the number of data points are much larger than a discrete time analogue. The data set used in this experiment is very rich containing thousands of observations.

The basic purpose of this paper is to develop a new experimental framework for studying behavior in a two person dynamic favor exchange game in continuous time. Some basic set of questions are as follows. What percentage of total favors are granted when individuals play the favor exchange game under different parameter scenarios. Do we see a significant increase in the rate of favor provision if the returns from such a relationship are higher? Does increasing the rate at which opportunities to provide a favor arrives make any difference in the overall rate of favor provision? Is efficiency significantly different from the level that would be generated if individuals implement a mechanism based

[2]See Camerer (1988), Landa (1994), Carmichael and MacLeod (1997), and Charness and Haruvy (2002) for studies and references on gift-exchange relationships.

on the net favors granted called as the chips mechanism (CM)? Under chips mechanisms, people keep track of the imbalances in number of favors provided between them and engage in score keeping, with the restoration of balance being a primary aspiration. As argued in Fiske (1992), this pattern of behavior provides a crucial description of a part of the human social interactions. This also leads to the following question: How much inequality in payoffs do individuals tolerate? Is that similar to the level that is expected to emerge if everyone follow the most efficient chips mechanism?

The next question posed is that whether letting individuals observe the opportunities received by her partner helps them achieve a higher efficiency. In the absence of such 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. Thus, the question asked is do we see a significant change in the overall behavior of individuals depending on the presence or absence of the informational constraint.

The contributions of the present study are three-fold. First, it contributes to the empirical literature on cooperation in dynamic games. Specifically, it is concerned with analyzing how cooperation is sustained in a long term relationship without immediate reciprocity. Second, it adds to the understanding of the effects of monitoring on cooperation in dynamic games. Third, at a more general methodological level, it contributes to the budding literature on dynamic game experiments conducted in a continuous time setting.

The remainder of the paper is organized as follows. Section I presents the model on which the experiment is based and discusses the chips mechanism in detail. Section II provides an overview of the theoretical literature on favor exchange alongwith the literature on continuous or near-continuous experiments in repeated games. Section III describes the experimental design, procedures and the testable hypotheses. The results are discussed in detail in section IV. Section V concludes.

I The Model

The model is identical to the 'Direct Favors' model as presented in Mobius (2001). Two individuals interact indefinitely in continuous time with both of them discounting future utility at a rate r. Each individual can provide the other person a favor at random times depending on whether or not she receives an opportunity to do so. These random opportunities arrive with probability p for each person, with a restriction that no two individuals can provide favors at the same time. Thus, the opportunities arrive according to a Poisson process and the inter-arrival time between opportunities follows an exponential distribution. No fractional favors are allowed, so each favor has size one. A granted favor has utility b to the receiver and it costs $c < b$ to provide a favor. Thus, it is socially optimal to always provide favors. Letting x_i represent the number of favors provided by person i and x_j as the number of favors received, the utility for person i is given by:

$$U_i(x_i, x_j) = bx_j - cx_i$$

If the opportunities to provide favors are observable by both individuals then a perfect public equilibrium[3] would exist that achieves the social optimum through a simple grim trigger 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

$$\frac{p(b-c)}{r} > c$$

However, when the opportunities to provide favors are privately observed, individuals now have to infer from the history of favor exchange to what extent their partner provided favor for them. A simple and intuitive accounting device as analyzed by Mobius is 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 an attractive choice to condition individual's strategies as it captures the common notion that someone 'owes' a favor (negative k) or is owned a favor (positive k). Thus, one can focus on equilibria where players' strategies only depend on the state variable k_i[4]. 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, denote a public history upto time t consisting of individuals' past favors by h^t. Attention is restricted to sequential equilibria in which 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). In fact the equilibria considered by Mobius is a form of PPE. However, it is further restricted to depend on only the state variable k_i. Thus, a strategy for an individual can be defined as follows.

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

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

[3]A strategy for person i is public if at every instant t, it depends only on the public history at t. A perfect public equilibrium is a profile of public strategies that forms a subgame perfect Nash equilibrium at any date, given any public history.

[4]Generally, these equilibria will not be optimal if individuals also have access to information about the order and the time with which they received past favors.

These equilibria are known as chips mechanisms (CMs). The rationale being the following. One can think of both persons starting with k^* poker chips each. Whenever one receives a favor, she gives the other person a chip. If a person runs out of chips, she receives no favor until she grants one to the other individual and obtains a chip in exchange. Under a CM, people keep track of the imbalance 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[5], that is, with a $k^* = 1$ can be termed the simple chips mechanism (SCM). Under a SCM, once an individual provides a favor, she is unwilling to provide any further favor until she receives one from her partner. This would induce a public history of alternating favors. 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)[6].

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:

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

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

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

The equation system defines a perfect public equilibrium if the incentive compatibility (IC) constraints and the individual rationality (IR) constraints are satisfied. Not only individuals have to gain from doing favors for $k < k^*$, the relationship between two people must have positive value for $k = -k^*$. Thus, we must have the following:

$$\begin{aligned} &V_{k^*}(k+1) - V_{k^*}(k) \geq c \quad for - k^* \leq k < k^* \quad (ICC) \\ &V_{k^*}(-k^*) \geq 0 \qquad\qquad\qquad\qquad\qquad (IRC) \end{aligned}$$

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

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

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

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

where α, β, x, and y are given by

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

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

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

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

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

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 increases 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[7]:

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

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 makes it an obvious choice as the benchmark in the private

[7] For a detailed proof, see Mobius (2001).

information environment. Fiske (1992) surveys ethnographic field work and experimental studies and argues that virtually all human social interactions can be described in terms of four patterns, each with a distinctive psychological basis. One of those patterns is chips mechanism[8]. These 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. Technically, a CM equilibrium has the properties of an ordered Abelian group. That is, the structure of CM relations exhibits all the properties of a linear ordering, and also entails the idea of an additive identity and an additive inverse. It also obeys the associative and commutative laws. Finally, addition under CM is order preserving.

Although the best CM is asymptotically efficient in p/r[9] 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, as emphasized by Hauser and Hopenhayn (2011). 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. Although the optimal PPE can achieve higher efficiency than the BCM, the first best is not achievable when the opportunities are private information. However, as mentioned earlier, this first best could be achieved in a simple equilibrium with grim trigger strategies (under a parametric restriction) if favor opportunities are publicly observed.

II Related Literature

This section has two parts. The first part presents a brief overview of the theoretical models developed in the favor exchange literature. Since the present paper contributes to a budding literature on repeated game experiments in continuous time framework, a short description of the studies under this literature is provided in the second part of this section.

[8]Chips mechanism can also be termed as equality matching.

[9]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.

A. Theoretical Literature on Favor Exchange Models

A two person model of favor exchange with private information was first studied by Mobius (2001). 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. Kalesnik (2005) shows that if one forbids partial favors, the game by Hauser and Hopenhayn generates the same set of expected payoffs as a good news prisoners' dilemma with two signals where, if a player shirks, his signal has zero probability of arriving.

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. Although Calvert introduces uncertainty and asymmetry in his model, there is no private information among actors.

Neilson (1999) examines the favors people perform for each other using a stochastic version of the infinitely repeated Prisoner's Dilemma game. He 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.

Nayyar (2009) develops a discrete-time model of favor exchange that allows for asymmetric opportunities for doing favors and finds cooperation among players can be supported asymptotically in equilibrium even under those asymmetries.

Kalla (2010) studies favor trading when players have privately known types. By introducing low and high types[10], he finds conditions under which the high type players are almost always able to separate themselves from the low type players.

Two other studies are related to the favor exchange game focusing on the perfect public equilibrium. 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 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 producing when the other firm has high costs. They characterize the set of perfect public equilibrium values. Abdulkadiroglu and Bagwell (2010) analyze a two person repeated trust game with private information which has 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.

[10] 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.

B. Experimental Literature on Games in Continuous Time

The experimental set-up in this paper is near-continuous. The literature on continuous/near-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 (2011) study prisoners' dilemma games played in continuous time with flow payoffs accumulated over 60 seconds time. They compare cooperation rates from these continuous time sessions to rates under control grid and one-shot sessions. Bigoni, Casari, Skrzypacz, and Spagnolo (2011) compare behavior in continuous time prisoner's dilemma games under deterministic and stochastic duration. Feeley, Tutzauer, Rosenfeld, and Young (1997) implement infinite-choice continuous time prisoner's dilemma games in the laboratory to observe what sorts of behavioral patterns emerge[11]. 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). Oprea, Wilson, and Zillante (2011) implement a series of wars 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 had to decide when to stop bidding. Also, Friedman, Henwood, and Oprea (2011) implement the standard Hawk-Dove bimatrix game in continuous time in the laboratory[12].

[11]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.

[12]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, Kahan, and Stein (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.

III Experimental Design, Procedures and Hypotheses

This section is divided into three parts. The first two parts discuss the design and procedures of the experiment. The last part lists the hypotheses that are formally tested in this study.

A. Design

In the experiment, the cost of providing a favor, c is fixed at 5 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 complete information environment). These are summarized in Table 1, along with the k^{max} corresponding to the BCM in each case for the private information environment. Also displayed are the intial 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 as such 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 complete information environment. The expected length of a match is 100 seconds. A shorter expected duration is avoided purposefully so as to get games with "high" 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".

Table 1: Treatments, parameters and the k^{max}

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;

B. Procedures

The experiments were all conducted at the Social Science Experimental Laboratory (SSEL) using the Multistage software package. Subjects were recruited from a pool of volunteer subjects, maintained by the SSEL. A total of six sessions were run, using a total of 84 subjects. No subject participated in more than one session. The experiment has a 2x2x2 factorial design. The treatment variables are the benefit of receiving a favor

Table 2: Treatments and Sessions

Date	Session	Information	Subjects	Exchange Rate	Treatment*	Matches	Games
5/9/11	1	Private	12	160	Low-Slow	4	24
5/9/11	1	Private	12	160	Low-Fast	4	24
5/9/11	1	Private	12	160	High-Fast	4	24
5/9/11	1	Private	12	160	High-Slow	4	24
5/9/11	2	Private	16	160	High-Slow	4	32
5/9/11	2	Private	16	160	High-Fast	4	32
5/9/11	2	Private	16	160	Low-Fast	4	32
5/9/11	2	Private	16	160	Low-Slow	4	32
5/10/11	3	Private	16	160	High-Slow	4	32
5/10/11	3	Private	16	160	Low-Slow	4	32
5/10/11	3	Private	16	160	Low-Fast	4	32
5/10/11	3	Private	16	160	High-Fast	4	32
5/12/11	4	Private	12	160	Low-Fast	4	24
5/12/11	4	Private	12	160	High-Fast	4	24
5/12/11	4	Private	12	160	High-Slow	4	24
5/12/11	4	Private	12	160	Low-Slow	4	24
11/17/11	5	Complete	12	160	Low-Slow	4	24
11/17/11	5	Complete	12	160	Low-Fast	4	24
11/17/11	5	Complete	12	160	High-Fast	4	24
11/17/11	5	Complete	12	160	High-Slow	4	24
11/19/11	6	Complete	16	160	High-Slow	4	32
11/19/11	6	Complete	16	160	High-Fast	4	32
11/19/11	6	Complete	16	160	Low-Fast	4	32
11/19/11	6	Complete	16	160	Low-Slow	4	32

Exchange Rate : US$ 1 was worth "Exchange Rate" number of 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.
*: in the order in which it was presented to participants.

(low and high), arrival rate of opportunities to provide a favor (slow and fast) and the information about the partner's opportunities (private and complete). Table 2 summarizes the characteristics of each session[13]. On arrival, instructions[14] were read aloud. Subjects interacted anonymously with each other through computer terminals. There was no possibility of any kind of communication between the subjects. 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[15]. A session lasted on average 1 hour and 5 minutes. Subjects earned on average US$15, and US$26, respectively[16].

[13]Additional four sessions were also run under the complete information environment with 52 subjects and with a different protocol (exchange rate and number of sessions). However, for the purposes of 'direct comparison', the data from these sessions are not reported and also not included in any of the statistical analyses. Using these data generates similar qualitative results but does change the significance level in some of the results. Throughout the current study data from these sessions serve as robustness checks.

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

[15]Appendix B discusses in detail about the user interface and screen display.

[16]Payoffs ranged from US$6 to US$23 in the private information treatments and from US$14 to US$31 in the complete information treatments. Each subject also earned an additional US$5 show-up fee.

In each match, the arrival of opportunities for each pair of participants were decided using a random draw of an integer from an uniform distribution after each second[17]. The discounted payoffs were induced by a random termination rule with the draw of an integer from an 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.

Subjects were randomly rematched with a new partner each match. Random rematching protocol allows for a large 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 problem. 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.

C. Hypotheses

The testable hypotheses are listed in this section. It must be noted that when a particular parameter or variable of interest is changed to examine the effect on favor provision rate or likelihood, all other parameters remain same to give rise to a 'ceterus paribus' change. The first hypothesis is that when opportunities to provide a favor are privately observed, increasing the benefit of receiving a favor makes the favor exchange relation more rewarding and valuable. Hence, the favor provision rate should be higher compared to the rate when the benefit is lower. Second, favor provision rate should be higher when the opportunities arrive at a faster rate. This is because a faster arrival rate of opportunities makes individuals more willing to help as cheating becomes harder. The third hypothesis compares the overall rate of favor provision under the complete and private information treatments. As discussed earlier, although the first best outcome can be supported under complete information (given the selected parameters for the experiment), it is not supported under the PPE in case of private information. Hence, one should expect that the rate of favor provision under complete information is strictly higher than the rate under private information. Fourth, when the opportunities to provide a favor are privately observed, the probability of granting a favor by an individual should go down as the net number of favors done by her goes up. It is also intuitive to assume that the likelihood of favor provision goes down as more time elapses after a favor has been provided by her partner. Fifth, in the absence of any informational constraints, an individual can condition rewards and punishments directly based on whether or not her

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

partner provided the favor at the last opportunity received. Thus, the hypotheses that are formally tested are:

Hypothesis 1. *In the case where 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 the benefit is low.*

Hypothesis 2. *In the case where the opportunities to provide a favor are privately observed, a faster arrival rate of opportunities results in a higher rate of favor provision than when arrival rate is slow.*

Hypothesis 3. *For each parameter configuration, the rate of favor provision under complete information is significantly higher than the rate under private information.*

Hypothesis 4. *In the case where the opportunities to provide a favor are privately observed, the likelihood of favor provision by an individual declines with an increase in net favors done by her as well as the time since last favor done by her partner.*

Hypothesis 5. *In the case where the opportunities to provide a favor are publicly observed by both players, 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.*

IV Results

A set of ten empirical results are reported in the following five subsections. The results obtained from varying the benefit of receiving a favor and the rate at which opportunities to provide a favor arrive for the private information treatments are presented in section A. Section B offers the results on the effect of monitoring partner's opportunities on the rate of favor provision. Findings on the inequalities in payoffs tolerated by the individuals are discussed in section C. Section D focuses on the variables that are likely to affect the decision to grant a favor and provides results on how the probability of favor provision responses to those variables on an average. Findings on volatility are presented in section E.

A. Favor Provision under Private Information

The first objective is to study the behavior in the environment when the opportunities to provide a favor are privately observed by the individuals. 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. For tractability, the reported data are taken at one second intervals and thus, the action of a subject each second is either being in a 'Do favor' mode or in 'Do not do favor' mode.

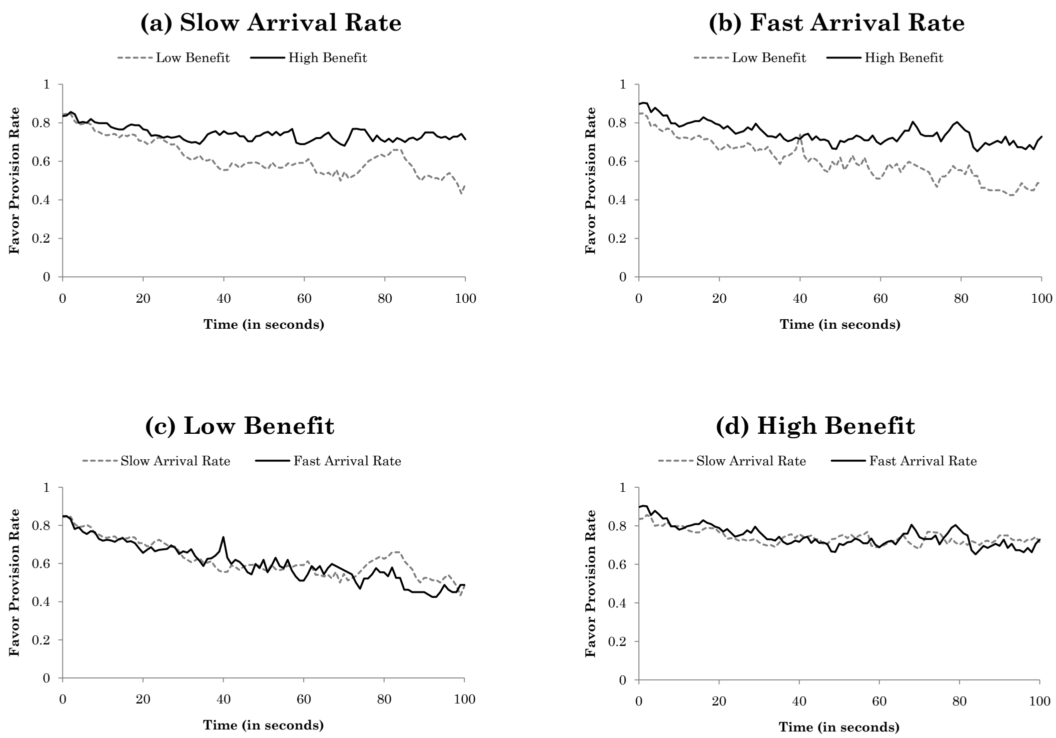

Figure 1: Rate of Favor Provision over time - Private Information Treatments

Table 3: Mean Favor Provision Rates

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 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 by treatment, aggregating over time and all subjects across sessions. As seen in the table, the levels of favor provision are higher in the high benefit 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. See Table 4 for the p-values of a difference in means t-test[18]. Further evidence is documented in panels (a) and (b) of Figure 1 that plots the instantaneous probability of granting a favor over time in these treatments [19]. The graphs collect the overall average behavior over time[20]. The time path for the high benefit scenario is always above the time path for low benefit one. So, supporting Hypothesis 1, the following result is obtained.

Result 1. *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 the benefit is low.*

Table 4: Private Info. treatments - Inter Comparisons

Comparison	p-value
Low-Slow versus Low-Fast	0.743
High-Slow versus High-Fast	0.734
Low-Slow versus High-Slow	0.052
Low-Fast versus High-Fast	0.001

p-values are reported from a t-test of difference in means using robust standard errors clustered at the pair level to acknowledge the fact that the observations within each dyadic game are dependent. It also includes 'time' as a covariate (time trend).

Table 3 also shows that changing arrival rate of opportunities generate similar levels of favor provision. The differences (0.95 points in low benefit treatment and 3.33 in the

[18]All p-values are calculated taking into account the lack of independence in the decisions of each pair of participants (group cluster robust standard errors).

[19]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 are shown only till 100 seconds.

[20]The behavior in the later time intervals (or seconds) are obviously dependent on the behavior of an individual prior to that time and also on the nature of play of her partner.

high benefit treatment) are statistically indistinguishable with p-values > 0.1 (see Table 4). Panels (c) and (d) of Figure 1 give further evidence of this. See that the time paths under slow and fast arrival rate cannot be distinuguished. Thus, the following conclusion is reached.

Result 2. *When the opportunities to provide a favor are privately observed, a faster arrival rate of opportunities generates similar rate of favor provision as when the arrival rate is slow. This rejects Hypothesis 2.*

It is important to note 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, the average favor provision declines for initial few seconds and then settles down. It can also be seen from Figure 1 that this decline is much steeper in the low benefit treatments as compared to the high benefit ones.

Next, the mean rate of favor provision predicted under the best CM is higher than what is observed in the data for each of the four treatments (see Table 3 for details). This is obviously not a test of whether an individual is following the best CM or not. Rather it just compares the overall rate of favor provision observed in the data to the (hypothetical, yet possible) situation where everyone follows the best CM. Testing for individual strategies is beyond the scope of the present study. 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 (see Table 5). This is also confirmed from the comparison of the time path of the probability of doing a favor in the data to the time path that would have been generated if everyone followed the best CM, as shown in Figure 2. Tables 3 and 5 together with Figure 2 support the following conclusion.

Result 3. *Favor provision is significantly less than the level under BCM*[21].

Table 5: Private Info. Treatments- Data Versus CM

Treatment	Data versus Simple CM	Data versus Best CM
Low-Slow	0.110	0.000
Low-Fast	0.043	0.000
High-Slow	0.177	0.000
High-Fast	0.008	0.000

p-values are reported from a t-test of difference in means using robust standard errors clustered at the pair level and including a time trend.

Figure 2 also shows that both the simple and the best chips mechanisms predict that

[21] This also implies less favor provision than under optimal equilibrium (PPE).

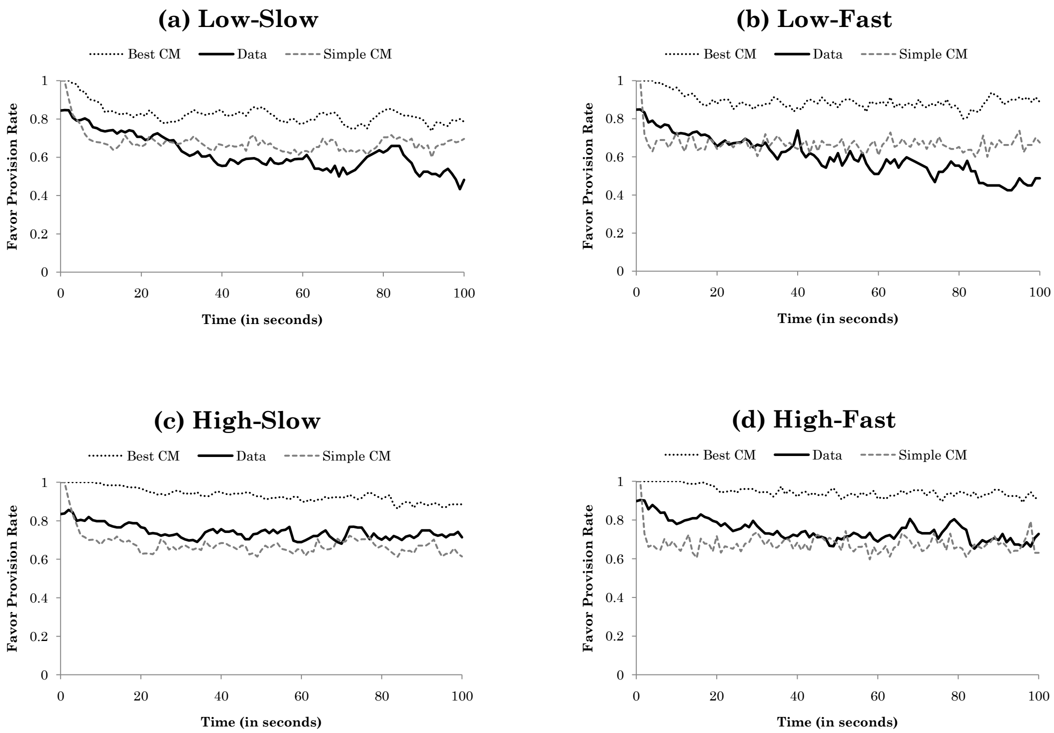

Figure 2: Rate of Favor Provision over time - Comparison with the Best and Simple Chips Mechanisms

the probability of granting a favor starts out at one[22] and then declines over the early few seconds, and then generating 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 decline in the probability as seen in the data is far less steep than the SCM. Table 3 also displays the favor provision rates as predicted by a SCM and Table 5 shows the p-values from a t-test of difference in means in the data and the SCM. The p-values from such a test are less than 0.05 in the treatments with a fast arrival rate of opportunities, showing that the favor provision as observed in the data is significantly different from the level that would have been generated if everyone followed the simple chips mechanism. It is 61.47 vs. 67.53 and 74.99 vs. 67.92 for data vs. SCM in the Low-Fast and High-Fast respectively. However, statistically indistinguishable rates are generated among the data and SCM in each of the treatments with slow arrival rate of opportunities (with p-values > 0.1). Thus, we have the following result.

Result 4. *Favor provision rate is consistent with the rate generated by the SCM when the arrival rate of opportunities is slow. When the arrival rate is fast, favor provision is significantly different from the level predicted by the SCM.*

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

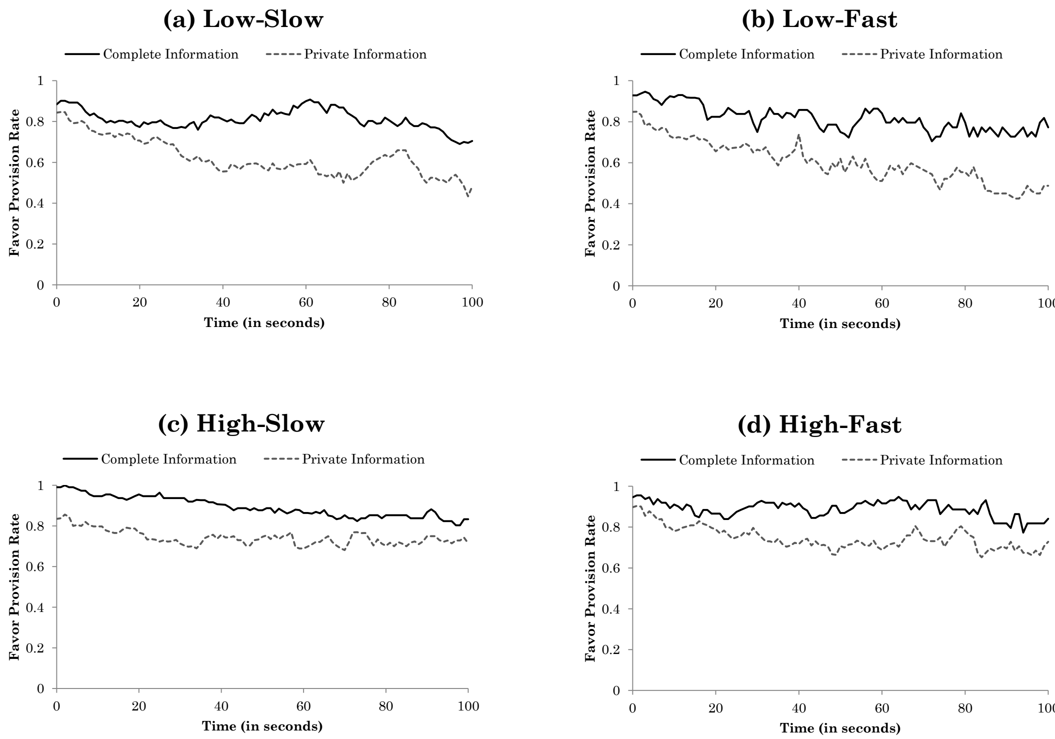

Figure 3: Rate of Favor Provision over time - Private versus Complete Information

B. Effect of Monitoring on Favor Provision

Most of the experimental studies on cooperation in repeated games analyze a complete information setting with perfect monitoring, barring a few exceptions. Since the interest of the present section is to explore the effect of monitoring on favor provision, it is worthwhile to discuss a few studies related to incomplete information or imperfect monitoring in repeated game experiments. Palfrey and Rosenthal (1994) compare the rate of contribution to a public good under incomplete information regarding the contribution cost when players meet only once with the case when they meet repeatedly. They find that repetition leads to more cooperation than one-shot games, but this increase is small. 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 prisoners 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. Aoyagi and Frechette (2009) study the infinitely repeated prisoner's dilemma games under imperfect monitoring and find that cooperation declines with the level of noise of the public signal.

In the present study, perfect monitoring means having the informational access to observe the opportunities received by the partner. This is referred to as the 'complete information' environment. The situation without the ability to monitor the partner's opportunities is referred to as the 'private information' environment or the situation with no monitoring. Focusing on the differences in the initial rate of favor provision[23] across the private and complete information environments shows that the initial favor provision rate is higher under complete information for all 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 6 shows the average percentage of favor provision by treatments in both the private and the complete 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 (p-value < 0.05) when the arrival rate of opportunities is slow and 21.38 points (p-value < 0.001) when the arrival rate is fast. And among the high benefit treatments, this difference is 17.52 percentage points (p-value < 0.001) when the arrival rate of opportunities is slow and 12.00 points (p-value < 0.01) when the arrival rate is fast.

Table 6: Mean Favor Provision Rates- Private Versus Complete Info.

Treatment	Private info. data	Complete info. data	p value
Low-Slow	62.42 [18406]	76.45 [13682]	0.032
Low-Fast	61.47 [16686]	82.85 [9482]	0.000
High-Slow	71.66 [23824]	89.18 [14532]	0.000
High-Fast	74.99 [17440]	86.99 [13992]	0.003

Number of observations in parentheses. p-values are reported from a t-test of difference in means using robust standard errors clustered at the pair level and including a time trend.

Figure 3 displays the overall rate of favor provision over time for the private information and complete information cases for each of the four treatments. The time path of average rate of favor provision under complete information is always above the time path under private information. Thus, we see that the informational constraint results in significant efficiency losses. Thus, Table 6 together with Figure 3 provide qualified support for Hypothesis 3 and the following result.

Result 5. *Favor provision is significantly higher if individuals can observe the opportunities received by their partners than under private information.*

Although the results for private information treatments do not lend themselves for any comparison with the results of the other studies on cooperation in continuous time, the results in the full information treatments can be discussed with reference to these studies[24]. Since, the horizon of interaction is stochastic in the current research, the

[23]It is defined as the proportion of time 0 'Do favor' choices to the total number of decisions. Each treatment has 224 (112) observations under the private (complete) information environment.

[24]The comparisons made in this section are only an attempt to relate to the only existing study for cooperation in continuous 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 prisoners' dilemma

only relevant study for camparison is the analysis of the treatments with stochastic duration as reported in Bigoni, Casari, Skrzypacz, and Spagnolo (2011). The 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 paper, the mean cooperation (or favor provision) rate across all treatments for the full information case is 84.06%. Thus, 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)[25]. 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.

C. Payoffs and Inequality

Using the instantaneous probability of doing a favor as observed in the data, a straightforward calculation gives the path of instantaneous expected payoffs over time for each of the four treatments under private and complete information environment. This is just the $(b-c)$ times the probability of granting a favor conditional on the arrival of opportunity to grant a favor. This is shown in the Figure 4. For each informational environment, the expected instantaneous profits for a participant are highest in the high benefit-fast arrival rate treatment followed by the high benefit-slow arrival rate treatment, which is in turn higher than the low benefit-fast arrival rate treatment. Profits are lowest in the low benefit-slow arrival rate treatment. A direct consequence of the result already established in the previous section is that the expected instantaneous profits for a participant are higher in the complete information case as compared to the private information case for all treatments (see Figure 4).

An important aspect of a favor exchange relation is that 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, inequality among a pair of individuals i and j is defined as the absolute difference of the share of either of the individual's payoff (π_{it}) in the total payoff at time t and 0.5[26], i.e.,

$$inequality_t^{ij} = \mid \frac{\pi_{it}}{\pi_{it} + \pi_{jt}} - 0.5 \mid = \mid \frac{\pi_{jt}}{\pi_{it} + \pi_{jt}} - 0.5 \mid$$

If within a pair, no one has still done a favor to the other person, then this measure

with flow payoffs.

[25]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".

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

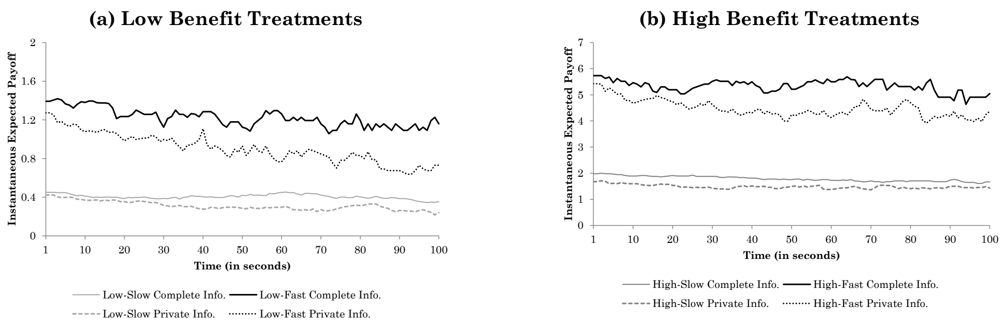

Figure 4: Instantaneous Expected Payoffs over time

is taken to be 0. Thus, it is perfectly possible for both the players to do no favors in the entire play and generate inequality level equal to 0 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 5 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[27] < 0.01 in all four treatments). Although the level of inequality in the data is significantly higher than the predicted level under the best CM (with p-value $<$ 0.06) in the low benefit treatments, the levels are statistically similar in the high benefit treatments with p-values > 0.1.

The inequality in the data is statistically indistinguishable from the level of inequality predicted by the situation when everyone granted all favors, that is, under the socially optimum situation (with p-values > 0.1 in all four treatments). It can also be noted from the figure that inequality is lower in the high benefit treatments than in the low benefit ones.

Next, one could compare the inequality levels tolerated in the private information environment with that in the complete information environment. Except for the High-Fast treatment, the inequality levels are statistically similar with p-values > 0.1. For the High-Fast treatment, inequality is lower (p-value $= 0.072$) in the complete information environment.

All of the above discussions lead to the following finding.

[27] These p-values are obtained from a difference in means t-test including time as a covariate (with a significant negative time trend) and are computed using robust standard errors with clustering at every pair of participants.

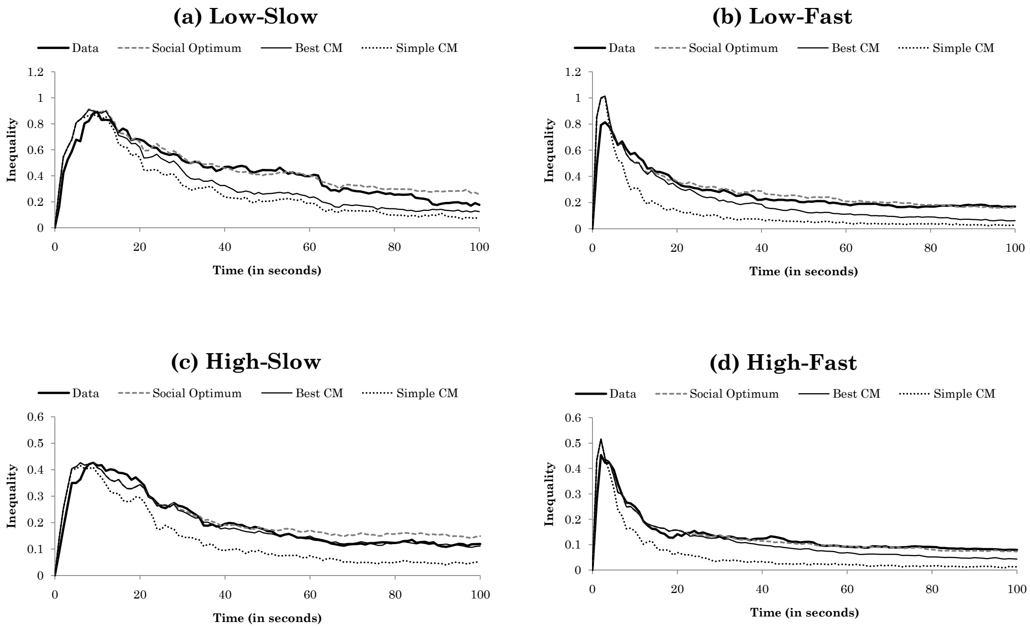

Figure 5: Inequality - Private Information Treatments

Result 6. *When opportunities are privately observed, inequality in payoffs is higher (similar) than the level as under the BCM for the low (high) benefit treatments. It is similar to the level as under the situation when everyone granted all favors. However, inequality is significantly higher than the level under a SCM. Lastly, overall, the inequality in payoffs is statistically indistinguishable among private and complete information environments.*

D. Response to State Variables

Participants were given real time information on the total number of favors done by them, the number of opportunities they received, the number of favors received (and also how many opportunities their partner had forgone in the complete information treatments), 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, 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. Inspite 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 also. 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 information on screen.

Table 7(a): Marginal Effects for Private Information: Behind and Tied

variable	Low-Slow	Low-Fast	High-Slow	High-Fast
$mode_{t-1}$	0.802371	0.6644397	0.7595886	0.6323295
	(0.01794)***	(0.01423)***	(0.01833)***	(0.02263)***
$netfavors_{t-1}$	-0.0635037	-0.0233226	-0.0260648	-0.0110786
	(0.00841)***	(0.0035)***	(0.00375)***	(0.00264)***
$timeelapsed_{t-1}$	-0.001743	-0.0061532	-0.0009457	-0.0033601
	(0.00029)***	(0.00083)***	(0.00019)***	(0.00054)***
$ownpayoff_{t-1}$	-0.0012236	-0.0007185	-0.0001531	-0.0001244
	(0.00027)***	(0.00012)***	(0.00003)***	(0.00003)***
no. of observations	11621	9556	14126	9617
no. of groups	56	56	56	56
pseudo-R^2	0.678	0.413	0.652	0.416

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

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

First, consider the treatments with private information. To understand how the state variables affect the probability of granting a favor, a probit regression is run with random effects at the subject level. The dependent variable is 'mode of person i at time t ($mode_{it}$)' which is a binary variable taking value 1 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 1 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 $t-1(timeelapsed_{it-1})$' and 'own payoff at time $t-1(ownpayoff_{it-1})$'[28].

The marginal effects are summarized in the Tables 7(a) and 7(b) for each of the treatments. The regression results are displayed for two different subcases 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 exactly the same number of favors as her partner till the previous second of play (or nobody has still done a favor from the start of the play) then she is termed as being '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 hugely by the choice at $t-1$ (persistence or inertia). Next, there is negative sign in front of the coefficients for net favors and the coefficients are significantly different from 0 with p-values < 0.01 in all cases. Thus, the probability of granting a favor declines with an increase in the net favors provided. Comparing the magnitudes of the coefficients, one can also conclude that this decline is steeper when the opportunities arrive at a slower rate as compared to the case

[28] Including 'time' as a covariate does not affect the analysis. In fact, the effect of 'time' is insignificant in all treatments.

Table 7(b): Marginal Effects for Private Information: Ahead

variable	Low-Slow	Low-Fast	High-Slow	High-Fast
$mode_{t-1}$	0.889326	0.7655228	0.8897911	0.7775418
	(0.00885)***	(0.00894)***	(0.00732)***	(0.01146)***
$netfavors_{t-1}$	-0.050599	-0.0340664	-0.0617863	-0.0278556
	(0.00932)***	(0.00358)***	(0.00831)***	(0.00426)***
$timeelapsed_{t-1}$	-0.0038308	-0.0098758	-0.0046424	-0.0090145
	(0.00083)***	(0.00158)***	(0.00075)***	(0.00172)***
$ownpayoff_{t-1}$	0.0040194	0.0007338	0.0003149	-0.00000192
	(0.00071)***	(0.00016)***	(0.00011)***	(0.00004)
no. of observations	6561	6906	9474	7599
no. of groups	55	54	55	55
pseudo-R^2	0.689	0.480	0.689	0.493

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

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

when they arrive at a faster pace. Decline is steeper when the benefit is low as compared to the case when the benefit is high[29].

Figure 6 displays the cumulative distribution of net favors observed in the data[30]. The proportion of higher values of net favors is higher, thus, individuals allow for a higher imbalance when the benefit is higher as compared to when it is low (and also when the arrival rate of opportunities is faster). The following result is obtained with respect to net favors as the state variable.

Result 7. *When opportunities are privately observed, the likelihood of favor provision by an individual declines with an increase in the net favors provided by her.*

Figure 7 depicts the moving-average of the probability of granting a favor as a function of time since last favor by partner[31]. It is clear that the probability of granting a favor falls as the time since last favor by partner goes on increasing[32]. The negative signs in front of the coefficients of the variable '$timeelapsed_{t-1}$' in the probit regressions summarized

[29]Except for the one case (out of 4) under the situation where the participant is ahead and comparison is between Low-Slow and High-Slow.

[30]As there are equal number of data points for net favor= k as for net favor= $-k$, only the positive values of net favors are considered.

[31]This also includes the initial data before a favor has been done though by definition there is no last favor before the first favor has been done. In such cases, the time since last favor is just the time since the start of the game.

[32]A closer look into the graphs suggests that there is a flat region initially or even a slight increment in the probability of granting a favor. This is interesting as this clearly shows that in continuous time experiments one needs to acknowledge the possible lags in decision making or that the reaction by participants is not instantaneous. The important result from the viewpoint of the favor exchange relation is however the fact that there is a decline in the likelihood to grant a favor and to continue with the relation as time passes without receiving a favor in return from the partner.

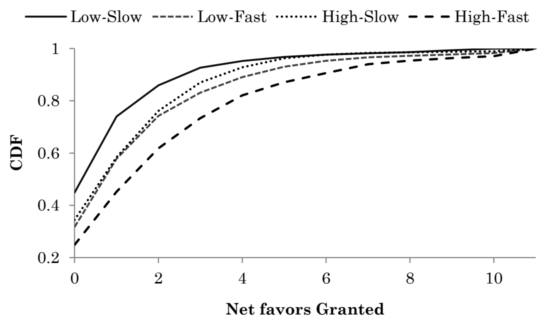

Figure 6: CDF of Net Favors Granted - Private Information Treatments

in Tables 7(a) and 7(b) supports this. The results 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. Also, the absolute magnitude of the marginal effect is higher in the fast arrival rate treatments as compared to the slow arrival rate treatments. This shows that the likelihood of providing a favor declines at a faster rate when the arrival rate of opportunities is higher as is also clear from Figure 7. This is very intuitive because it becomes more likely for a participant to believe that her partner has foregone an opportunity to grant a favor when its a while that she had received a favor in the situation when opportunities arrive at a faster rate. The basic result with respect to the time since last favor as the state variable is as follows.

Result 8. *When opportunities are privately observed, the likelihood of favor provision by an individual declines with an increase in the time since last favor provided by her partner.*

The effect of the variable 'own payoff' on the likelihood of granting a favor is , at best, inconclusive. The signs are positive for some of the coefficients and negative for others and also that the coefficients are an order of magnitude too small. It is also insignificant in one of the regressions.

When the opportunities to provide a favor are observed by both the players, then apart from the state variables already mentioned, players can condition their strategy on whether or not their partner provided them a favor at the last received opportunity. Again, for better understanding of the effect of state variables, a probit regression with random effects at subject level is run for each treatment and 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[33], where $nice_{it-1}$ is a binary variable taking value 0 (1) if

[33] Again including 'time' as a covariate does not affect the analysis.

Table 8(a): Marginal Effects for Complete Information: Behind and tied

variable	Low-Slow	Low-Fast	High-Slow	High-Fast
$mode_{t-1}$	0.7866978	0.5646562	0.5529992	0.5509232
	(0.0386)***	(0.03653)***	(0.06317)***	(0.03201)***
$nice_{t-1}$	-0.0648936	-0.0168244	-0.037256	-0.033072
	(0.01767)***	(0.01009)*	(0.01349)***	(0.01055)***
$netfavors_{t-1}$	-0.0145328	-0.0016817	-0.0006784	-0.0027572
	(0.00453)***	(0.00129)	(0.00031)**	(0.00125)**
$timeelapsed_{t-1}$	-0.0004554	-0.001486	0.00000259	-0.0003636
	(0.00024)*	(0.00051)***	(0.00002)	(0.00037)
$ownpayoff_{t-1}$	-0.0002252	-0.0000473	-0.00000350	-0.0000108
	(0.00016)	(0.00005)	(0.00000)	(0.00001)
no. of observations	8187	5230	8426	7587
no. of groups	28	28	28	28
pseudo-R^2	0.716	0.455	0.754	0.409

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

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

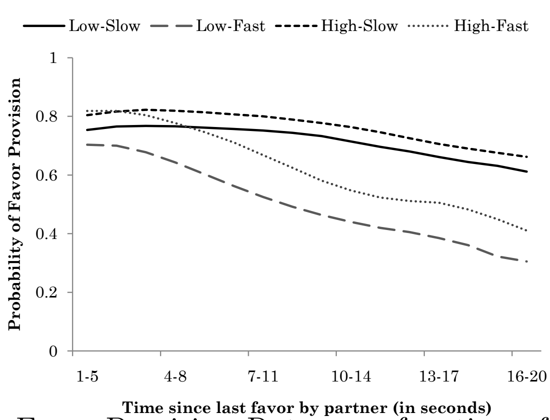

Figure 7: Favor Provision Rate as a function of Time Since Last Favor by Partner - Private Information Treatments

Table 8(b): Marginal Effects for Complete Information: Ahead

variable	Low-Slow	Low-Fast	High-Slow	High-Fast
$mode_{t-1}$	0.9096734	0.8031664	0.8112257	0.6596043
	(0.01207)***	(0.03604)***	(0.03705)***	(0.04533)***
$nice_{t-1}$	-0.134904	-0.0326393	0.0039377	-0.0887263
	(0.03393)***	(0.01985)*	(0.00549)	(0.02413)***
$netfavors_{t-1}$	-0.0070354	-0.0071307	-0.0066474	-0.0041348
	(0.00742)	(0.00226)***	(0.0023)***	(0.00144)***
$timeelapsed_{t-1}$	-0.0008807	-0.0063092	-0.0005093	-0.0003658
	(0.00051)*	(0.00183)***	(0.00021)**	(0.00062)
$ownpayoff_{t-1}$	0.0022634	0.0002482	0.00000912	0.0000313
	(0.00054)***	(0.00009)***	(0.00002)	(0.00001)**
no. of observations	5383	4140	5994	6293
no. of groups	27	27	28	28
pseudo-R^2	0.729	0.638	0.764	0.532

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

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

partner provided the favor (declined to grant a favor) at the last received opportunity[34].

Tables 8(a) and 8(b) collect the results of the probit regression for all the cases. 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 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 complete information than under private information for each treatment. Finally, the effect of 'own payoff' is still inconclusive. This gives rise to the following conclusion.

Result 9. *In the case where the opportunities to provide a favor are publicly observed by both players, 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.*

All the results of this section support hypotheses 4 and 5.

[34]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'.

Table 9: Average Switch Frequencies

Treatment	Pvt. Info. Data	Simple CM	Best CM	Complete 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

E. 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 9. It can be readily inferred that people switch less on an 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 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%). With more opportunities the activity level increases and it becomes possible to forego occasional favors and save on the cost of providing a favor. All these differences[35] are statistically significant with p-values < 0.01 (using robust SEs clustered at each pair). Thus, we have the following result.

Result 10. *When opportunities to provide a favor is 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 10. 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 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 underpredicts the differences quantitatively. Third, although the simple CM is able to correctly predict about the qualitative difference in changing the arrival rate (quantitatively, it overpredicts), 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 complete information environments suggests that monitoring the opportunities received by the partner significantly reduces the volatility of play in all four treatments.

[35]Except the High-Slow versus Low-Slow case which has a p-value of 0.38.

V Conclusion

Analyzing the exchange of favors in the context of a dynamic two player game with private information generated several results. Individuals provide a larger fraction of favors when the return from receiving a favor is higher. However, overall favor provision is invariant to the changes in frequency with which individuals receive an opportunity to grant a favor. Although the level of inequality in payoffs among a pair of individuals is similar to that of the level that would have been generated if everyone followed the best chips mechanism, the efficiency level achieved by the individuals is far lower than the corresponding best chips mechanism in each of the parameter configurations considered. The favor exchange relation is more volatile if opportunities are more frequent and also if benefits are lower.

Next, the study compared the behavior in the case where the opportunities to provide a favor are privately observed to the case where they are observed by both the individuals. Findings suggest that relaxing the informational constraint does help achieve a higher level of efficiency in every situation.

Under informational constraints, the likelihood of favor provision by an average individual declines as the net favors granted by her increases and also as the time since last favor received by her partner goes up. However, when the constraint is lifted, people condition their decision to grant a favor on whether or not their partner provided them a favor at the last received opportunity. Also, probability of favor provision is now less responsive to net favors granted and time since last favor by partner than the private information counterparts. This suggests that there are substantial differences in the behavior of individuals based on whether or not they are able to monitor the opportunities received by their partner.

There are several interesting directions to pursue in future research. First, testing for specific strategies was beyond the scope of the present study. In future research it would be interesting to study in detail about the strategies employed by each individual. One could track the changes in strategies used by individuals over time within a match. It might also be worthwhile to conduct a thorough analysis of learning with more matches under a particular parameter configuration and analyze the evolution of strategies in the current setting. However, as time is continuous, one must be careful in testing for strategies based on state variables. There can be lags in switching from one mode to another while reacting to a change in a state variable[36]. This lag may also vary depending on whether the individual is in a 'Do Favor' or 'Do not do favor' mode. Also, heterogeneity in reaction times among participants might be an important issue to address.

A second line of research might consider reversing the direction of private information about opportunities. The setting in this paper was such that the players never had to ask for favors. The favors just became available when granted by her partner. It might

[36] One could implement $t > 1$ second (instead of 1 second) as the period so as to provide ample time for participants to react to a change in a state variable.

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. That is, do players indulge in cheap talk? These phenomena like signaling and reputation building add more realistic features to a model of favor exchange but at the same time complicates the decision making choices of the players involved.

Third, 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. Nayyar (2009) develops a model of favor exchange that allows for these asymmetries and finds that cooperation among players can be supported in equilibrium even under those asymmetries. It would be worthwhile to see whether this remains true in the laboratory also. The role of partial favors could be important in these situations.

Finally, the interest in experiments in dynamic games conducted in a continuous time setting is very recent. The experimental framework developed in this paper 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[37]. 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 the desired quantities to produce over time for each firm and also whether the monopoly output is produced too often or not. As mentioned earlier, this new framework for experimental studies of dynamic games generates a lot more data than a discrete time analogue by not requiring subjects to decide every period. Similarly, one could study dynamic games where the stage game is not repeated every period, rather, it arrives according to a Poisson process.

References

[1] Abdulkadiroglu, A. and Bagwell, K. (2010), "Trust, reciprocity and favors in cooperative relationships," *Working paper.*

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

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

[37] A slide bar can be displayed on the screen for each participant. Each subject can choose their desired level of production and change the desired level by changing the position on the slide bar.

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

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

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

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

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

[9] 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.

[10] 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.

[11] 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.

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

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

[14] 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.

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

[16] Fiske, A.P. (1992), "The four elementary forms of sociality: framework for a unified theory of social relations," *Psychological Review*, 99(4), 689-723.

[17] 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*.

[18] Friedman, D. and Oprea, R. (2011), "A continuous dilemma," *American Economic Review, forthcoming.*

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

[20] 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.

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

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

[23] Kalesnik, V. (2005), "Continuous time partnerships with discrete events," *Working paper.*

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

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

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

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

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

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

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

APPENDIX A-B: FOR ONLINE PUBLICATION

A. Instructions to subjects

The following is the instructions from one of the sessions in the private information treatment. The only difference in the other sessions under the private information treatment was in the sequence of the treatments. The instructions in the complete information treatment closely follows the private information treatment differing in the informational aspects wherever necessary. An exact copy of the instructions in the complete information treatment is available upon request.

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 4 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 purpose. 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 4 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 4 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 participant 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 3 investment opportunities and the person you are matched with will get 3 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 4 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 4 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 participant 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.

B. User Interface

Figures B1-B3 show the user interface. Figure B1 displays the initial choice screen. Each subject has to choose from either of the two options labelled 'INVEST' and 'DO NOT INVEST'. To keep the experimental language neutral[38], 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 costs them c points but pays 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'[39]. Once all the subjects in the room made an initial decision, the game began in real time.

Figure B2 (Figure B3) displays a sample screen from the play of a game in the private information treatment (complete information treatment). Each subject can freely switch between the two options by clicking on the respective buttons as many time as they wish[40]. The current option selected by the subject is 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. 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. This feature makes the methodology different from a discrete time analogue in this kind of a stochastic setting with jump signals. In a discrete time setting of the same experiment subjects would have to decide every period whether or not to invest after knowing whether they had an opportunity to do so. This not only slows down the experiment but also requires each subject to decide and choose an action every period. Instead just by letting subjects decide the intent or wish to invest gives them the freedom to switch their actions asynchronously. Thus, not only a lot of data can be generated in a short span of time, it is also easier for a subject to play the game in a continuous time setting rather than in discrete time.

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

[39]The current paper focuses on full favors. So, when an opportunity arrives, one can grant either the full favor or grant nothing. One might be interested to conduct an experiment 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 then would represent the fraction of a favor that a subject would wish to grant and would range from 0 (no favor) to 1 (full favor).

[40]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, Casari, Skrzypacz, and Spagnolo (2011) use touch screen PCs in their experiments.

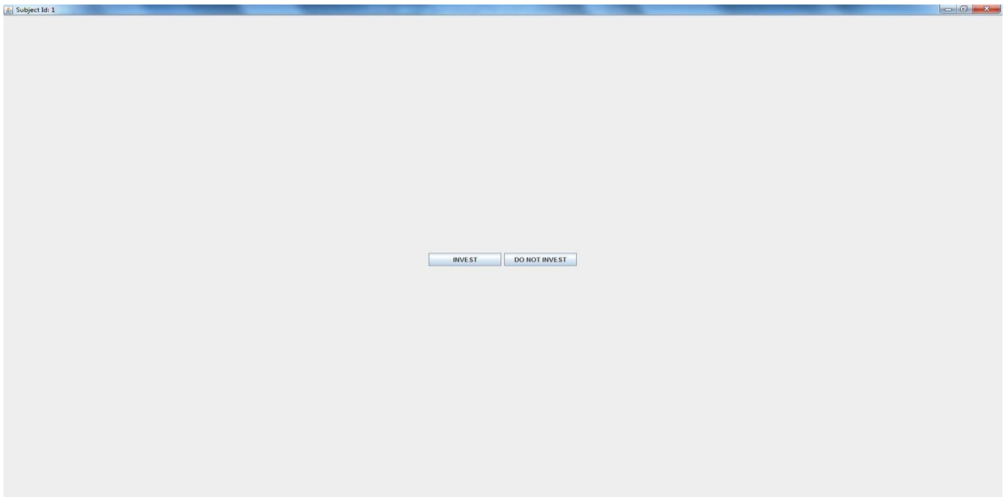

Figure B1: Initial Choice Screen

The screen also had the basic parameter information alongwith the match number. There are two tickers corresponding to the labels ‘You’ and ‘Other’. A black dot represents an investment alongwith the time (in seconds) at which it was undertaken, whereas a white dot stands for an investment opportunity that was not undertaken alongwith the time of the opportunity arrival. 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 complete 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 is still not displayed anywhere on the screen. The only difference in the private and complete information treatments is over the observability of the opportunities and not over the current selected option by partner. Finally, the screen also displayed the real time information on 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 complete 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 a ‘INVEST’ mode when she received it.

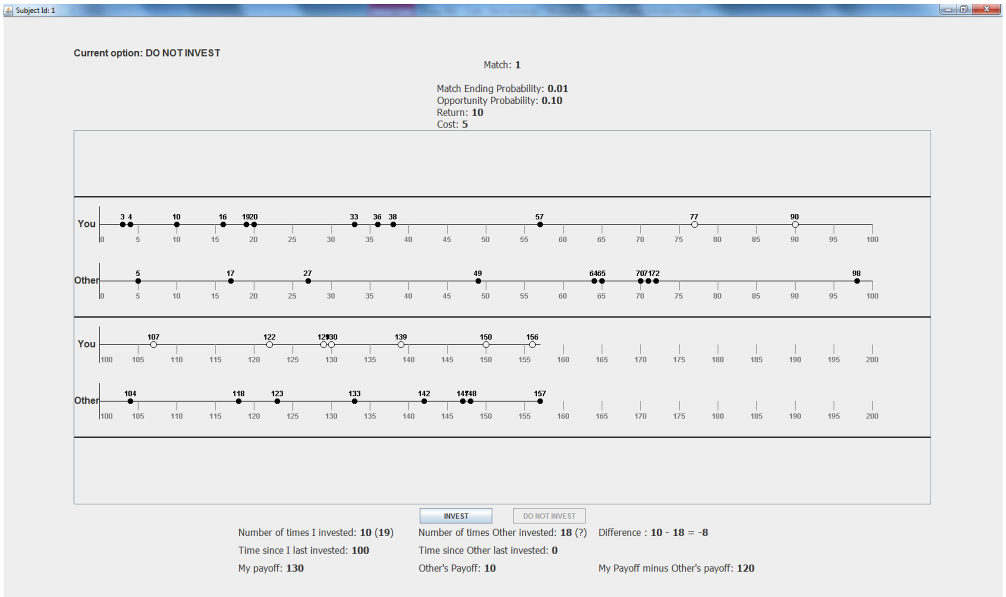

Figure B2: Subject Screen for Private Information Treatments

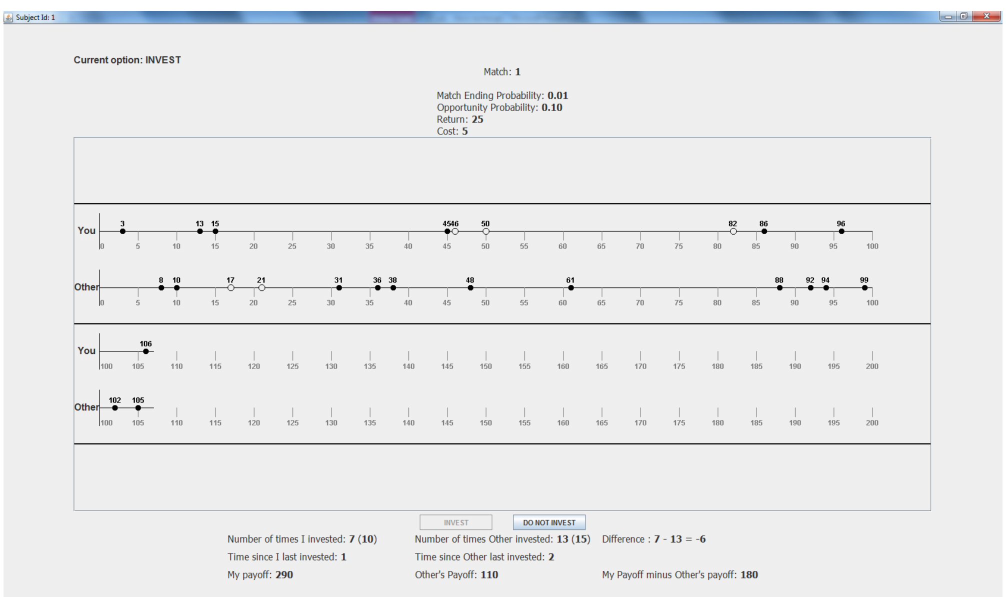

Figure B3: Subject Screen for Complete Information Treatments