Reputation Without Repeated Interaction:

A Role for Public Disclosures

by

Steven T. Schwartz *
Richard A. Young **
Kristina Zvinakis **

July 2000

* Binghamton University
** The Ohio State University
# Contact author:
Department of Accounting and MIS, The Ohio State University, 2100 Neil Avenue, Columbus, OH 43210-1144. Phone: (614) 292-0889. E-Mail: young.53@osu.edu.

This paper reports on an experiment designed to investigate the economic impact of public disclosures. The setting is a modified repeated Prisoner's Dilemma in which subjects are randomly assigned a new, anonymous partner at the start of each round. We demonstrate that public disclosure of past choices increases the frequency of cooperation which, as implied by the Prisoner's Dilemma setting, improves efficiency. We further demonstrate that even untimely disclosures can be effective in inducing cooperation among agents.

The accounting and economics literatures have identified settings where increases in efficiency are achieved through coordination. First, agents may have the ability to credibly commit to their future behavior, perhaps by signing contracts that are (somehow) enforceable. Second, agents may be in a situation where they transact with each other directly and repeatedly, so that future decisions may be made implicitly conditional on past behavior.

While efficiency gains might be achievable through contracting or repeated interaction, there are many instances where contracting is costly or illegal, and where transactions are not repeated among the same set of agents. The contribution of this paper is it demonstrates that, in the absence of contracting or repeated transactions, public disclosures can be useful for increasing efficiency.

The basic structure of our experiment is a modified Prisoner's Dilemma game with randomly paired subjects and a stochastic endpoint. There are three experimental treatments that differ only by the type of disclosure. In the first treatment, no public disclosures are made (*no disclosure*). In the second, a player receives the history of how his or her partner played in previous rounds with a lag (*delayed disclosure*). In the third, a player receives the complete history of how his or her partner played without a lag (*immediate disclosure*).

In order to guide the experiment, we conduct a formal analysis that derives conditions under which the addition of public disclosures allows for Nash equilibria wherein players choose to cooperate in each round of the game.[1] Under *no disclosure*, there exist no equilibria in which players cooperate. In contrast, under both *immediate* and *delayed disclosure*, there may exist equilibria in which players cooperate along the equilibrium path. Cooperative play is enforced by the credible (off-equilibrium) threat that if a player is identified as having defected in the past, future partners will punish him or her for one or more consecutive rounds. However, under delayed disclosure, the conditions under which such equilibria exist are more stringent than under immediate disclosure.

Even with the addition of public disclosures, "always defect" equilibria continue to exist. While the "always cooperate" equilibrium is Pareto optimal, it may nevertheless be difficult to achieve:

> ". . . . , the resulting strategy coordination problem is more difficult than those arising in repeated bargaining and coordination games. Not only are there multiple equilibria in repeated cooperation games, but the equilibria with desirable strategic properties, like efficiency, are constructed using history contingent strategies. Consequently, history must be used not only to focus expectations on an equilibrium assignment, but also to monitor and enforce compliance with the equilibrium assignment." (Van Huyck et al. (1997), p. 3)

Therefore, consistent with Van Huyck et al.'s appraisal of such settings, public disclosures and repetition transform the *cooperation* problem into a *coordination* problem. In response to the coordination problem, players may wish to form *reputations* for choosing the strategy which, when used by all players, constitutes the socially preferred Nash Equilibrium.[2] Repeated play allows others to learn how one will play the game, and hence one can build a reputation for playing a particular way. Because extant theory does not provide definitive guidance as to how players will coordinate their strategies, we adopt an experimental approach to the question.

The results of our experiment show that subjects exhibit a greater frequency of cooperation in the disclosure treatments than in the non-disclosure treatment. After the initial stages of the experiment, delay in the disclosure does not significantly inhibit cooperation. Further, subjects use the disclosures they receive in an intuitive manner. For example, subjects generally cooperate more often with those who have cooperated in the past.

The remainder of the paper is organized as follows. Section 1 discusses background literature. Section 2 presents the model. Section 3 describes the experimental design. Section 4 presents the hypotheses. Section 5 examines the results and Section 6 summarizes and concludes the paper.

1. Background Literature

Because we explore public disclosures in the context of a game similar to a repeated Prisoner's Dilemma *(PD),* we briefly review the theoretical and experimental literature in that area. Given the volume of such work and the heterogeneous nature of the implementation of the experiments, this review is admittedly incomplete.[3]

In a *one-shot PD* game, it is well known that the unique Nash equilibrium is (Defect, Defect). This equilibrium is characterized by lesser joint and individual payoffs than (Cooperate, Cooperate). This equilibrium is somewhat compelling as a predictor of behavior, because it is attained with dominant strategies.

In the *finitely-repeated PD* game, backward induction yields (Defect, Defect) as the unique

Nash equilibrium. However, in the repeated *PD* game with *unknown horizon*, many equilibria may exist, depending on payoffs and the probability the game continues. Equilibria in which cooperation is observed in every round are supported by the threat of punishment in the future for defection in the current round. This threat holds at every point in the game due to its stochastic endpoint. Thus, feedback inherent in a repeated pairings setting allows one agent to hold another accountable for prior defection (Axelrod, 1984; Roth and Murnighan, 1978).

Numerous experiments have explored how individuals play *PD* games (Rapaport and Chammah, 1965; Roth and Murnighan, 1978; Roth, 1988; Dawes, 1990; Dawes and Thaler, 1988). Surprisingly, in these experiments, cooperation in one-shot *PD* games is far from negligible. Many experiments have followed, attempting to identify factors that contribute to this counterintuitive result. Major factors that have been identified include relative payoffs (Butler, 1992) and experience with the game (Andreoni and Miller, 1993). We note that early experimentation often did not provide pecuniary remuneration for the subjects.

Of greatest relevance to our work are experiments on *repeated PD* games. In repeated *PD* games, involving both known and unknown horizons, significant cooperation is found when subjects are continuously paired. However, when subjects are randomly paired and subjects' prior actions are not disclosed, cooperation is greatly decreased.[4]

We adopt the perspective that cooperation in the repeated *PD* is motivated by a desire to focus attention on, and enforce an efficient equilibrium. In this sense, we propose that subjects in our study may wish to develop reputations as "cooperators". However, other behavioral rationales for cooperation have been identified in the literature. For example, Dawes and Thaler (1988, p. 190) note, "a cooperative act itself - or a reputation for being a cooperative person -

may with high probability be reciprocated with cooperation, to the ultimate benefit of the cooperator." Kahneman, Knetsch, and Thaler (1986) and Berg et al. (1995) provide evidence consistent with this proposition. Dickhaut and McCabe (1997) take this idea one step further, conjecturing that the threat of explicit retaliation may not be necessary. They argue that the mere act of recording a steward's exchanges, which they consider central to the accounting discipline, may induce behavioral modifications even in the absence of repeated play.

2. Model

In this section, we present the game used in the experiment. We first consider the benchmark case where there are no disclosures of prior strategic choices. Two identical agents, hereafter referred to as managers, are each confronted with the following choice: invest funds, F, in either a *private investment* or a *joint investment*. F may not be split between the two investments. The private investment earns a return of R_P. The revenue from the private investment is $(1 + R_P)F$, all of which is kept by the investing manager. The joint investment earns a return of R_J, where $R_J > R_P$. The revenue from the joint investment is split evenly between the two managers, regardless of their respective investments. The total revenue from the joint investment is equal to $(1 + R_J)F$ if the joint investment is chosen by only one manager, and $2(1 + R_J)F$ if joint investment is chosen by both managers. The net proceeds to a manager are the revenues they receive minus their cost, F. Figure 1 depicts this game in strategic form.

[Insert Figure 1 about here.]

We further assume that $R_J < 2 (R_P + .50)$, which implies the game is a *PD*, where joint investment is analogous to *cooperation* and private investment is analogous to *defection*. In a single play of such a game one would expect managers to choose Private, since (Private, Private) is the unique, dominant strategy, Nash equilibrium.

2.1 Multi-period Setting - No Disclosure

We consider a multi-period setting where, after each round, play is continued with probability p. At the beginning of a round, each manager has the opportunity to interact with exactly one other randomly chosen, anonymous, member of the population of managers. That is, the managers are re-paired at the beginning of each round.

Similar to the "work alone" option found in Frank (1987), we provide each manager with the option of *refusing to participate* in a joint investment with the other manager.[5] Thus, the manager now chooses one of: {Private, Joint, Refuse}. If *either* manager chooses Refuse, both managers obtain the payoff that would have been received if both had chosen Private. Given the above parameter restriction, the resulting single-period game is a *PD* augmented by a weakly dominated strategy (Refuse). Figure 2 illustrates the strategic form of the single-period game where F = 100, R_P = 60% and R_J = 120%. We conclude with the observation that in this multi-period version of the game all Nash equilibria are characterized by Private or Refuse in every round.[6]

[Insert Figure 2 about here.]

2.2 Multi-period Setting - Disclosure

Consider a similar multi-period setting, but where disclosures of each manager's past choices are made to their current partner before investment decisions are made. Further, suppose the information technology that discloses past strategic choices is not necessarily timely. Formally, at the beginning of round t, each manager observes the prior strategic choices of the other manager for rounds 1 through $t - \delta$, $\delta \geq 1$. The case of $\delta = 1$ corresponds to immediate disclosure.

We present necessary and sufficient conditions in the multi-period disclosure game for the existence of a Nash equilibrium in which only cooperation is observed. For two reasons we propose an equilibrium that describes how the disclosures of past behavior can be used to enforce cooperation. First, the proposed equilibrium seems like an intuitive one, motivated by prior theoretical and experimental studies. Second, proposing an equilibrium lends structure to the experiment and aids interpretation of the experimental results.

It now becomes useful to formalize the players' strategy space. We assume player i's strategy in round *t* is *solely* a function of the disclosure received in that round. Formally, $f_i(\theta)$ is a function that completely defines player i's strategy over θ, the disclosure received. Let s_{it} denote player i's investment choice in round *t* so that $s_{it} \equiv f_i(\theta_t)$, where θ_t = the disclosure received in round *t*.

The proposed equilibrium may be described as follows: players enforce cooperation along the equilibrium path by refusing to participate for *m* consecutive subsequent rounds with partners who have defected in round *t*.

Definition: *Let p =the probability of continuing to the next round,* δ *= the disclosure delay, m = the number of punishment rounds, and t = the current round. The following strategy is defined as an* ***m-punishment strategy*** *in the repeated modified PD:*

(1) a manager ***cooperates*** *if paired with another manager who has not defected in rounds* $t-\delta-m+1$ *to* $t-\delta$,

(2) a manager ***refuses to participate*** *if paired with another manager who has defected at least once in rounds* $t-\delta-m+1$ *to* $t-\delta$.

Figure 3 illustrates the response of other players under an *m-punishment* (*m-p*) *strategy* to a particular deviation by a hypothetical player, M*. Specifically, Figure 3 assumes: (1) M* defects in rounds 1 through δ, and cooperates thereafter and (2) all other players respond as prescribed by the *m-p* strategy.

[Insert Figure 3 about here.]

The Proposition below provides conditions under which *m-p* is a Nash equilibrium and so cooperation is observed everywhere along the equilibrium path.

Proposition. *Define* $f_i^{m-p}(\theta)$ *as the m-p strategy for player i. Define* Φ *as the vector of all N player strategies. Using the notation in Figure 1,* $\Phi = (f_1^{m-p}(\theta), f_2^{m-p}(\theta) \ldots f_N^{m-p}(\theta))$ *is a Nash equilibrium if and only if:*

$$\frac{b-a}{a-c} \leq p^{\delta} \left[\frac{1-p^{\delta+m-1}}{1-p^{\delta}} \right]. \qquad (1)$$

Proof. See Appendix A.

The basic idea of the proof is as follows. We first show that if condition (1) does not hold, the proposed defection strategy shown in Figure 3 is a better response than *m-p* given all other N-1 players are using *m-p*. This proves necessity. We then show that under condition (1) no other unilateral defection strategy can improve upon that proposed in Figure 3, and further that *m-p* beats the proposed defection strategy. This proves sufficiency.

In our experiment the parameters values are $a = 120$, $b = 170$ and $c = 60$, and either $\delta = 1$ (immediate disclosure) or $\delta = 3$ (delayed disclosure). Substitution of these parameters into equation (1) reveals that an *m-p* equilibrium exists when $m \geq 1$ for immediate disclosure and when $m \geq 2$ for delayed disclosure. Under immediate disclosure, only one round of punishment is necessary to enforce an efficient equilibrium. Under delayed disclosure, two rounds of punishment are needed to enforce efficient play.

The model's main role in our paper is to demonstrate that cooperation can be rational behavior within the structure of our experiment. However, when using the model to predict or explain the behavior that occurs in our experiments, care must be taken. First, the model assumes subjects are risk neutral and have (their own) pecuniary remuneration as the only argument in their utility function. It is unlikely that this is completely descriptive of the population in our experiment.[7] Second, we only consider strategies that are contingent on current round disclosures. Evolutionary models of similar settings have found that players' entire experience to each point in the game may be relevant. Given that the primary objective of our research is to explore the effect of public disclosures on the overall level of cooperation and efficiency, it does not seem necessary that we adopt the (complex) evolutionary approach.

3. Experimental Design

In order to study the effect of disclosure on efficiency, we employ a 1 X 3 design. The three *cells* are: no disclosure (ND), delayed disclosure (DD) and immediate disclosure (ID). Each cell consists of three experimental *sessions*. The labeling of the sessions is provided in Figure 4.

[Insert Figure 4 about here.]

Students were recruited from undergraduate business courses at University of Arizona. Students were told the approximate length of the experiment and the range of payments they might expect to receive from participating.[8] All tasks involved in the experiment were computerized. The structure of the game was common knowledge to all participants, and points were determined as in Figure 2. At the end of the experiment, subjects were remunerated with $1 for every 384 points earned, in five-cent increments. There was no show-up fee. However, subjects received additional compensation for a record-keeping task, as explained below. Gross of the record-keeping, expected remuneration from an *m-p* equilibrium was $19.69, while the expected remuneration for the defect equilibrium was $9.84. The difference would appear to be salient. All subjects were paid in cash, privately, at the end of the experiment.

3.1 Non-Disclosure (Cell ND)

Subjects were assigned the role of project manager and were randomly paired at the beginning of each round. There were five training rounds in which subjects' point balances did not count towards cash payments. Subsequently, the subjects played for pay. Figure 5 describes the sequence of events for the three disclosure treatments.

Subjects in each cell were informed at the beginning of the experiment that they would play seven *games*. A game consists of a randomly determined number of *rounds*. Each game lasted for at least one round. For every round reached after the first, there was an 8/9 probability that play would continue for at least one more round.[9] Instructions for Cell ID are found in Appendix B.

Each round consisted of two *stages*. In Stage 1, subjects chose Participate or Refuse. If both subjects chose Participate, they moved to Stage 2. In Stage 2, each subject simultaneously invested 100 points in either Private or Joint. Points could not be split between the two investments. Private earned a return of R_P = 60%, all of which was kept by the investing manager. Joint earned a return of R_J = 120%. All revenue from Joint was split evenly between the two subjects, regardless of their respective investments. If either subject chose Refuse, both subjects received the payoffs they would have received if they had both chosen Private.

Subjects' point balances were updated by their net proceeds (revenue less amount invested) earned that round. These points were converted to cash at the end of the experiment.

[Insert Figure 5 about here.]

3.2 Delayed Disclosure (Cell DD)

In Cell DD, subjects played the same game as described above (Cell ND) except for the following. In Stage 1, prior to deciding whether to Participate or Refuse, subjects received a disclosure of the choices made by their current partner in Rounds 1 through t-3, where t is the current round. The disclosures were either Participate (which identifies a subject who chose Participate but whose partner chose Refuse), Refuse, Joint, or Private. This cell corresponds to $\delta = 3$.

3.3 Immediate Disclosure (Cell ID)

In Cell ID, subjects played the same game as described above (Cell ND) except for the following. In Stage 1, prior to deciding whether to Participate or Refuse, subjects received a disclosure of *all* the prior choices that were made by their current partner. This cell corresponds to $\delta = 1$.

3.4 Record-Keeping

In all three cells, subjects were required to keep a record of specific information they observe on the computer screen. In Cells DD and ID subjects recorded the disclosure information that appeared on their screen. The purpose of this requirement was to slow the pace of play and encourage subjects to consider the information contained in the disclosures.[10] In order to make the cells as comparable as possible, record-keeping was also required in Cell ND. Since no disclosures were provided in Cell ND, subjects instead merely recorded the strategy used by their partner in the prior round. All record keeping took place before subjects made their current investment decisions. To compensate them for record keeping, subjects in Cells DD and ID were paid three dollars in addition to their regular earnings, and those in Cell ND were paid an additional four dollars.

4. Hypotheses

Hypothesis 1 concerns the effect of disclosures on efficiency. In the non-disclosure treatment there are several Nash equilibria, but all are characterized by Private (or Refuse) in each round. Given the parameter values we chose for our experiment, the addition of the investment histories in both the immediate disclosure and delayed disclosure settings allows for additional equilibria not present in the non-disclosure setting, including those where only Joint is observed (*Joint Only*).

Although *Joint Only* equilibria are more efficient and Pareto dominate those equilibria that are payoff-equivalent to *Private Only*, it is not clear that subject play would converge to such an equilibrium. *Private Only* may be considered focal, as it is the dominant strategy solution to the one-shot game. Further, *Private Only* is the maximin strategy in the repeated game. In order to reach a *Joint Only* equilibrium such as *m-p*, subjects must be aware of it and must believe others are aware of it. Mistakes and experimentation by subjects may draw attention away from this equilibrium.

The structure of the payoffs implies the frequency of Joint perfectly proxies for any efficiency measure, since a subject's choice of Joint yields an incremental increase of 60 points in social welfare, regardless of the other subject's choice. Therefore Hypothesis 1, shown in alternative form, is stated in terms of the relative frequency of Joint.

H1: Subjects choose Joint relatively more frequently when an investment history disclosure is present than when no disclosure history is present.

As described in Section 3, the range of parameters that give rise to efficient equilibria is narrower in a delayed disclosure setting than in the immediate disclosure setting. Given the above-noted difficulties in reaching an efficient equilibrium, one might expect these difficulties would be exacerbated in the delayed disclosure setting relative to the immediate disclosure setting. Hypothesis 2 is stated below in its alternative form.

H2: Subjects choose Joint relatively more frequently in the immediate disclosure treatment than in the delayed disclosure treatment.

In theory, in order to reach efficient outcomes using investment histories, subjects must be willing to punish defectors. If all subjects are playing *m-p*, we would observe no defections along the equilibrium the path. However, there are multiple equilibria and we would not necessarily expect convergence to *m-p* play. If subjects are attempting to focus attention on and enforce *m-p*-like play, we would expect investment choices to be a function of the disclosures received. Specifically, we would expect subjects who defect more frequently to receive less cooperation and more refusals. Therefore, Hypothesis 3 and 4 are stated probabilistically (again, in their alternative form).

H3: In both of the disclosure cells, the relative frequency with which a subject chooses Joint in the current round is positively associated with disclosures that their current partner previously chose Joint, Refuse or Participate.

H4: In both of the disclosure cells, the relative frequency with which a subject chooses Refuse in the current round is positively associated with disclosures that their current partner previously chose Private.

5. Results

The relative frequency of investment choices is shown in Table 1, which reveals that Joint occurs most frequently with immediate disclosure and least frequently with no disclosure. Further, Private occurs most frequently with no disclosure and least frequently with immediate disclosure.

[Insert Table 1 and 2 about here.]

In order to observe the effects of subject learning, and the development of play in general, we divide the experiment into three segments: Segment 1 (Games 1 and 2), Segment 2 (Games 3 and 4), and Segment 3 (Games 5 through 7). This partition is somewhat arbitrary, but the results we present are not highly dependent on our choice. Table 2 presents the relative frequency of Joint by session and segment. Within a treatment, there is some variability in the relative frequency of Joint across sessions, but some general observations can be made. In all three ND sessions the relative frequencies of Joint is less than 18% by Segment 3. In both the ID and DD cells, two of three sessions (ID-2, ID-3, DD-1, DD-2) maintain a high level of efficient play throughout the experiment. With the exception of Segment 2 in Session DD-1, the relative frequency of Joint exceeds 40% in every segment in these four sessions. Session ID-1 shows a significant decline in efficiency as the experiment progressed, while DD-3 shows continuously low levels of efficiency. Figure 6 presents a graph summarizing the relationship between relative frequency of Joint and Game.

[Insert Figure 6 about here.]

In order to test Hypothesis 1, observations are classified as either NO-DISCLOSURE (Cell ND) or DISCLOSURE (Cells ID and DD, combined). If strategies were solely a function of current disclosures, as in the theory we develop, observations between subjects but within a session are independent. Therefore, we conduct the main test of Hypothesis 1 using an individual subject as the unit of observation. A Wilcoxon Rank Sum test for differences in location is performed.[11] The results, shown in Table 3, strongly support Hypothesis 1: disclosure appears to increase the frequency of joint investment. In each segment the null hypothesis is rejected at $p < .01$.[12]

We test Hypothesis 2 by conducting Wilcoxon Rank Sum tests for ID vs. DD analogous to those used to test Hypothesis 1. The results are also shown in Table 3. There is support for Hypothesis 2 in the early part of the experiment. However, the medians are not significantly different in Segment 3, when presumably subject learning is substantially complete. Recall, though, that efficient equilibria exist with these parameter values in both the immediate and delayed disclosure settings.

[Insert Table 3 about here.]

Hypotheses 3 and 4 are investigated using Tables 4 and 5, which present the relative frequencies of a subject's investment choices, conditional on the disclosures about the subject's partner. The data is presented separately for Cells ID and DD, and partitioned by segment, except we pool across Segments 1 and 2.[13] We should point out that the support of Hypotheses 3 and 4 does not imply uniform adoption of or convergence to the *m-p* equilibrium.

A cursory inspection of the data would indicate that this clearly has not occurred. Rather, these hypotheses are concerned with whether disclosures are used to monitor and enforce cooperation.

It seems apparent from Table 4 that subjects are conditioning their investment choices on the disclosures received, and doing so in an intuitive manner. Consider the ID cells, Segments 1 and 2, Panel A1. The relative frequency of Joint, when it is disclosed that the current subject chose Private (denoted D in Table 4) in the preceding round, is 28.44%, whereas it is 50.79% if it were disclosed that the current subject chose Joint, Private or Participate (denoted C). In addition, the relative frequency of Private, when it is disclosed that the current subject chose D in the preceding round, is 48.00%, whereas it is 31.01% if it were disclosed that the current subject chose C. This pattern is consistent in both disclosure treatments, in all segments, and for disclosures about the current subject's partner going back as much as three rounds from the most currently disclosed round.

[Insert Tables 4 and 5 about here.]

The response to a given disclosure may depend on the other disclosures. To investigate this question, Table 5 presents the conditional relative frequency of investment on the combined three-round history: Rounds *t-1, t-2* and *t-3* under ID and *t-3, t-4* and *t-5* under DD. Inspection of Table 5 reveals that the reaction to a given disclosure generally does, in fact, depend on the other disclosures received. For example, the relative frequency of Joint conditional on a disclosure of CCC vs. CCD are 42.89% and 28.23%, respectively, indicating that the disclosure in Round *t-3* is conditionally informative given a disclosure of CC for Rounds *t-1* and *t-2*, respectively. The only noticeable exception to this observation is that, in the DD cell, disclosures in Round *t-5* do not appear to be conditionally informative regarding investment choice given a disclosure of CC in Rounds *t-3* and *t-4*, respectively. The reason for such a departure is not readily apparent. Collectively, though, the data in Tables 4 and 5 strongly support Hypothesis 3.

Recall from the theory section that the Refuse option was introduced to give subjects the opportunity to punish prior defections and, at the same time, publicly signal that they are not attempting to exploit their partner. Thus, Hypothesis 4 is derived from the theory. We investigate Hypothesis 4 using Tables 4 and 5. For example, consider Segments 1 and 2 in the ID cell (Table 4, Panel A1). The conditional frequency of Refuse when the *t-1* disclosure is D is 17.78%, whereas it is 6.07% when the disclosure is C. This pattern is consistent in both disclosure treatments, and for disclosures about the current subject's partner going back as much as three rounds. Inspection of Table 5 reveals that, with few exceptions, disclosures for each of the prior three disclosed rounds are conditionally informative given the other two disclosures. For example, in the DD treatment the relative frequency of Refuse conditional on a disclosure of DCC is 7.79%, whereas the relative frequency of Refuse conditional of a disclosure of DCD is 14.71% (Table 5, Panel A2).

[Insert Figures 7 and 8 about here.]

One might be interested in the increment in the relative frequency of Joint and Refuse if the most recent disclosure was **C** vs. **D,** as the experiment progressed. Figures 7 and 8 display the relative frequencies of Joint and Refuse conditional on the *t-1* (*t-3* in DD) disclosure, for each game in the experiment. Figure 7 indicates that the difference in the conditional frequencies of Joint given **C** and given **D** increases through time. What is most striking is that, in general, the relative frequency of both Joint and Refuse is conditional to a greater degree on the most recently disclosed round as the experiment progressed. The apparent exception is Game 7 in DD, which is somewhat deceiving, as there are very few observations comprising this data point. We readily admit that care must be taken in considering the conditional relative frequencies of Joint and Refuse individually, because they would be used jointly by a strategic subject. However, in general, the strength of enforcement of efficient play appears to be increasing over time. If this trend were to continue into the future (hypothetical Games 8 and beyond), this could indicate that cooperation would likely not dissipate as seen in many voluntary contribution games where "contribute only" is not part of a Nash equilibrium. However, with regard to long-term trends, the evidence may best be categorized as suggestive.

We summarize our results as follows. The introduction of public disclosures appears to promote greater efficiency in our setting. Further, the disclosures are used somewhat intuitively, seemingly to monitor and enforce efficient play.

6. Conclusion

This paper describes an experiment that investigates the ability of subjects to form useful reputations in an adapted version of a multi-period Prisoner’s Dilemma game. Instead of repeated play with the same partner, we examine recurring play with a randomly chosen partner. Theoretically, we demonstrate that reputations can still be formed and exploited to facilitate efficient play, when we assume a record of subject’s past strategic choices is disclosed to a subject’s current partner. Empirically, our findings are consistent with the hypothesis that public disclosures increase efficiency. Further, there is evidence that subjects condition their strategies on disclosures in an intuitive way. Subjects who chose to cooperate in previous rounds were more likely to receive cooperation in the current round. Further, subjects in the disclosure cells appeared to begin employing the refuse option in a manner consistent with the theory.

In addition to the role of public disclosures in affecting efficiency, our work bears on the importance of timely disclosures. Often, reports considered reliable through verification (e.g. those produced by an accounting system) are notoriously late. This motivates our manipulation of the timing of disclosure. Within our model, immediate disclosure allows for cooperation to be equilibrium behavior in a broader set of circumstances than does delayed disclosure. Delay may be costly in that it provides weaker incentives for efficient behavior. This may provide an explanation regarding concerns over the timeliness of reporting. However, in our experiment, delayed disclosures appeared to be only slightly less effective than immediate disclosures. Still, this experimental result is consistent with our theory, since the parameter values used in the experiment allow for efficient equilibria in both the disclosure settings. Future experiments may manipulate parameter values in such a way that the existence of efficient equilibria depends on the length of the delay.

Appendix A: Proof of the Proposition

Suppose that every manager, except M*, is using an *m-p* strategy. We show that M*'s best response is to cooperate in every round. Without loss of generality, we consider rounds of the game which begin where M* first defects, which we label as $t=1$. First, assume M* defects in rounds 1 through δ. We first prove that (1) is necessary. If play is as indicated in Figure 3, M* receives a payoff of b for each round 1 through δ, due to the delay in disclosure. For the next $(m+\delta-1)$ rounds, M*'s partners will refuse to participate, yielding a payoff of c in each round. Subsequently, M* partners and M* choose cooperate, yielding a perpetuity from that round forward that pays a. The payoffs are written in the table below.

Strategy	Expected payoffs to M*
Defect δ times consecutively	$b\,(1+p+p^2+\ldots+p^{\delta-1}) + c\,p^{\delta}(1+p+\ldots+p^{\delta+m-2}) + a\,p^{2\delta+m-1}\left(\frac{1}{1-p}\right)$
m-p strategy	$a\,(1+p+p^2+\ldots+p^{\delta-1}) + a\,p^{\delta}(1+p+\ldots+p^{\delta+m-2}) + a\,p^{2\delta+m-1}\left(\frac{1}{1-p}\right)$

Subtracting the *m-p* payoff from the defect payoff, simplifying, and setting greater than or equal to zero contradicts (1). This proves necessity.

We next prove that (1) is sufficient. Consider M* defecting only in round 1. For M* to prefer defecting only in round 1 to cooperating the entire game, the following condition must be true.

$$(b-a) > (a-c)\,p^{\delta}\,[1+p+\ldots+p^{m-1}] = (a-c)\,p^{\delta}\,\frac{1-p^{m}}{1-p} \qquad (2)$$

Now consider M* switching from cooperation to defection in round 2, given that M* is defecting in round 1 (and $\delta \geq 2$). This would provide a net benefit of $(b\text{-}a) - (a\text{-}c)\, p^{(\delta+m-1)}$, which is also positive by condition (2). Similarly, if M* prefers to defect in rounds 1 and 2 rather than always cooperate, M* prefers to defect in rounds 1, 2 and 3 rather than always cooperating (assuming $\delta \geq 3$). By applying the same logic as above, it can be shown that if M* prefers to defect in round 1 rather than cooperate the entire game, then M* prefers to defect in rounds 1 through δ rather than cooperating the entire game. This yields a net benefit of:

$b - a - (a-c)p^{\delta}\left[\frac{1-p^{\delta+m-1}}{1-p^{\delta}}\right] \geq 0$, which contradicts (1). Therefore, under condition (1) M* prefers cooperating the entire game to defecting in any number of rounds to any other strategy.

APPENDIX B: INSTRUCTIONS for Cell ID

INSTRUCTIONS

YOU MAY NOT SPEAK TO ANY OTHER PARTICIPANT ONCE THE EXPERIMENT HAS STARTED.

General. Thank you for participating in this experiment. Throughout the experiment, you will accumulate points that will be converted to cash at the end of the experiment. You will receive $1.00 for every 384 points. The decisions you make will determine your point total (and thus, the cash payment you receive).

Your Task. You will be assigned the role of a project manager. Each period (or round), you will be randomly paired with another participant playing the role of a project manager. Each of you face the same set of choices. First you must choose whether or not to participate with your current partner. If both you and your partner choose to participate, you have 100 points to invest in *one* of two investments: private investment or joint investment. If either you or your current partner choose not to participate that period, you are both automatically invested in the private option. At the end of each period, you will keep the **net proceeds** generated by your investment. Net proceeds are computed as follows:

(Investment + Return on Investment) - Cost of Investment = Net Proceeds

DESCRIPTION OF INVESTMENT CHOICES. Note that you will not know the investment decision of your current partner *prior* to making your investment choice. Similarly, your partner will not know your investment decision *prior* to making his/her investment choice.

Joint Investment. To invest in the joint investment, enter **J**. The joint investment is available to both managers and earns a return of 120%. One of three things can happen: neither manager invests in it, both managers invest, or only one manager invests. Total proceeds from the joint investment are split evenly between both managers, even if only one manager invests in it.

Private Investment. To invest in your private investment, enter **P**. The private investment earns a return of 60%, and you keep all of the net proceeds. Therefore, if you choose to invest in your private investment, you will earn 60 points [(100 + 60) - 100] plus one-half the net proceeds from the joint investment, if applicable.

Summary

Your Investment	Your Partner's Investment	Your Proceeds	Your Partner's Proceeds	Total Proceeds
Joint (**J**)	Joint (**J**)	120	120	240
Joint (**J**)	Private (**P**)	10	170	180
Private (**P**)	Joint (**J**)	170	10	180
Private (**P**)	Private (**P**)	60	60	120
Refuse (**N**)	ANY	60	60	120

Figure 1
(Your net proceeds are in the upper left hand corner; your partner's net proceeds are in the lower right hand corner)

Your Partner's Investment

		Joint	Private	Refuse
Your Investment	Joint	120 / 120	10 / 170	60 / 60
	Private	170 / 10	60 / 60	60 / 60
	Refuse	60 / 60	60 / 60	60 / 60

Your Net Points in a Period. You begin the experiment with 0 points. Your net proceeds for any one period can range from 10 to 170 points. Each period, your proceeds are added to the proceeds you've earned in all previous periods (that is, your points are accumulated).

Disclosure Information. Before play begins in each round you will receive certain information about your current partner. Specifically, you will be informed how they have played the game up until the current round. They will receive the same information regarding yourself. Remember, you are paired with different people each round and therefore the disclosures will be different each round.

Record Keeping. You are required to keep a record of the disclosure information you are given. A record sheet will be given to you at the beginning of the experiment. Each round, including training rounds you must fill out this sheet. An additional three dollars will be paid to you for the proper completion of this sheet.

Duration of Experiment. There are five (5) training rounds wherein your point balances do not count towards your cash payment. After the training rounds, your point total will be reset to zero. The experiment consists of seven games. Each game will last at least one round. At the end of the first round, a pair of dice is rolled. If a 5 is rolled, the game will end. Otherwise, the game continues for at least one more round, after which, another die roll takes place. This continues until a roll of 5 is observed.

Review of Sequence. The following events will occur each period (or round):

1. You are randomly paired with another participant.
2. You observe how your current partner has played previously.
3. You decide whether or not to participate with your current partner.
4. If you both decide to participate you make investment decision (**P** or **J**).
5. You observe your current partner's investment choice and net proceeds in addition to your net proceeds.

Cooperating With the Other Manager. Some subjects have found that in similar situations they do better if they cooperate. You may find the choice of private tempting. Remember both you and your partner face the same situation. You will note, if both you and your partner choose private, you will each earn 60 points, for a total of 120. However, if both you and your current partner choose joint, you will each earn 120 points, for a total of 240. This is the maximum *combined* earnings for you and your partner. REMEMBER, YOU MAY CHOOSE ANY STRATEGY YOU WISH.

Payment. At the end of the experiment, please remain seated until someone has recorded your earnings. You will be called individually to receive your payment, in private. The amount you receive will not be disclosed by the experimenter to any other participant.

Questions. If, at any time during the experiment, you have a question, please raise your hand and someone will assist you.

REMEMBER, ONCE YOU ARE SEATED,
YOU MAY NOT COMMUNICATE WITH ANY OTHER PARTICIPANT.

Notes:

1) In the ***delayed disclosure*** instructions, the Disclosure Information paragraph appears as follows:

Disclosure Information. Before play begins in each round you will receive certain information about your current partner. Specifically, you will be informed how they have played the game up until three rounds prior. That is, in round 6 you will see how your current partner played in rounds 1 through 3. They will receive the same information regarding yourself. Remember, you are paired with different people each round and therefore the disclosures will be different each round.

In the ***no disclosure*** instructions, the Disclosure Information paragraph is omitted.

2) In the ***no disclosure*** instructions, the Record Keeping paragraph is modified as follows:

Record Keeping. You are required to keep a record of the choice made by the manager you are paired with each round. An additional four dollars will be paid to you for the proper completion of the record sheet.

Appendix C: Tables

Table 1: Relative Frequency of Strategies

Treatment	**Total Number of Strategies**	**Relative Frequency of Strategies**			
		Joint	**Private**	**Participate***	**Refuse**
ID	1,826	41.7 %	34.2 %	11.7 %	12.4 %
DD	1,814	36.2 %	47.5 %	7.9 %	8.4 %
ND	2,032	19.2%	70.6 %	4.9 %	5.3 %

ND = No Disclosure
DD = Delayed Disclosure
ID = Immediate Disclosure

* Participate refers to circumstances where the subject chose Participate but the subject's partner chose Refuse.

Table 2: Efficiency Levels Measured by Relative Frequency of Joint Investment

Treatment	Segment 1 (Games 1 - 2)	Segment 2 (Games 3 – 4)	Segment 3 (Games 5 – 7)	Overall
ID-1	52.00%	34.55%	18.29%	27.50%
ID-2	47.37%	64.29%	44.69%	47.93%
ID-3	63.27%	42.44%	48.74%	46.36%
ID-All Sessions	**52.58%**	**44.02%**	**35.49%**	**41.70 %**
DD-1	50.89%	34.03%	41.91%	41.67%
DD-2	52.35%	43.50%	51.88%	48.87%
DD-3	27.08%	21.30%	23.48%	23.54%
DD-All Sessions	**43.43%**	**32.50%**	**35.02%**	**36.20%**
ND-1	12.50%	17.00%	17.00%	14.38%
ND-2	18.10%	11.90%	7.88%	11.52%
ND-3	45.83%	36.40%	13.43%	32.46%
ND-All Sessions	**25.07%**	**22.41%**	**10.80%**	**19.20%**

ND = No Disclosure
DD = Delayed Disclosure
ID = Immediate Disclosure

Table 3: Wilcoxon Rank Sum Test of Differences in Medians

Treatment	Segment 1 (Games 1 - 2)	Segment 2 (Games 3 – 4)	Segment 3 (Games 5 – 7)
Disclosure	3,820	3,745	3,973.5
Non-disclosure	1,230	1,305	1,076.5
p-values	.0001	.0002	.0001
Immediate Disclosure	1,252	1,292	1,095
Delayed Disclosure	828	788	985
p-values	.0485	.0120	0.8983

Unit of observation is subject.
Wilcoxon Rank Sum statistic is presented for each sub-sample.

Table 4: Frequency of Investment Choice Conditional on Disclosure Received

Panel A1:	**Cell ID** **Segments 1 - 2** (Games 1 through 4)							
	Disclosures							
Strategy	*t-1* = **D**	*t-1* = **C**		*t-2* = **D**	*t-2* = **C**		*t-3* = **D**	*t-3* = **C**
Joint	28.44% (64)	50.79% (226)		28.95% (55)	46.20% (170)		35.76% (54)	41.02% (121)
Private	48.00% (108)	31.01% (138)		48.95% (93)	32.61% (120)		40.40% (61)	35.93% (106)
Refuse	17.78% (40)	6.07% (27)		17.89% (34)	7.07% (26)		17.22% (26)	8.81% (26)
Participate	5.78% (13)	12.13% (54)		4.21% (6)	14.13% (37)		6.62% (6)	14.24% (42)

Panel A2:	**Cell ID** **Segment 3** (Games 5 through 7)							
	Disclosures							
Strategy	*t-1* = **D**	*t-1* = **C**		*t-2* = **D**	*t-2* = **C**		*t-3* = **D**	*t-3* = **C**
Joint	14.80% (45)	39.74% (244)		15.71% (41)	33.33% (185)		12.95% (29)	30.61% (150)
Private	49.67% (151)	29.32% (180)		45.21% (118)	34.23% (190)		51.34% (115)	32.24% (158)
Refuse	27.30% (83)	11.89% (73)		30.65% (80)	12.07% (67)		24.55% (55)	16.53% (81)
Participate	8.22% (25)	19.06% (117))		8.43% (22)	20.36% (113)		11.16% (25)	20.61% (101)

Cells contain the relative frequency of each strategy in round *t* given a particular disclosure.
The number of observations is in parentheses.
Disclosure is equal to **D** if strategy chosen by subject's partner in round *t-i* is Private and **C** otherwise.

Table 4: Relative Frequency of Strategy Conditional on Disclosure Received

Panel B1:	**Cell DD** **Segments 1 - 2** (Games 1 through 4)							
	Disclosures							
Strategy	*t-3* = **D**	*t-3* = **C**		*t-4* = **D**	*t-4* = **C**		*t-5* = **D**	*t-5* = **C**
Joint	19.75% (62)	36.73% (119)		17.67% (47)	36.36% (96)		19.63% (43)	35.29% (78)
Private	63.69% (200)	47.22% (153)		62.78% (167)	45.45% (120)		62.10% (136)	46.61% (103)
Refuse	13.06% (41)	4.01% (13)		15.04% (17)	4.55% (12)		15.07% (33)	4.07% (9)
Participate	3.50% (3)	12.04% (22)		4.51% (12)	13.64% (36)		3.20% (7)	14.03% (31)

Panel B2:	**Cell DD** **Segment 3** (Games 5 through 7)							
	Disclosures							
Strategy	*t-3* = **D**	*t-3* = **C**		*t-4* = **D**	*t-4* = **C**		*t-5* = **D**	*t-5* = **C**
Joint	11.65% (31)	38.41% (116)		12.02% (28)	34.12% (87)		11.98% (23)	31.70% (71)
Private	57.52% (153)	36.42% (110)		56.22% (131)	48.43% (98)		57.81% (111)	39.73% (89)
Refuse	24.06% (64)	5.96% (18)		25.75% (60)	5.49% (14)		21.88% (42)	9.38% (21)
Participate	6.77% (18)	19.21% (58)		6.01% (14)	21.96% (56)		8.33% (16)	19.20% (43)

Cells contain the relative frequency of each strategy in round *t* given a particular disclosure.
The number of observations is in parentheses.
Disclosure is equal to **D** if strategy chosen by subject's partner in round *t-i* is Private and **C** otherwise.

Table 5: Relative Frequency of Strategy Conditional on Three-Round History

Panel A1:	**Cell ID All Games**							
	Disclosure History of Partner *t-1, t-2, t-3*							
Strategy	**CCC**	**CCD**	**CDC**	**CDD**	**DCC**	**DCD**	**DDC**	**DDD**
Joint	42.89% (172)	28.23% (35)	36.28% (41)	23.40% (22)	28.85% (45)	19.18% (14)	11.30% (13)	14.29% (12)
Private	24.19% (97)	41.94% (52)	30.97% (35)	46.81% (44)	45.51% (71)	47.95% (35)	53.04% (61)	53.57% (45)
Refuse	6.48% (26)	14.52% (18)	19.47% (22)	26.60% (25)	15.38% (24)	21.92% (16)	30.43% (35)	26.19% (22)
Participate	26.43% (106)	15.32% (19)	13.27% (15)	3.19% (3)	10.26% (16)	10.96% (8)	5.22% (6)	5.95% (5)

Panel A2:	**Cell DD All Games**							
	Disclosure History of Partner *t-3, t-4, t-5*							
Strategy	**CCC**	**CCD**	**CDC**	**CDD**	**DCC**	**DCD**	**DDC**	**DDD**
Joint	38.26% (88)	36.73% (18)	36.84% (14)	20.21% (19)	38.96% (30)	19.12% (13)	17.00% (17)	8.00% (16)
Private	33.91% (78)	48.98% (24)	52.63% (20)	50.00% (47)	48.05% (37)	58.52% (40)	57.00% (57)	68.00% (136)
Refuse	2.17% (5)	2.04% (1)	7.89% (3)	21.28% (20)	7.79% (6)	14.71% (10)	16.00% (16)	22.00% (44)
Participate	25.65% (59)	12.24% (6)	2.63% (1)	8.51% (3)	5.19% (4)	7.35% (5)	10.00% (10)	2.00% (4)

Cells contain the relative frequency of each strategy in round *t* given a particular disclosure.
The number of observations is in parentheses.
Disclosure is equal to **D** if strategy chosen by subject's partner in round *t-i* is Private and **C** otherwise.

Appendix D: Figures

	Joint	Private
Joint	$(1+R_J)F-F$ $= \mathbf{R_J\,F \equiv \mathit{a}}$ $\mathbf{R_J F \equiv \mathit{a}}$	$.5(1+R_J)F–F$ $= \mathbf{[.5(1+R_J)–1]\,F \equiv \mathit{d}}$ $\mathbf{[R_P + .5(1+R_J)]\,F \equiv \mathit{b}}$
Private	$(1+R_P)F+.5(1+R_J)F–F$ $= \mathbf{[R_P +.5(1+R_J)]\,F \equiv \mathit{b}}$ $\mathbf{[.5\,(1+R_J)–1]\,F \equiv \mathit{d}}$	$= (1+R_P)\,F–F$ $= \mathbf{R_P\,F \equiv \mathit{c}}$ $\mathbf{R_P\,F \equiv \mathit{c}}$

Figure 1: Strategic Form Representation of One-period Investment Game

Payoffs in upper left-hand (lower right-hand) corner are for the row (column) player

Game is a Prisoner's Dilemma if $a > c$ and $b > a$

$R_J > R_P$ implies $a > c$ and $2R_P + 1 > R_J$ implies $b > a$

	Joint	Private	Refuse
Joint	120 120	10 170	60 60
Private	170 10	60 60	60 60
Refuse	60 60	60 60	60 60

Figure 2: Strategic Form Representation of Modified Prisoner's Dilemma under Experimental Parameter Values

Payoffs in upper left-hand (lower right-hand) corner are for row (column) player

Round	1	· · ·	δ	δ+1	· · ·	2δ+m-1	2δ+m	· · ·
M*	P	· · ·	P	J	· · ·	J	J	· · ·
Partner	J	· · ·	J	R	· · ·	R	J	· · ·

Figure 3: M* defects δ times in succession and partners play the *m-p* strategy

J = Joint, P = Private, R = Refuse

Treatment	Session	Number of Subjects	Total Number of Rounds
Immediate Disclosure (ID)	ID-1	10	56
	ID-2	10	58
	ID-3	14	49
Delayed Disclosure (DD)	DD-1	8	66
	DD-2	10	53
	DD-3	12	49
Non-Disclosure (ND)	ND-1	10	48
	ND-2	14	62
	ND-3	12	57

Figure 4: Experimental Design

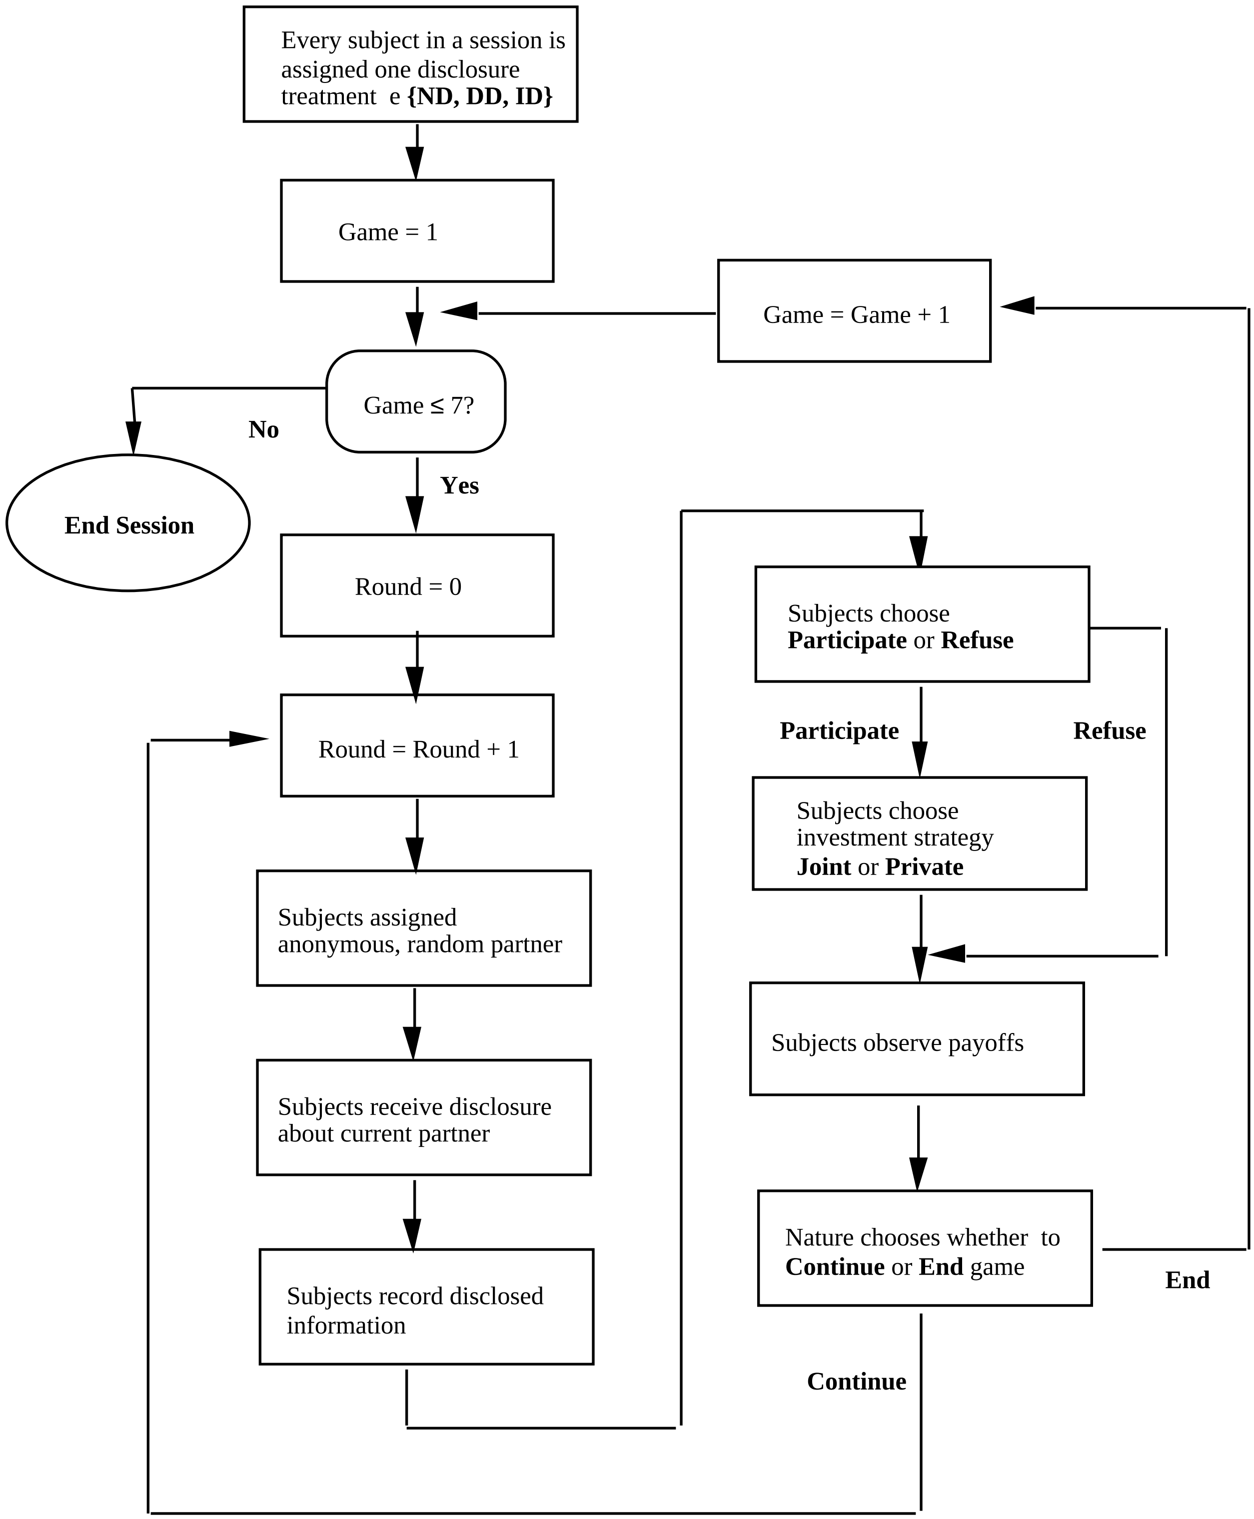

Figure 5: Sequence of Events in One Session

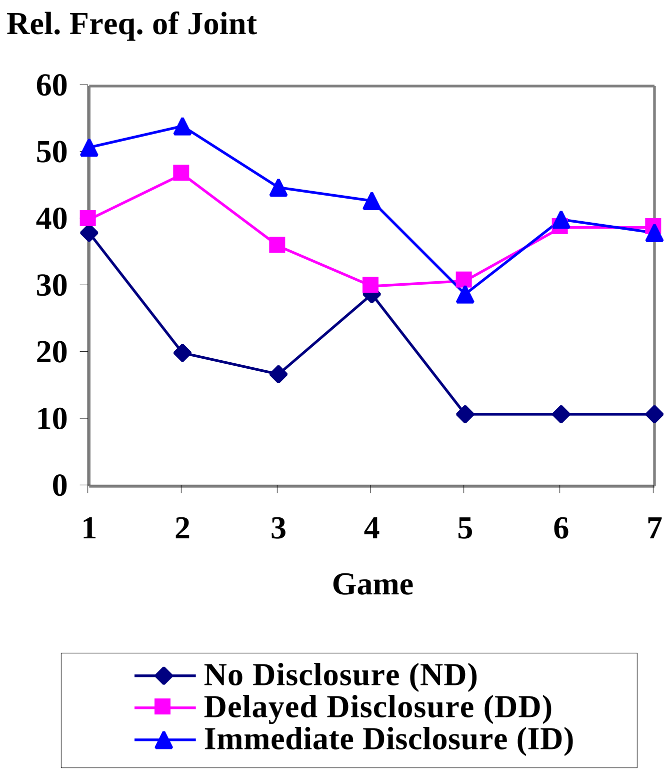

Figure 6: Relative Frequency of Joint Investment by Game (pooled across sessions)

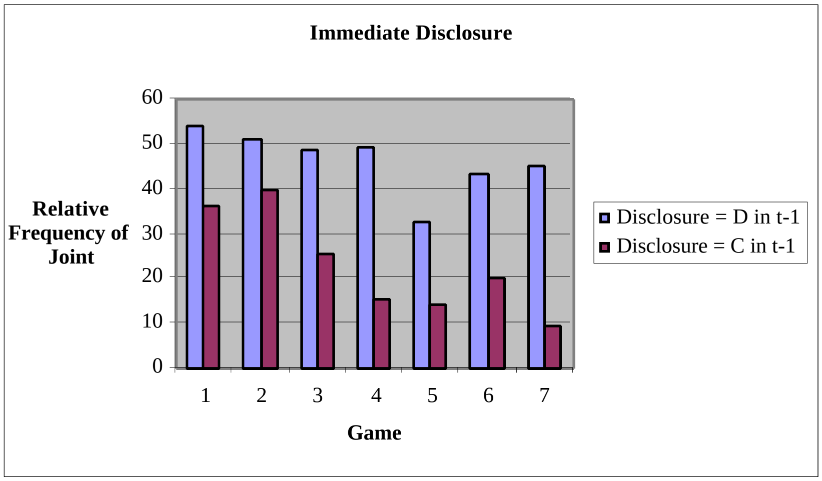

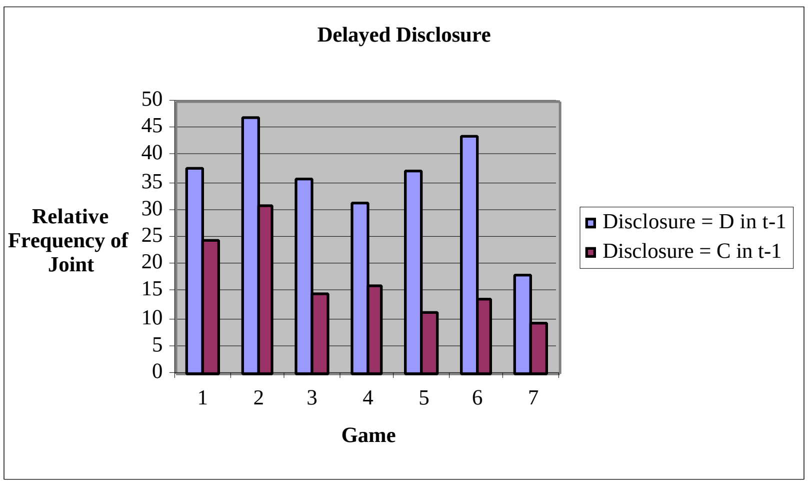

Figure 7: Effect of Disclosure Delay on Joint Investment

Disclosure refers to the disclosed strategy of a subject's partner for round *t-1* in ID and *t-3* in DD. Disclosure is equal to **D** if disclosed strategy is Private and **C** otherwise.

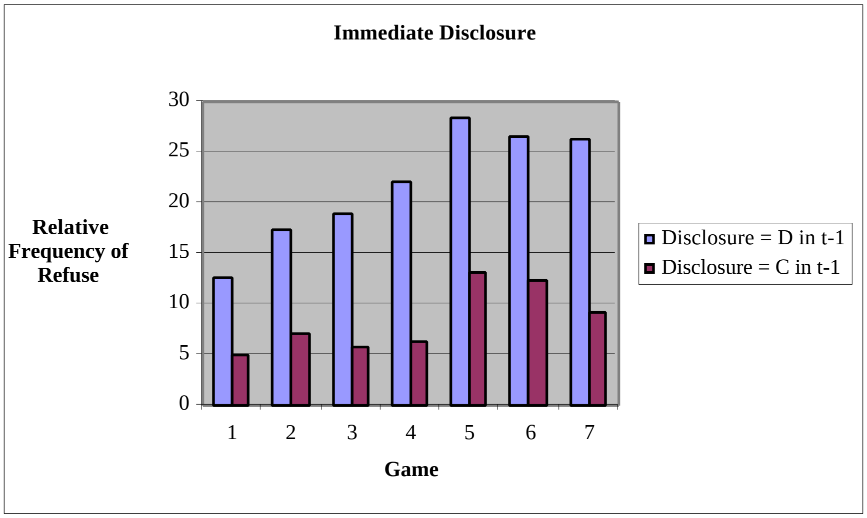

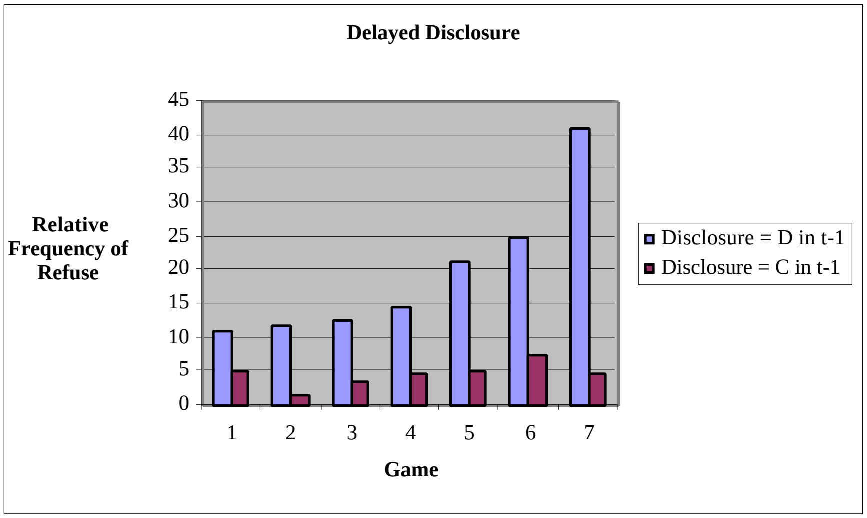

Figure 8: Effect of Disclosure Delay on Refuse

Disclosure refers to the disclosed strategy of a subject's partner for round *t-1* in ID and *t-3* in DD. Disclosure is equal to **D** if disclosed strategy is Private and **C** otherwise.

Acknowledgements

The authors gratefully acknowledge the helpful comments of John Dickhaut and an anonymous reviewer, as well as Greg Allenby, Anil Arya, Jake Birnberg, Harry Evans, Esther Gal-Or, Jon Glover, Vicky Hoffman, Steve Kachelmeier, John Kagel, Dan Levin, Pierre Liang, Kevin McCabe, Don Moser, Nandu Nagarajan, Doug Schroeder, Shiva Shivaramakrishnan, Nicole Thibodeau-Morin, and workshop participants at Arizona, Arizona State and Ohio State. We also are grateful to the University of Arizona for the use of its experimental laboratory.

References

Andreoni, J. and J. Miller. (1993). “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experiment Evidence.” *The Economic Journal.* May, 570-585.

Antle, R. (1982). “The Auditor as an Economic Agent.” *Journal of Accounting Research*. 20 Autumn, 503-527.

Arya, A., J. Glover and R. Young. (1996). “Capital Budgeting in a Multidivisional Firm.” *Journal of Accounting, Auditing and Finance*. Fall, 519-533.

Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.

Berg, J., J. Dickhaut and K. McCabe. (1995). "Trust, Reciprocity, and Social History." *Games and Economic Behavior.* 10 122-142.

Butler, D. (1992). "An Experimental Investigation into the Effects of Uncertainty on Rational Behaviour in Two-person Symmetric Games." *Journal of Behavioral Decision Making*. 5 4, 283-301.

Cooper, R., D. DeJong, R. Forsythe and T. Ross. (1996). “Cooperation without Reputation: Experimental Evidence from Prisoner’s Dilemma Games.” *Games and Economic Behavior*. 12, 187-218.

Dawes, R. M. (1990). “Social Dilemmas.” *Annual Review of Psychology*. 169-193.

______ and R. Thaler. (1988). “Anomalies: Cooperation.” *Journal of Economic Perspectives*. 2 3 Summer, 187-197.

Demski, J. and D. Sappington. (1983). “Multi-Agent Control in Perfectly Correlated Environments.” *Economic Letters.* 13, 325-330.

__________ (1984). “Optimal Incentive Contracts with Multiple Agents.” *Journal of Economic Theory*. 33, 152-171.

Dickhaut, J. and K. McCabe. (1997). “The Behavioral Foundations of Stewardship Accounting and a Proposed Program of Research: What is Accountability?” *Behavioral Research in Accounting*. 9, 1997.

Frank, R. (1987). "If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?" *American Economic Review*. 77, 593-604.

Kahneman, D., J. Knetsch, and R. Thaler. (1986). "Fairness and the Assumptions of Economics." *Journal of Business.* 59, S285-S300.

Kreps, D. (1990a). *A Course in Microeconomic Theory.* Princeton University Press.

______ (1990b). "Corporate Culture and Economic Theory." In *Perspectives on Positive Political Economy*. Edited by J. Alt and K. Shepsle. New York: Cambridge University Press.

______, P. Milgrom, J. Roberts, R. Wilson. (1982). "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma." *Journal of Economic Theory*. 27, 245-252.

Mookherjee, D. and S. Reichelstein. (1992). "Dominant Strategy Implementation of Bayesian Incentive Compatible Allocation Rules." *Journal of Economic Theory*. 60, 161-192.

Rapaport, A. and A. Chammah. (1965). *Prisoner's Dilemma*. Ann Arbor: University of Michigan Press.

Roth, A. (1988). "Laboratory Experimentation in Economics: A Methodological Overview." *Economic Journal.* 98, 974-1031.

______ and J. Murnighan. (1978). "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma" *Journal of Mathematical Psychology*. 17, 189-198.

Van Huyck, J., R. Battalio and M. Walters. (1997). "Is Reputation a Substitute for Commitment in the Peasant-Dictator Game?" Working paper. Texas A & M University. October.

[1] Because we model and conduct an experiment on an *unknown horizon* game, we are able to derive cooperation as equilibrium behavior, assuming fully rational self-interested agents. There is another stream of literature that attempts to explain observation of regular patterns of cooperative behavior in *n-period* finite games (Kreps et al., 1982; Cooper et al., 1996).

[2] In this context, we use the term *reputation formation* to mean a self-interested individual creating the impression that he or she will choose a strategy that belongs to a particular Nash equilibrium. Kreps (1990b) asserts that *ex post* observability of relevant actions is a prerequisite for the formation of valuable reputations: ". . . reputations for behaving in a particular way work more efficiently the more deviations from that behavior are observable" (p. 95). He notes further that entities need not interact repeatedly with the same individuals in order to reap the benefits of a valuable reputation, as long as the entities' actions are made public.

[3] In addition to its relevance to basic research on PD games, our work may bear on the accounting and economics literature devoted to mechanism design. Generally, a "tacit collusion" (i.e., multiple equilibria) problem arises in multi-agent settings. Remedies sometimes involve structuring the contract so as to place the agents in a PD in their subgame (Antle, 1982; Arya et al., 1996; Demski and Sappington, 1983, 1984; Mookherjee and Reichelstein, 1992). Kreps (1990a, Ch. 18) has an excellent illustration of implementation with multiple agents, based on Demski and Sappington (1983). "The incentive scheme in the proof of Proposition 1 places the agents in a Prisoner's Dilemma game wherein any agent

who attempts to misrepresent private information will end up worse off than if he told the truth." It is interesting to note that in much of this literature "cooperation" among the two agents in the subgame would not be socially efficient with regard to the three persons (principal and two agents). In fact, the optimal solutions *rely on* the agents playing non-cooperatively in a PD.

4 See Andreoni and Miller (1993) and Cooper et al. (1996) for a comparison of random and repeated pairing settings.)

5 The Refuse option is important when we introduce disclosures. One concern that may arise is whether the threat necessary to hold the cooperative equilibrium together in the disclosure game is credible. Without the Refuse option, defections would be the only available punishment mechanism. Hence, defection may be interpreted as either exploitation or punishment by a subsequent partner. The inability to discern the current partner's motivation for choosing to defect in the past would eliminate equilibria characterized by cooperation from the set of subgame perfect equilibria. In place of providing the opportunity to refuse, one could disclose the entire history involving a particular player to their partner, including how a particular player *and his partner* had played in each round. This approach would require huge amounts of information and be difficult to operationalize in an experiment.

6 Inspecting Figure 2, the single-period equilibria are: (Private, Private), (Refuse, Private), (Private,Refuse) and (Refuse, Refuse). If prior strategic choices are not made public, a manager's current strategic choice cannot be a function of his/her current partner's past

choices. No credible threat exists to enforce cooperation, so the multi-period Nash equilibria are payoff-equivalent to the repeated, single-period Nash equilibria.

7 In the stage game, subject choices are relevant with respect to the level but not the spread of remuneration. The choice of a multi-period strategy, e.g. *m-p* vs. "always defect" may introduce some risk (an early vulnerability to exploitation as opposed to a maximin strategy). However, given the length of the game, we anticipated that the effect of risk preferences would be negligible.

8 Subjects spent an average of one hour and 15 minutes and earned an average of $16. The range provided to the subjects encompassed the payments they would have received from any strategy for any plausible length of play. This information was provided to help generate student interest in acting as subjects.

9 After one round of play, a pair of dice is rolled. If the outcome is five, the game ends; otherwise, the game continues for at least one more round, after which, another roll of the dice takes place. This continues until a five is rolled. While the time set aside for laboratory use potentially constrains the ability to carry out the game in this fashion, this constraint never was binding. That is, this procedure was not violated in any of the nine sessions.

10 Results from pilot experiments that did not ask for record keeping indicated that subjects cooperated more in the disclosure cells than in the non-disclosure cells. However, there was little evidence the subjects considered the disclosures they received in choosing a strategy.

11 An analysis using the Kolmogorov-Smirnov test for broad differences in distributions was also performed. The results do not vary significantly from those reported herein.

[12] We also conducted a conservative analysis using session as the unit of observation. Given the small number of sessions, three in each treatment, the power of any statistical test is severely limited. Performing a Wilcoxon Rank Sum test on this data, the p-values for Segments 1, 2 and 3 are .0528, .1556 and .0282, respectively. This may still be considered strong support for Hypothesis 1, given the small sample size.

[13] We pool across Segments 1 and 2 in order to preserve data points. We must discard as many as 3 (5) rounds a game in the ID (DD) treatment, given the absence of disclosures in these rounds. This led us to arbitrarily limit our attention to the prior three disclosed rounds.