Can Agents Be Trusted? Experimental Evidence From The Monitoring Game

Yoav Wachsman, Coastal Carolina University, USA

ABSTRACT

Most principals can monitor their agents, but monitoring is usually costly and imperfect. This paper reports the experimental results of a Monitoring Game. In this sequential game, each principal decides whether to monitor an agent that he is randomly paired with and then the agent decides whether to cheat or be honest. Monitoring is costly, but it increases the probability that the agent will get caught. The experiment shows that cheating is commonplace, although most participants do not cheat if they are monitored. Additionally, cheating is more common when neutral terms are used in the instructions.

Keywords: Principal-Agent Game; Trust; Monitoring; Cheating

INTRODUCTION

Employers typically are unable to fully observe some aspects of their employees' behavior, such as effort level. Consequently, employees might have the incentive to shirk or cheat (e.g., steal from the employer). Employers can increase the probability that employees get caught if they cheat by monitoring them; however, monitoring is typically costly to both the employer and the employee.

This paper reports the experimental results of a sequential, perfect-information game called the Monitoring Game. In the Monitoring Game, a principal decides whether to monitor or not monitor an agent that he is randomly paired with and then, upon learning the principal's decision, the agent decides whether to cheat or be honest. If the agent is not caught, then cheating increases her payoff and reduces the principal's payoff. Monitoring is costly to both the principal and the agent, but it increases the probability that the agent will get caught.

The main contribution of this paper to the literature is that, to the best of my knowledge, it is the first paper to examine the strategic interactions between a principal and an agent when the principal can monitor the agent but monitoring is imperfect (i.e., the probability of catching a cheating agent is less than one, even if the agent is monitored). Most research on principal-agent problems investigates the optimal contract that the principal can design and whether participants reach that contract in a controlled environment. However, principals can control other aspects of their interaction with agents, such as whether to monitor them or not. For example, employers can monitor their employees by hiring more middle-level managers, using software to track their employees' activities online, or requiring periodic reports from their employees. These actions, however, are costly to employers. Employees usually also incur a cost from monitoring since monitoring makes their working environment less comfortable and may subject them to additional requirements, such as reporting their activities.

The experiment shows that cheating is common and occurs in higher frequency when agents are not monitored. Furthermore, despite the frequency of cheating, monitoring is uncommon. Finally, using non-neutral terms, such as cheating and monitoring, instead of neutral terms significantly reduces the frequency of cheating but does not affect the frequency of monitoring.

The Monitoring Game is a type of a principal-agent game, a game in which a principal, such as employers, wants his agent (e.g., employee) to behave in a costly way that will benefit the principal, but is unable observe some aspect of the agent's behavior, such as effort. In the Monitoring Game, the principal wants the agent to be honest, but he cannot directly observe whether the agent is cheating or not. Several researchers studied principal-agent

interactions. Marris and Mueller (1980) provide a summary of the early literature on principal-agent dilemmas. Besanko, Regibeau and Rockett (2005) use a theoretical model to illustrate the principal-agent dilemma, while Strausz (2006) compares a principal-agent model under a deterministic and a stochastic environment.

Others examine principal-agent dilemmas in a controlled experiment (Guth et al, 1998; Frederickson and Waller, 2005; and elsewhere). Casadesus-Masanell (2004), argues that external incentives are not always necessary to elicit a desirable behavior out of an employee. In fact, research has shown that reciprocity is common in both controlled and uncontrolled environments. Schechter (2007) finds that farmers in Paraguay often give potential thieves a gift so that they will not steal as much from their farm. Fehr, Kirchsteiger and Riedl (1993) find that most employers were willing to pay above market wages and most employees reciprocate by exerting more effort. **Fehr** and Gachter (1998 & 2000) demonstrate that many people have a proclivity to reciprocate when others act toward them with kindness and to retaliate when they think that they are being treated unfairly. Employees may have an intrinsic desire to be trustworthy based on social norms, ethical standards, and altruism (see Camerer 2003 for review of literature).

Königstein, Kovács and Zala-Mezö (2003) find through surveys that players typically agree on what is considered fair. Büchner et al (2004) utilize a three-person ultimatum game in which one person, a principal, makes an offer on how to split an award, while each of the two responders have a veto power to decide whether the offer is acceptable. They find that mutual trust can be effective in inducing costly but desirable behaviors. Fehr and Schmidt (2004) show that most employees prefer a bonus contract that does not guarantee them a payoff, but implies that they will receive a bonus if they work hard. In their experiment, bonus contracts outperformed piece-meal contracts that paid per unit of effort.

Trust is important in most interactions and is often achieved despite theoretical predictions to the contrary. Ben-Ner et al (2004) find that cooperation is common in a two-part dictator game when one participant dictates how much money she will share with another participant and then their roles are reversed. They find that there is a strong correlation between the amounts sent in the first round and the amounts sent in the second round when the roles are reversed. Trust, however, is common even when there is no opportunity for reciprocity. Berg et al (1995) examine an Investment Game where a first mover can send any portion of $10 to another participant. The amount given is tripled and then the second mover chooses how much money to share with the first mover. They find that first movers send $5.16, on average, and many send $10 even though the unique Nash equilibrium is for the first mover to send nothing. Finally, Eckel and Wilson (2004) find that trust does not relate to one's risk aversion and that some individuals are trusting even if they are paired with a stranger in a one-shot game. Therefore, it is possible that principals in the experiment reported here will trust the agents that they are paired with and not monitor them, even if theory predicts otherwise.

Nonetheless, trust is not universal. Bohnet and Zeckhauser (2004) find that 40% of all participants only trust women, while 9% only trust men in a two-stage trust game. Alesina and La Ferrara (2002) find that minorities that have been traditionally discriminated against and people who had a traumatic experience are less trusting.

The Monitoring Game examines whether employers trust employees in a controlled environment and whether that trust is reciprocated. As the previous experiment shows, individuals are often willing to trust others even when they do not know who they are paired with, and that trust is usually reciprocated. Furthermore, the experiment investigates what affect, if any, does framing have on participants' behavior.

GAME STRUCTURE AND PREDICTIONS

The Monitoring Game is a sequential, perfect information, one-shot game between two players - a principal and an agent. The principal, who moves first, decides whether to monitor (M) or not monitor (NM) an agent; then, upon learning the principal's decision, the agent decides whether to cheat (C) or be honest (H). Once the agent makes her decision, nature determines through chance whether the agent is caught (CA) or not caught (NC).

If the principal does not to monitor the agent there is a probability p that the agent will get caught if she cheats. If the principal monitors the agent there is a probability q that the agent will get caught if she cheats, whereas

q > p. Thus, monitoring increases the probability that the agent would get caught, but it also imposes a monitoring cost of c on the principal and the agent. The principal and the agent can each receive a benefit of b if they cooperate – cooperating means that the principal does not monitor and the agent does not cheat. If the agent successfully cheats (cheats and does not get caught), her payoff increases by t and the principal payoff falls by t. t could represent the amount that the agent steals from the principal. If the agent cheats and is caught, she gets nothing.

The Monitoring Game is illustrated as a game tree in Figure 1. Figure 1 shows both the general payoff for the Monitoring Game and the specific payoffs given the b, c, and t that were chosen for this experiment. In the experiment discussed here, b = 15, c = 3, t = 6, p = ¼, and q = ½. Under these parameters, a rational, risk neutral agent will choose to cheat if she is not monitored and to be honest if she is monitored. The payoffs are in Experimental Pesos (EP), whereas each EP was worth 10 cent at the end of the experiment.

Figure 1: Game Tree for the Monitoring Game
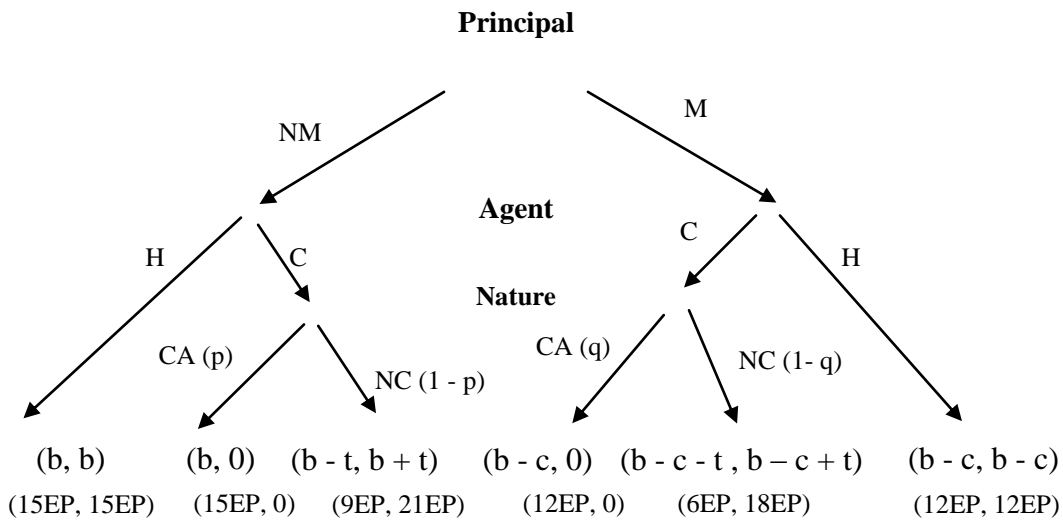
Principal

If the agent cheats, her expected value is $p(0) + (1 − p)(b + t)$ if the principal does not monitor her and $q(0) + (1 − q)(b − c + t)$ if he does. Under the experimental parameters, the agent's expected value from cheating is 15.75 EP if she is not monitored and 9 EP if she is monitored. A rational, risk-neutral agent would choose to cheat if she is not monitored (NM; C) and to be honest if she is monitored (M; H). Thus, if both rationality and risk neutrality are common knowledge, the principal will choose to monitor the agent knowing that it would lead to outcome (M; H) with a payoff of 12 EP for him, which is better than outcome (NM; C) with an expected payoff of 10.5 EP.[1]

As long as the agent is risk neural, or more accurately not very risk loving or very risk averse, and the principal is rational and assumes that the agent is rational and risk neutral, the unique subgame perfect equilibrium for this game is for the principal to monitor the agent and for the agent to be honest (M; H). However, this conclusion might not hold for another set of parameters. For example, for a given set of b, p, and c, if t is sufficiently high or q is sufficiently low the agent will cheat even if she is being monitored. Similarly, for a given set of p, q, t, and b, if c is sufficiently high the principal will opt not to monitor.

The assumption of risk neutrality may not be correct, although, it is an assumption that is commonly made in theoretical models (Cohen, Jaffary and Tanios, 1985). Several studies on risk aversion find that people tend to be risk averse when there is some risk of a significant punishment and are risk-loving when there is some probability

---

[1] The principal expected payoff if he does not monitors the agent and the agent decides to cheat is $p(b) + (1 − p)(b − t)$. Under the parameters of the experiment, the expected payoff is $.25(15) + .75(9) = 10.5$ EP.

(even a very small probability) of receiving a large award (Cohen, Jaffary and Tanios, 1985; Friedman and Savage 1948; and elsewhere). For instance, many people buy lottery tickets even though their expected value from the ticket is negative.

Consequently, for any positive p (p > 0) agents who are sufficiently risk averse will opt to be honest (H) when they are not monitored even if they are only concerned with their own payoff. On the other hand, as long as monitoring is imperfect (q < 1), rational agents who are very risk loving might opt to cheat when they are being monitored since there is some probability that cheating will lead to a higher payoff. Therefore, we can make game theoretical predictions for the parameters of this experiment if and only if we assume that the agent is risk neutral and rational and that the principal is rational and assumes that the agent is risk neutral and rational.[2]

There are three reasons why participants might deviate from the subgame perfect equilibrium (M; H). First, some participants may not fully understand the game; though, the instructions were read out loud and participants completed a couple of exercises to assure that they understood how the game worked. Secondly, as previous experiments suggest, participants may not be exclusively concerned with their own payoff. Particularly, agents who are being trusted (not monitored) might reciprocate by being honest even though cheating gives them a higher expected payoff. Finally, some agent may be sufficiently risk loving and choose to cheat no matter what the principal does or may be sufficiently risk averse and choose to always be honest no matter what the principal does. Similarly, some principals might assume that the agent that they are paired with is very risk averse or very risk loving and, therefore, opt to not monitor.[3]

One weakness of this experiment is that it is not clear if a deviation from the subgame perfect equilibrium is due to unselfish behavior on the part of the participants or the possibility that certain participants are significantly risk averse or risk loving. This problem can be further explored in future experiments by setting p equal to zero.[4] This experiment, however, is designed to simulate a more realistic situation where agents can still get caught even if the principal does not make special effort to monitor them. For instance, an employee may still get caught stealing by her co-workers or managers even if she is not directly monitored. Even though it is not possible to distinguish whether deviations from the subgame perfect equilibrium strategies is due to selflessness or risk aversion, it is probably reasonable to assume that for at least some of the principals the decision not to monitor constitutes an act of trust whereas for some of the agents the decision not to cheat is a way to reciprocate for the principal's trust.

EXPERIMENTAL DESIGN AND PROCEDURES

Participants in the experiment were students from the E. Craig Wall Sr. College of Business Administration at Coastal Carolina University who were solicited from various undergraduate classes. Upon arriving, students were randomly divided into eight "employers" and eight "employees." Although the author refers to one group as employers and the second group as employees in this paper, the results of this experiment are relevant to other types of interactions amongst principals and agents. The two groups were seated facing opposite walls and were not allowed to communicate. Then, the instructions for the experiment were read out loud and the participants were asked to complete a couple of exercises to assure that they understood how the experimental payoffs are calculated.[5]

Each session consisted of eight decision periods. During each period each employer was randomly paired with an employee, but the employer did not know who he was paired with. Each employer was paired with each employee once and only once, much like Andreoni's (1995) strangers treatment. The employers had two minutes to

---

[2] Since the game is a sequential, perfect-information game, it is not necessary for the agent (the second mover) to know that the principal knows that she is rational and risk neutral or to know that the principal is rational and risk neutral. Thus, common knowledge or rationality and risk neutrality is not necessary to solve this game.

[3] If the principal assumes that the agent is very risk averse or risk loving, the principal should choose not to monitor the agent in order to avoid the monitoring cost since in either case the principal's decision to monitor or not will not affect the agent's decision to be honest or cheat.

[4] If p = 0, then theory predicts that even a very risk averse agent will cheat if she is not monitored.

[5] Instructions for the experiment are available online at:
http://dl.dropbox.com/u/10046738/Instructions%20for%20Experiment%20%28trust%29.doc

indicate on a slip of paper whether they were going to or not going to monitor the employee that they were paired with. Monitoring would cost both the employer and the employee he was paired up with three Experimental Pesos (EP), the equivalent of 30 cent.

Each slip was handed to a researcher who inputted the information into an Excel spreadsheet and then distributed it to the employee that the employer was paired with. After an employee discovered whether she was being monitored or not, she had two minutes to decide whether to cheat or be honest. If an employee decided to be honest, she would place a blue chip into an opaque envelope. If an employee decided to cheat, she would place a red chip into the envelope. Then, the envelopes were collected and the chips were discretely placed into the box of the corresponding employers without anyone, including the experimenter, seeing what the chips' colors were. The chip that each employee selected was mixed with three blue chips that were already inside the box. If an employer decided not to monitor his employee, then the researcher would take one chip out of his box. If an employer decided to monitor his employee, then the researcher would take two chips out of his box. After the chips were taken out, the experimenter examined the remaining chips in the boxes to see if each agent cheated or not.

If an employee decided not to cheat, then both players would receive 15 EP if the employer decided not to monitor and 12 EP if he decided to monitor (since monitoring cost both the employer and the employee 3 EP). Cheating would increase the employee's payoff by 6 EP and reduce the employer's payoff by 6 EP if the employee was not caught. The employee was considered caught if she placed a red chip in the box and that chip was randomly taken out. If the employee was caught, she would get 0 EP and the employer would get 12 EP if he elected to monitor and 15 EP if he elected to not monitor. The payments are shown in Table 1. As discussed in the previous section, if rationality and risk-neutrality are common knowledge, then economic theory predicts that employers would always opt to monitor their employees and employees would not cheat.

Table 1: Payoff Table

Employer's Decision	Employee's Decision	Prob. of being caught	Result	Employer's Payoff	Employee's Payoff
Monitor	Cheat	2 in 4 (.50)	Caught	12 EP	0 EP
			Not Caught	6 EP	18 EP
	Not Cheat	0 in 4 (.00)	–	12 EP	12 EP
Not Monitor	Cheat	1 in 4 (.25)	Caught	15 EP	0 EP
			Not Caught	9 EP	21 EP
	Not Cheat	0 in 4 (.00)	–	15 EP	15 EP

Both the employers and the employees were privately informed on a slip of paper what their payoff for the period was. Therefore, agents learned at the end of every period whether they were caught or not and employers learned whether the agent that they were paired with cheated or not (and whether she was caught if she cheated). At the end of the session, all the participants were paid their earnings plus a $5 fee for showing up on time in cash. Final payments were made in private in a separate room.

The experiment was repeated in 8 independent sessions. Four of the sessions were conducted as a No Framing Treatment and the other four sessions were conducted as a Framing Treatment. In the No Framing Treatment, the instructions and the communication throughout the experiment utilized neutral terms. Specifically, in the No Framing Treatment "employers" were called "observers" and employees were called "agents." Similarly, "cheating" and "being honest" were referred to simply as "Action A" and "Action B," respectively. Finally, monitoring and not monitoring was referred to as "Action 1" and "Action 2" respectively. Other than these basic differences in terminology used in the instructions the two treatments were identical.

RESULTS

The Monitoring Game was used to investigate how principals and agents behave when principals can opt to use costly but imperfect monitoring. The experiment also sought to investigate how using a non-neutral

language affects participants' behavior. The experiment yielded several interesting results. The frequency of each outcome is shown by treatment in Table 2 and by session in Table 3.

Table 2: Occurrence of Outcome by Treatment (in percent)

	Treatment		
	No Framing	*Framing*	*Total*
Monitor & Cheat	9.77	2.34	6.05
Monitor & Honest	22.66	34.38	28.52
Not Monitor & Cheat	41.80	27.34	34.57
Not Monitor & Honest	25.78	35.94	30.86

Table 3: Occurrence of Outcome by Session (in percent)

	Session							
	NF1	*NF2*	*NF3*	*NF4*	*F1*	*F2*	*F3*	*F4*
Monitor & Cheat	15.63	6.25	4.69	7.32	11.44	1.56	1.56	4.69
Monitor & Honest	18.75	10.94	39.06	61.04	95.37	53.13	23.44	26.56
Not-Monitor & Cheat	40.63	40.63	29.69	46.39	72.48	31.25	21.88	20.31
Not-Monitor & Honest	25.00	42.19	26.56	41.50	64.85	14.06	53.13	48.44

Result 1 –Principals usually opt not to monitor the agent that they were paired with.

Principals elected to monitor agents 34.6% of the time. Interestingly, only one out of 64 principals in the experiments monitored the agents that he was paired with every single time, even though monitoring the agent was the sub-game perfect strategy (as explained earlier). The median number of times that a given principal opted to monitor agents was 3, the mean was 2.44, and the mode was 4. Figure 2 shows the frequency (number of times out of the eight sessions) that employers opted to monitor the employee that they were paired with. Note that 9 employers did not monitor their employee once throughout the experiment and few employers monitored their employee more than 4 out of 8 times.

Figure 2: Frequency of Monitoring by Principals
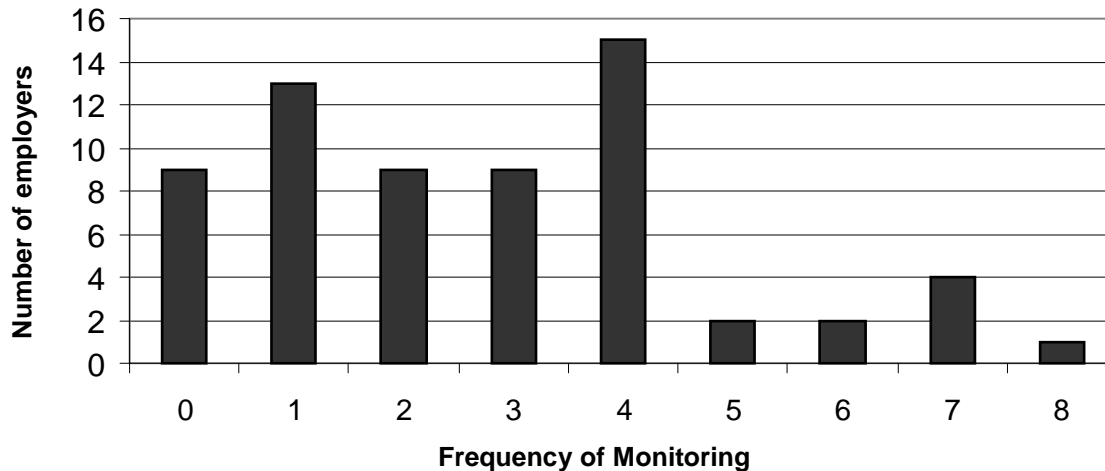
Result 2 – Cheating was fairly common and occurred at a significantly higher rate when agents were not monitored.

There was a considerable amount of cheating in the experiment. Overall, the agents cheated 40.6% of the time. As economic theory predicts, cheating was far less common when agents were monitored because the probability that they would get caught was higher. Agents cheated 52.8% of the time when they were not monitored and only 17.5% of the time when they were monitored. In fact, agents cheated with a lower frequency when they were monitored in all eight sessions. Based on a Wilcoxon-Mann-Whitney rank-sum test, there is only a .4% chance ($p = .004$) that the lower frequency of cheating when the agents were monitored was a random coincidence; therefore, the author rejects the hypothesis that monitoring does not reduce the agents' cheating with a .05 level of significance. This result indicates that agents were more concerned about their expected payoff than about reciprocating the principal's trust.

Interestingly, agents who cheated and got caught were, on average, were 21.4% less likely to cheat the following period than agents who cheated and did not get caught. Thus, agents were less likely to cheat if they were caught the in the previous period even though the fact that they were caught did not affect the probability that they would get caught in the following period since the experiment consisted of 8 independent, one-shot games with consistent probabilities.

Result 3 – Cheating was less common when non-neutral terms were used, but the amount of monitoring was not significantly affected by framing.

Cheating was less common in the Framing Treatment, where agents cheated 29.7% of the time, than in the No Framing Treatment, where agents cheated 51.6% of the time. Furthermore, agents cheated only 6.4% of the time when they were monitored in the Framing Treatment compared to 30.1% of the time in the No Framing Treatment. Framing the instruction also decreased the portion of instances in which agents cheated when they were not monitored from 61.8% to 43.2%.

When the sessions are ordered from the session with the highest amount to the session with the lowest amount of cheating, as shown in Figure 3, the three sessions with the lowest amount of cheating are all Framing sessions. A Wilcoxon-Mann-Whitney rank-sum test shows that there is only a .0143 probability that framing the experiment did not reduce cheating. Therefore, the alternative hypothesis that framing reduces cheating is accepted with a .05 level of significance. Figure 3 shows the frequency of cheating by session.

Figure 3: Frequency of Cheating by Session
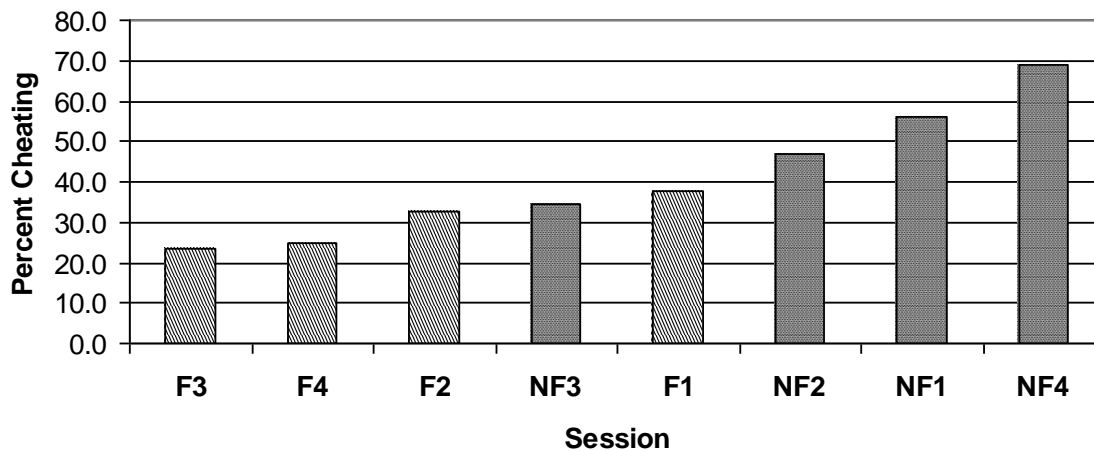
Framing did not seem to significantly affect the proclivity of principals to monitor their agents. Principals opted to monitor their agents 32.4% of the time in the No Framing Treatment and 36.7% of the time in the

Framing Treatment. When the sessions are ordered by frequency of monitoring, a Wilcoxon-Mann-Whitney rank-sum test cannot reject the null hypothesis that framing has no effect on the amount of monitoring.

CONCLUDING REMARKS

The experiment reported here was designed to simulate a typical interaction amongst principals and agents such as employers and employees. Employers in the Monitoring Game chose to either engage or not engage in costly but imperfect monitoring and then the employees that they were paired with chose either to cheat or to be honest. The experiment yielded several interesting results. First, most employers are reluctant to monitor employees even though theory predicts that they would monitor their employees if rationality and risk-neutrality are a common knowledge.

The experiment did not examine the risk neutrality of the participants, which is a very difficult variable to estimate. Therefore, some employers might have opted not to monitor the employee that they were paired with because they assumed that the employee was very risk averse or risk loving and, thus, will always make the same decision no matter what employer did. However, a more likely explanation is that employers were reluctant to monitor because they wanted to avoid the monitoring cost or because they preferred to trust the employees.

Secondly, employees commonly cheated and were more likely to cheat when they were not being monitored. This indicates that most employees do not reciprocate the trust that employers place in them. However, the experiment reported here is composed of one-shot games, whereas actual interaction amongst employers and employees is typically repeated, which makes it easier to build trust.

Thirdly, using non-neutral terms significantly reduced cheating but did not change the frequency of monitoring. Therefore, employees are a less likely to engage in an undesirable behavior when that behavior is labeled as cheating (or another negative term) and are more likely to engage in a desirable behavior when that behavior is labeled as positive. The experiment stresses the importance of correctly labeling desirable and undesirable behaviors, even when labeling the behaviors does not affect the expected payoffs for the individuals involved.

The experiment elucidates the difficult dilemma that employers face. Based on the experiment, employees clearly respond to monitoring by cheating less; nonetheless, monitoring is costly and most employers seem to avoid it. However, the combination of monitoring employees and clearly setting expectations (i.e., framing the interaction) may almost completely eliminate cheating. Cheating occurred only 6.4% of the time when the instructions were framed and the employees were monitored. Thus, clearly setting expectations and monitoring employees can be effective even if monitoring is imperfect.

AUTHOR INFORMATION

Dr. Yoav Wachsman is an Associate Professor of Economics at the E. Craig Wall Sr. College of Business Administration at Coastal Carolina University. His areas of specialty are Microeconomics with Game Theory, Natural Resource Economics and Tourism Economics. Several of his papers investigate reciprocity and cooperation using theoretical models and controlled experiments.

REFERENCES

1. Alesina, A. & La Ferrara, E. (2002). Who trusts others? *Journal of Public Economics*, 85(2), 207-34.
2. Andreoni, J. (1995). Cooperation in public-goods experiments: kindness or confusion? *American Economic Review*, 85(4), 891-904.
3. Ben-Ner, A. et al. (2004). Reciprocity in a two-part dictator game. *Journal of Economic Behavior and Organization,* 53(3), 333-52.
4. Berg, J., Dickhaut, J.W. & McCabe, K.A. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1), 122-42.

5. Besanko, D., Regibeau, P. & Rockett, K.E. (2005). A multi-task principal-agent approach to organizational form. *Journal of Industrial Economics*, 53(4), 437-67.
6. Bohnet, I. & Zeckhauser, R. (2004). Trust, risk and betrayal. *Journal of Economic Behavior and Organization*, 55(4), 467-84.
7. Büchner, S. et al. (2004). Incentive contracts versus trust in three-person ultimatum games: an experimental study. *European Journal of Political Economy*, 20(3), 673-94.
8. Camerer, C.F (2003). *Behavioral game theory: experiments in strategic interaction*. Princeton University Press, Princeton, NJ.
9. Casadesus-Masanell, R. (2004). Trust in agency. *Journal of Economics and Management Strategy*, 13(3), 375-404.
10. Cohen, M, Jaffary, J-Y & Tanios, S. (1985). Individual behavior under risk and under uncertainty: an experimental study, *Theory and Decision*, 18, 203-28.
11. Eckel, C.C. & Wilson, R.K. (2004). Is trust a risky decision? *Journal of Economic Behavior and Organization*, 55(4), 447-65.
12. Fehr, E., Kirchsteiger, G. & Riedl A. (1993). Does fairness prevent market clearing? An experimental investigation. *The Quarterly Journal of Economics*, 108(3), 437-59.
13. Fehr, E. & Gachter, S. (1998). Reciprocity and economics: the economic implications of homo reciprocans. *European Economic Review*, 42(3-5), 845-59.
14. Fehr, E. & Gachter, S. (2000). Fairness and retaliation: the economics of reciprocity. *Journal of Economic Perspectives*, 14(3), 159-81.
15. Fehr, E. and Schmidt, K. M. (2004). Fairness and incentives in a multi-task principal-agent model. *Scandinavian Journal of Economics*, 106(3), 453-74.
16. Frederickson, J.R. & Waller, W. (2005). Carrot or stick? Contract frame and use of decision-influencing information in a principal-agent setting. *Journal of Accounting Research*, 43(5), 709-33.
17. Friedman, M. & Savage, L. (1948). The utility analysis of choices involving risk. *Journal of Political Economy*, 56, 279-304.
18. Guth, W. et al. (1998). An experimental study of a dynamic principal-agent relationship. *Managerial and Decision Economics*, 19(4-5), 327-41.
19. Königstein, M., Kovács, J. & Zala- Mezö, E. (2003). Fairness in a one-principal-two-agents game − a post-experimental questionnaire analysis. *Journal of Economic Psychology*, 24(4), 491-503.
20. Marris, R. & Mueller, D.C. (1980). The corporation, competition, and the invisible hand. *Journal of Economic Literature*, 18(1), 32-63.
21. Schechter, L. (2007). Theft, gift-giving, and trustworthiness: honesty is its own reward in rural Paraguay. *The American Economic Review*, 97(5), 1560-82.
22. Strausz, R. (2006). Deterministic versus stochastic mechanisms in principal-agent models. *Journal of Economic Theory*, 128(1), 306-14.

<u>NOTES</u>