

An Experimental Investigation of Fairness and Reciprocal Behavior in a Triangular Principal–Multiagent Relationship

Alessandro Rossi*, Massimo Warglien†

June 28, 2001



*Department of Computer and Management Science and ROCK (Research on Organizations, Coordination and Knowledge), University of Trento, Via Inama 5, I-38100 Trento; ph: +39-0461-882101, fax: +39-0461-882124, email: arossi@cs.unitn.it.

†Department of Business Economics and Management, Ca'Foscari University of Venice and Cognitive Science Laboratory, University of Trento, email: warglien@unive.it.

Abstract

Issues of fairness in hierarchies have been mostly investigated – both theoretically and experimentally – within dyadic principal–agent relationships. In this paper we consider triangular principal–multiagents structures, integrating vertical hierarchical relationships with horizontal agent–to–agent ones. We explore in the laboratory a game that allows to investigate how principal’s fairness affects cooperation between two interdependent agents performing a simple production task. Our experimental findings show that perceived fairness of principal’s actions may trigger reciprocation in agents’ behavior, affecting how agents play the production game.

JEL Classification: C72, C92.

Keywords: Principal–agent theory, Prisoner’s Dilemma, Reciprocity, Fairness, Experimental economics.

1 Introduction

Issues of equity and fairness in hierarchical organizations have been widely recognized as a key problem since many years by human resource practitioners (see Council [1991] for an review). In the last decade, they have also gained the spotlight in empirical economic research, typically in the frameworks of agency theory and contract design. Such developments have been fostered mostly by experimental studies investigating, in laboratory conditions, the behavioral consequences of alternative types of compensation

schemes and the effectiveness of economic and non economic contract enforcement devices on work effort levels.¹

Much of the evidence gathered by experiments has shown that the behavior of subjects in agency relationships is significantly affected by relative and distributive concerns. Agents seem to take into account the way principals behave and perform systematic comparisons of payoffs. Agents' concerns for principals' fairness result in agents' reciprocating behavior (costly punishment of principals' unfair behavior and costly reward of principals' fair behavior). In turn, principals' are often affected by fairness considerations, offering "fair contracts".

The scope of these analyses is usually restricted to bilateral, "vertical" relationships between a principal and an agent, reflecting the dyadic orientation of both organizational economics and theories of fairness. However, most organizational contexts imply pyramidal, multi-agent structures. In such contexts, "vertical" fairness considerations become often inextricably intermingled with concerns for "horizontal" equity between agents. Studies on team compensation and peer-to-peer working relationships show that such considerations may be of crucial importance in affecting job performances (consider, for instance, the impact of relative evaluation or group incentive schemes, or the effect of information about peers compensation on job performance).

Fairness issues in hierarchies are thus at the crossroads of both horizontal and vertical relations. Nevertheless, very little research has jointly addressed

¹For a comprehensive survey of experimental research on these topics see, for instance, [Gächter and Fehr, 1999, Rossi, 1999].

these two dimensions of interaction. Not only is empirical research on “triangular” principal–agent relationships substantially absent, theory is missing as well. To our knowledge, only a few theoretical studies [Mookherjee, 1984, Itoh, 1994] have developed the principal–agent framework in a multi–agent setting. Triangular features are similarly overlooked by economic theories of reciprocity.²

This paper may be regarded as an exploratory attempt to blend the vertical and horizontal agency relationships in a three–person game. In contrast with previously existing three–person games (such as McCabe et al. [2000], Kagel and Wolfe [1999], Camerer and Knez [1995]), here a hierarchical structure in players role is introduced (one principal and two agent). We explore whether and to what extent fairness in the principal’s behavior affects cooperation between two interdependent agents performing a stylized production task.

In order to make it easier to interpret the experimental results, we have kept the experimental scheme as simple and familiar as possible. The experiment consists in the iteration of a two–stage game. In the first stage the principal decides which share of the pie that will be generated by his agents he will keep for himself and which share will correspondingly be distributed to the agents, according to a piece rate scheme. In the second stage, agents generate the pie by playing a production game in which the relative payoffs of the agents have a Prisoner’s Dilemma structure, but their absolute value is determined by the piece rate unilaterally determined by the principal in the first stage.

Our results highlight that the principal fairness strongly affects agents’

²By the way, see a short discussion in the concluding section of Rabin [1993].

behavior. Kind principals foster mutual cooperation between agents, while greedy ones induce more joint defection.

The paper is organized as follows: next section summarizes the main experimental evidence on fairness in agency relationships and contract design. Section 3 introduces our experiment. Section 4 presents the main experimental results. Some comments on the results and further developments of our research are shortly discussed in Section 5.

2 Fairness and agency: previous studies

The “Gift Exchange Game” is the most extensively studied experimental paradigm for vertical fairness ([Fehr et al., 1993, 1998b, Fehr and Falk, forthcoming, Fehr et al., 1998a, Fehr and Tougareva, 1995]). It is a two-stage game similar to a sequential social dilemma: the first-stage is a wage determination game in which workers (agents) and firms (principals) trade for stipulating job contracts with each other (according to a particular labor market structure); in the second-stage, workers who successfully concluded a contract with a firm must choose an effort level. Theoretical predictions suggest that workers should exhibit minimal effort levels, no matter which wage they receive. Firms, anticipating this, should respond by paying the competitive (zero rent) wage corresponding to the minimum effort level. Experimental findings, however, show that average wages are substantially above the competitive wage corresponding to the agent’s minimum effort, and agent’s effort levels are higher than the minimum. Moreover, workers’s wages contain substantial amounts of rent (wages are much higher than the

competitive wage corresponding to the workers' observed effort levels). These results seem to suggest that principals actually do take into account fairness motives when offering a contract to agents, and that agents react to fair wages showing working efforts higher than the minimum level. Further studies have then investigated to what extent the reciprocal attitude of principals and agents may be able to mitigate the *contract enforcement problem* [Fehr and Gächter, 1998, Fehr et al., 1997]. Behavioral evidence confirms that reciprocal motives within the agency relationship may be regarded as a successful device in raising effort levels above the Nash equilibrium ones [Fehr and Gächter, 1998], and that reciprocity alone may be more effective than many traditional contract enforcement devices such as incentive contracting, fines and monitoring [Fehr et al., 1997].

More standard agency settings have also been recently explored. Keser and Willinger [2000] implemented in the laboratory a standard textbook principal–agent game with hidden action. In this setting, theory predicts that the principal should be able to have the agent accepting a contract that gives all the rent to the principal and leave to the agent the efficiency salary. Experimental evidence, on the contrary, shows that principals' offers are, on average, much more fair than the predicted ones, because principals seem to anticipate that (as in the ultimatum game) unfair offers may be rejected by agents.

Finally, Anderhub et al. [1999] investigate a principal–agent game with no hidden action and deterministic profit function where the agent's contract consists in a fixed component (base pay) and a return share on firm's profits. They show that agents tend to reject unfair contracts and that fair contracts

are reciprocated (efforts level are higher than the optimal ones conditional to the accepted contract).

3 A triangular Principal–Multiagent game

3.1 The model

Consider this simple production setting with one principal and two agents: the two agents are involved in a simple production task where each of them has to decide on the allocation of his working effort. More precisely, each agent has to decide whether he is going to help (or collaborate with) the other agent or not. The decision of one agent affects both his production level and the one corresponding to the other agent: while helping efforts of one agent increase the other's production level, on the other side they decrease the agent's own production amount. Moreover, if both the agents decide to provide help, they are both better off (with respect to their production levels) but, regardless of what the other agent is going to do (providing help or not), the production level for one agent is always higher when he is not providing help because he can concentrate more effort on his own production task.

Produced units are placed in a market with excess demand by the firm owner, the principal. Without loss of generality we can assume that each produced unit is worth 1 experimental currency unit for the principal. He is the residual claimant of the value of units produced by the two agents.

Agents' remuneration is governed by a simple piece rate rule, whose rate per unit is identical for the two agents and is arbitrarily decided by the principal.

Agents cannot decide to terminate the contract with the firm (this means that no side market option is introduced in the model and the participation of workers in the firm is not investigated here).

The described production task is modeled as a two-stage game, that runs as follows: in the first stage the principal publicly announces which share $1 - (W/100)$ of the output value (that will be produced in the second stage by the two agents) is to be attributed to himself as his own payoff in the round. Alternatively, one can interpret $W/100$ as the piece rate, the per unit remuneration assigned to agents by the principal.³ The domain for W is any integer number between 1 and 100. In the second stage, then, each agent has to decide between two alternative strategies (H and L) that result in different individual output values, as shown in Figure 1.

(q_1, q_2)	H	L
H	60, 60	10, 70
L	70, 10	20, 20

Figure 1: The firm production function: relationship between agents’ decisions and agents’ individual output levels.

This structure of output may be thought as the simplest way to model task interdependency of two agents in a production setting: if they both choose to help each other (strategy H) they are both better off, while restraining from helping the other agent (strategy L) is the dominant strategy (for an agent concerned in maximizing his own production level).

³Language in subjects’ instructions was kept as neutral as possible and we explicitly avoided terms as “piece rate” or “remuneration”.

Hence, the agents' relative payoff structure in the game clearly recall a prisoner's dilemma game where absolute payoffs depend on W as depicted in Figure 2.

Payoff(A_1), Payoff(A_2)	H	L
H	$(W/100)60, (W/100)60$	$(W/100)10, (W/100)70$
L	$(W/100)70, (W/100)10$	$(W/100)20, (W/100)20$

Figure 2: Agents' payoff conditional to piece rate W and agents' behavior.

Finally, the principal's payoff, depending on the agents' strategies, is determined as in Figure 3.

	H	L
H	$(1 - W/100) \times 120$	$(1 - W/100) \times 80$
L	$(1 - W/100) \times 80$	$(1 - W/100) \times 40$

Figure 3: Principal's Payoff conditional to piece rate W and agents' behavior.

The game extensive form is reported in Figure 4. Using a standard backward induction argument it is clear that, no matter what the principal decides in the first stage of the game, in the second stage the agents should restrain themselves from helping the other, since they face a standard prisoner's

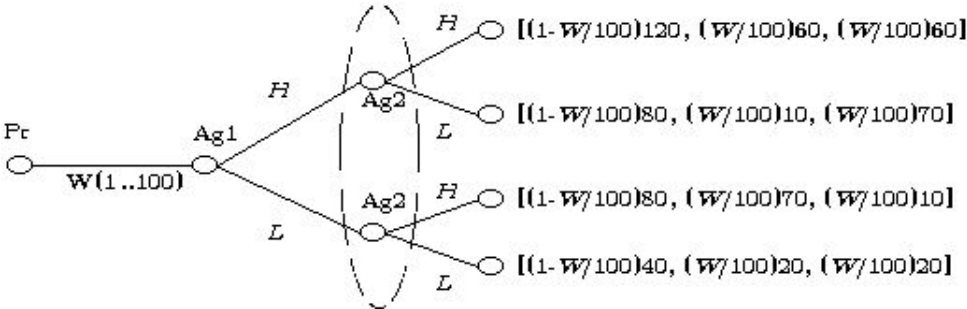


Figure 4: The game extensive form representation.

dilemma game (whose payoffs are a linear transformation of the production levels of Figure 1). Thus, a principal anticipating this, should decide to retain for himself, the largest possible share of the pie. Hence, the unique Nash equilibrium for the one shot game is for the principal to choose $W = 1$ and for the two agents to play strategy L .

3.2 The experimental design

The experimental design consists in a series of two experiments that were played sequentially by a population of 54 college undergraduates recruited at the University of Trento (Italy) during July 1999 (30 of them were undergraduates in Economics). Subjects were recruited through announcements on bulletin boards in the Faculty of Economics and were asked to show up at the Computable and Experimental Economic Laboratory. The announcements claimed that subjects would have been engaged in an experiment lasting about 1 hour and would have been able to gain up to a maximum of 50000 Italian liras (approximately equal to 25 US dollars). During the experiment subjects earned experimental points that were finally converted in Italian liras at the rate of 15 Italian liras per experimental point and were paid to the subjects. In addition, all subjects received a 10000 Italian liras (approximately equal to 5 US dollars) show up fee. The exchange rate was known in advance by all subjects. Their average final payoff was of about 34000 Italian liras (approximately equal to 17 US dollars) for subjects in the role of principals and of about 16000 Italian liras (approximately equal to 8 US dollars) for subjects in the role of agents,

amounts which seemed more than sufficient to motivate them during the experiment.

The subjects were randomly divided in groups of 3 who remained anonymously grouped during the entire experiment; the role of principal or agent was also randomly assigned. Subsequently, subjects were seated in front of computer terminals. After that an experimental administrator had read the experiment instructions⁴ and answered aloud to any question,⁵ the experiment begun. Interaction between subjects were reduced to the minimum during the experiment: each subject could see some two other participants in the room but not their terminal monitors and verbal communication was not allowed at all. Since one group could finish the experiment earlier than the others, participants were asked to remain quietly seated at their desk and to fill a payment form needed for the payment of the experiment.

Each of the two experiment consisted in the repetition of 15 identical rounds of the game presented in Section 3.1. The number of repetitions were considered a reasonable length of time to allow learning to take place (if any was to occur). Each round was thus organized:

- *First stage.* The subject in the role of the principal is asked to type a number between 1 and 100, corresponding to the value to assign to variable W ;

⁴A translation from Italian of instructions is given in Appendix A.

⁵Each subject was revealed his role in the experiment, the principal role or the agent role, only after all questions were answered, so that, in asking questions to the administrator, subjects could not signal to other participants their role.

- *Second stage.* Each of the two subjects having the role of agent are communicated the value of variable W and are asked to choose between strategy H or L ;
- *End of round.* Each of the three subjects is given full information on decisions taken and payoffs earned by all the participants in the group.

At the end of the 15 rounds, subjects were told that they had to participate to another experiment (experienced treatment), where groups were randomly reshuffled while everyone kept the role held in the previous experiment (novice treatment).

The total payoff of each subject at the end of the experiment was then equal to the sum of the payoffs earned by the subject during the 30 rounds, plus the show up fee.

4 Experimental Results

Figure 5 plots overtime the average piece rate W chosen by principals and the observed frequency of agents helping each other (henceforth “cooperating”) in each of the two treatments. Table 1 presents some summary statistics on piece rates and on cooperation rates for each session.⁶

The equilibrium prediction is fulfilled in a relative low number of observations (around 12% in the novice treatment and 20% in the experienced treatment). The time series of piece rate do not show any

⁶In session 5, novice treatment, the data file was accidentally overwritten.

Table 1: Descriptive Statistics per session

			novice treatment	experienced treatment
session	groups	subjects	av. W	coop. rate
1	4	12	23.5	0.492
2	4	12	20.5	0.375
3	4	12	36.7	0.533
4	3	9	23.6	0.555
5	3	9	26.6	0.411
ALL	18	56	26.3	0.484

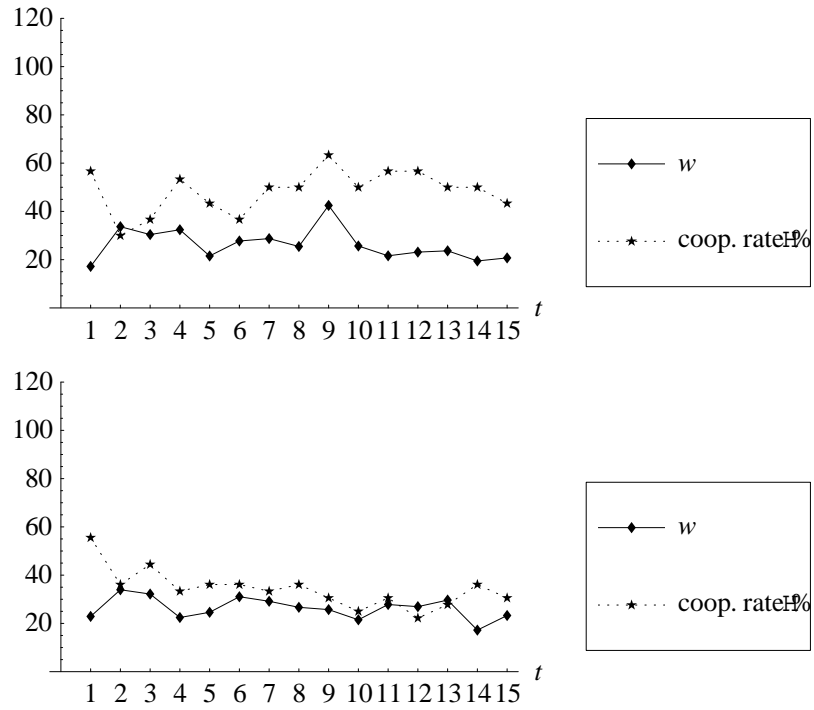


Figure 5: Average piece rate W and average cooperation rates for the novice (upper plot) and the experienced (lower plot) treatments.

significant trend towards the equilibrium in both treatments. Both treatments presents also the same average piece rate value ($\bar{W} = 26.3$).

The cooperation rate exhibited by the agents (as shown in Figure 5) appears to be slightly decreasing overtime in both novice and experienced treatments: cooperation rates decline from 56.6% (first period) to 43.3% (last period) in the novice treatment and from 55.5% to 30.5% in the experienced treatment. Statistical analysis although shows that in both cases the decrease is not significant (randomization test, $\alpha = 0.05$).⁷ The average cooperation rates are, respectively, about 48% in the novice treatment and about 34% in the experienced treatment.

Behavior in both treatments somehow contrasts with the typical pattern of declining cooperation rates that is observed overtime in most experiments on iterated prisoner's dilemma [Andreoni and Miller, 1993], team production Nalbantian and Schotter [1997] and public goods provision games Isaac et al. [1985]).

As it is reasonable to expect, a close investigation of data reveals that the experimental behavior of agents is heavily affected by the behavior of the principal. In accordance with common sense, but contradicting equilibrium predictions, the level of cooperation between agents responds to the kindness of the principal. When the principal increases the piece rate, agents do not decrease the joint production level in 84% (81%) of the observations in the novice (experienced) treatment, and when the principal decreases the piece rate, agents do not increase the joint production level in the 82% (91%) of

⁷Although In the experienced treatment this result may depend from the limited number of independent observations.

the observations in the novice (experienced) treatment.

This pattern of behavior may be viewed, in both novice and experienced treatments, in Figure 6. When the principal tends to choose values of W near the equilibrium, agents coordinate on (L,L) in the majority of observations. As the principal selects higher values of W , more pairs of agents tend to coordinate on (H,H) (in this case, all three subjects are better off than in equilibrium). Thus, principal's fairness matters and affects the mutual relationships of agents. For values of W lesser than 65-75, there is a neat monotonic mapping from the piece rate W to the cooperation rate achieved by the agents. Notice, however, that for high levels of W the correlation between piece rate and output breaks down. This effect may be not significant since there are very few observations in the right tail of the histograms plotted in Figure 6 ($W > 70$ only in 8% (4%) of the observations in the novice (experienced) treatment). Moreover, as the history of individual sequences of runs reveals, most "ultrafair" piece rates have been offered by principals after a sequence of highly unfair moves – which might impair their effectiveness.

The overall link between agents' performance and principals' fairness is significant for both the novice and expert populations of players: the Spearman rank order correlation coefficient is equal to $r_s = 21\%$ for novices and equal to $r_s = 36\%$ for experts (both significant at the $\alpha = 0.001$ level). The increased rank order correlation between the treatments is due to a large extent to the increased frequency of (L,L) responses to low piece rates. Thus, it seems that players have learned a strategy of coordinated reciprocation to the principal unfair moves. This suggests that the impact of fairness

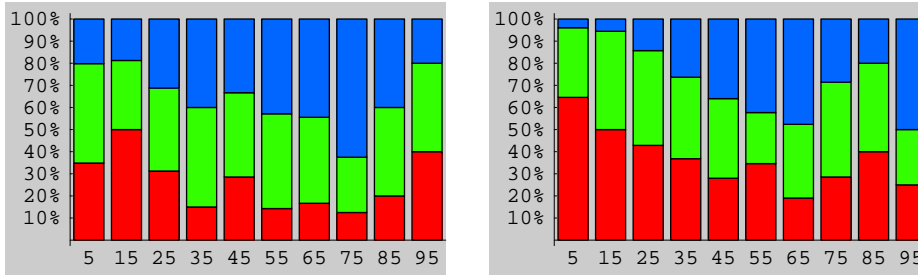


Figure 6: Histogram plot of the frequency of productivity of agents corresponding to different classes of piece rate W decided by the principal in the novice (left side) and in the experienced treatment (right side). Top, middle and bottom bars are respectively related to the frequency of observations where both the agents cooperate, one agent only cooperates and the other one defects, both agents defect.

considerations is not a temporary phenomenon, to be dissolved by a better understanding of the game structure.

One may argue that the link between the level of piece rates W and cooperation rates could find an alternative explanation to the one suggested here. More precisely, it could be argued that the behavior of agents may be strongly affected by absolute size effects of the payoff structure, rather than by fairness considerations. In other words, it could be that increases in cooperation rates may be just due to increases in the scale of the payoff structure, regardless on any perception of fairness. If this was the case, substituting the principal with an automated random device governing the piece rate W , should not change the observed agents' behavior.

A control treatment was then designed to test this competing explanation, based on size effects. The control treatment consists in a two-person

prisoner's dilemma game corresponding to the production subgame of the basic treatment. Payoffs for the two agents of the control game are, once again, the ones depicted in Fig. 2, but this time agents a computerized device, rather than another experimental subject, chooses the value of W sampling at random from the uniform distribution $U \sim [1, 100]$ (it is common knowledge between the agents that values of W are randomly selected by an artificial device).⁸

The results of the control treatment are shown in Fig. 7. The comparison of Fig. 7 with Fig. 6 clearly shows that, while the experimental outcomes of a prisoner's dilemma game are not invariant to scale effects in the payoff structure, these effects of size on cooperation rates go in the opposite directions to the ones exhibited in the baseline treatment. The control treatment shows that, as the magnitude of the payoffs increases, more and more agents defect, while cooperation is a more frequent outcome when payoffs are smaller. Thus, in this treatment subjects seem to be very sensitive to the absolute temptation to defect, despite the fact that the relative payoff

⁸The experimental design closely followed the one described for the baseline treatment, with the following differences: 36 first year undergraduate students (with no previous knowledge of game theory) were recruited; subjects were divided in three cohorts of 12 participants and then randomly matched in pairs; no second experience treatment was run, since we didn't want to test for the role of experience; experimental points were converted at the rate of 40 Italian liras per point, and subjects earned on average 22000 Italian liras (approximately equal to 11 US dollars) for an experiment lasting, on average, around 35 minutes. Each round of the experiment run as follows: in the first stage the computer program extracts the random value of W and sends it to the two agents, in the second stage each of the two agents plays a prisoner's dilemma game with the payoff depicted in Fig. 2.

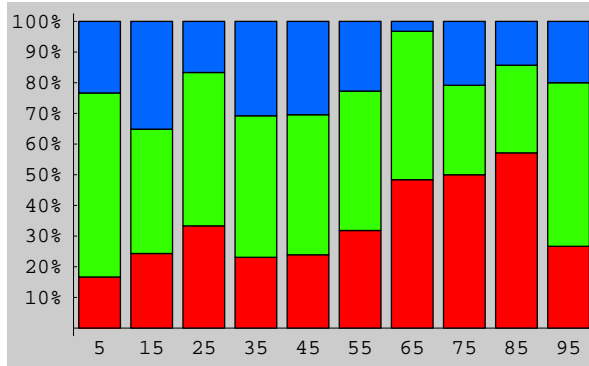


Figure 7: Histogram plot of the frequency of cooperation of agents corresponding to different classes of the random scale value W in the control treatment. Top, middle and bottom bars are respectively related to the frequency of observations where both the agents cooperate, one agent only cooperates and the other one defects, both agents defect.

structure of Fig. 2 is such that the additional gain from defection is relatively small compared to the payoff corresponding to mutual cooperation.

This result strongly corroborates the hypothesis that agents' perception of principal's fairness, and not scale effects in the payoff structure, accounts for the pattern of cooperation displayed by agents in the baseline treatment.

Despite the tendency of agents to retaliate against unfair piece rates, the data shows persisting high levels of unfairness in principals' behavior (see Figure 8). If for example one takes as a benchmark usual laboratory behavior in bargaining games [Roth, 1995], the behavior of principals in our experiment seems unusually greedy. In about 50% of cases, principals take as much as possible. Actually, our principals' behavior bears more resemblance to that of players of a dictator game. In part, this may be explained by the persistence of some miscoordination in agents' responses, that makes

unfairness paying off on average. Given the observed distribution of agents' responses, taking as much as possible is still the best move for principals even in the experienced treatment (Figure 8). Although experienced agents succeed in reducing the steepness of the principals' average payoff curve, the level of coordinated retaliation by agents is not enough to transform such curve in a parable (even in the case of minimal piece rates, there is still a 30% of eveniences in which at least one agent plays H). We also suspect that, as the number of other players increases, considerations of unilateral fairness (pure altruism) may dilute.

5 Discussion and further research

While the experiment shows that vertical and horizontal fairness interact in hierarchical triangles, much needs to be done to better understand the nature of such interactions.

More specifically, understanding how principals' fairness affects relationships between agents deserves some caution. Reciprocity has two faces, one positive and the other negative [Rabin, 1993]: I may be willing to sacrifice my own material well-being to help the kind other or I may be willing to sacrifice my own material well-being to punish the unkind other. In dyadic relationships, these two faces can be easily distinguished. This may not be the case with triangular relationships. In particular, our experiment clearly shows that fair principals tend to generate positive reciprocity between pairs of agents – they act in each other's favor, as well as in favor of the principal, and show higher and more persistent cooperation rates than in a conventional

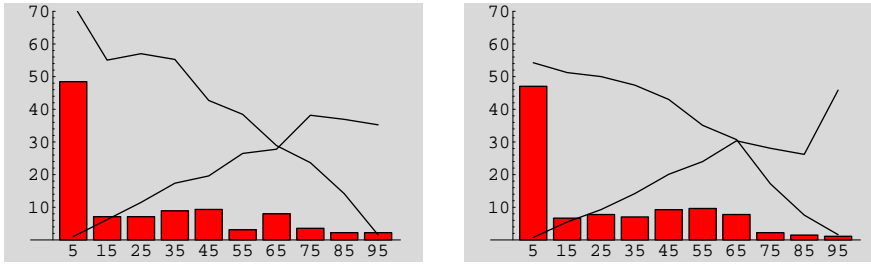


Figure 8: Relative frequencies of piece rate W and average payoff for a principal (downward data plot) and an agent (upward data plot) in the novice and in the experienced treatment.

iterated prisoner's dilemma. On the other hand, interpreting how principals' unfairness affects relationships between agents is much harder. Experimental data suggest that unfair principals induce less cooperation between agents.

But whether this may be due to the fact that greedy principals generate greedy agents, or, conversely, to the fact that agents unite their purposes in retaliating the principal, still remains an unsolved research question. In other words, does hierarchical unfairness induce unkindness or mutualism between agents? Unfortunately, the structure of the game doesn't help much in directly discriminating between these two hypotheses. Neither aggregate data nor individual analysis of single play sequences give any incontrovertible evidence. By deciding to produce L , an agent hurts the principal but at the same time makes the other agent worse off. There is no way to infer an agent's intention within this design and a more accurate experimental design might be devised to separate those two effects.

Furthermore, we think that less symmetric situations are worth exploring.

For example, the principal might be able to differentiate agents' rewards, introducing asymmetries in incentives; asymmetries in agents' capabilities are

case of interest as well. Also, the effects of information asymmetries deserve further investigation: fairness considerations may be significantly affected by different distributions of information among players.

Finally, we claim that our experiment suggests more prudence in the use of standard game-theoretic concepts in organization theory. While the use of non-cooperative games as a tool for modeling organizational phenomena has become widespread, little or no attention at all has been accorded to how behavior in such games may change when they are embedded in a hierarchical context. Our experiment shows that even when equilibrium predictions do not change, the hierarchical context may deeply affect actual agents' strategies. We think that much useful understanding might be gained by systematically exploring how well-known games are played in hierarchical contexts.

Appendix: Experimental Instructions

Introduction

You are participating to an economic experiment. You are kindly asked to carefully read the instructions. Then you will be able to ask questions that will be openly answered. This experiment will last about one hour. If you follow the instructions closely and make decisions carefully, you can earn a considerable amount of money. During the experiment you can earn experimental points that at the end of the experiment will be converted into Italian liras (1 experimental point = 15 Italian liras) and will be added to the fixed amount of 10000 Italian lira. This will be your monetary payment for participating in the experiment.

Instructions

During the whole experiment you are anonymously matched with other two players in this room. One of the players is called *player 1* (from now on, P_1) and the other two players are called *player 2* and *player 3* (P_2 and P_3).

Matching will be performed at random by the computer program at the beginning of the experiment and will not be revealed. Your identity during the experiment (P_1 , P_2 or P_3) will be revealed after reading the instructions and after that all questions will have been answered.

The experiment involves the repetition for 15 times (rounds) of two stages, that will be described in a moment. At the end of each round, payoffs will be announced and then the next round will start. Your final payoff will be equal

to the sum of payoffs earned in each of the rounds, plus the fixed payment of 10000 Italian lira.

Each round runs as follows:

First phase

Player P_1 decides and sends to players P_2 and P_3 the value to be assigned to W . W is a percentage number that can be chosen among all integer numbers between 1(%) and 100(%).

Second phase

Players P_2 and P_3 decide, independently and simultaneously, whether to undertake action A or action B .

End of round

The experimental software computes the quantities Q_2 and Q_3 , produced by players P_2 e P_3 , on the basis of their choices during the second phase according to the following table:

P_3 's action

P_2 's action

	A	B
A	$Q_2 = 60, Q_3 = 60$	$Q_2 = 10, Q_3 = 70$
B	$Q_2 = 70, Q_3 = 10$	$Q_2 = 20, Q_3 = 20$

Finally, the experimental software computes and sends to everyone the payoff earned by each player. Payoffs are computed according to the following

formulas:

$$P_1\text{'s Payoff} = (100 - W)(Q_2 + Q_3);$$

$$P_2\text{'s Payoff} = WQ_2;$$

$$P_3\text{'s Payoff} = WQ_3.$$

Acknowledgements

Authors would like to thank all the participants at the 1999 ESA Regional Meeting in Grenoble and at CEEL 2000 Workshop in Trento, and Rachel Croson, Manfred Königstein, Teck Ho, Alessandro Narduzzo, Tibor Neugebauer, Paolo Patelli, Karim Sadrieh, Rupert Sausgruber and Enrico Zaninotto for their helpful comments and suggestions on earlier drafts of this paper. Michele Lorenzini and Marco Tecilla provided unvaluable help during the development of the experimental software. Support from Enel and Murst (Cofin 99) is gratefully acknowledged.

References

- Anderhub, V., Gächter, S., Königstein, M., 1999. Efficient contracting and fair play in a simple principal–agent experiment. Discussion Paper.
- Andreoni, J., Miller, J.H., 1993. Rational cooperation in the finitely repeated prisoner’s dilemma: Experimental evidence. *Economic Journal*, 103, 570–585.
- Camerer, C., Knez, M.J., 1995. Outside options and social comparison in three-player ultimatum game experiments. *Games and Economic Behavior*, 10(1), 65–94.
- Fehr, E., Falk, A., forthcoming. Wage rigidities in a competitive incomplete contracts market. *Journal of Political Economy*.
- Fehr, E., and Gächter, S., 1998. How effective are trust– and reciprocity–based incentives? In A. Ben-Ner and L. Putterman (Eds.), *Economics, Values and Organizations*. Cambridge: Cambridge University Press.

- Fehr, E., Gächter, S., Kirchsteiger, G., 1997. Reciprocity as a contract enforcement device. *Econometrica*, 65(4), 833–860.
- Fehr, E., Kirchler, E., Weichbold, A., Gächter, S., 1998a. When social norms overpower competition: Gift exchange in experimental labor markets. *Journal of Labor Economics*, 16(2), 324–351.
- Fehr, E., Kirchsteiger, G., Riedl, A., 1993. Does fairness prevent market clearing? an experimental investigation. *Quarterly Journal of Economics*, 108(2), 437–460.
- Fehr, E., Kirchsteiger, G., Riedl, A., 1998b. Gift exchange and reciprocity in competitive experimental markets. *European Economic Review*, 42(1), 1–34.
- Fehr, E., Tougareva, E., 1995. Do high stakes remove reciprocal fairness? – evidence from russia. Discussion Paper, University of Zürich.
- Gächter, S., Fehr, E., 1999. Experiments and labour economics. Discussion Paper, University of Zürich.
- Isaac, R.M., McCue, K., Plott, C., 1985. Public goods provision in an experimental environment. *Journal of Public Economics*, 26, 51–74.
- Itoh, H., 1994. Job design, delegation and cooperation: A principal–agent analysis. *European Economic Review*, 38(3–4), 691–700.
- Kagel, J.H., Wolfe, K., 1999. A new three-person ultimatum game. Discussion Paper.
- Keser, C., Willinger, M., 2000. Principals’ principles when agent’s actions are hidden. *International Journal of Industrial Organization*, 18, 163–185.

- McCabe, K.A., Mukherji, A., Runkle, D.A., 2000. An experimental study of information and mixed-strategy play in the three-person matching-pennies game. *Economic Theory*, 15(2), 421–462.
- Mookherjee, D., 1984. Optimal incentive schemes with many agents. *Review of Economic Studies*, 51, 433–446.
- Nalbantian, H.R., Schotter, A., 1997. Productivity under group incentives: An experimental study. *American Economic Review*, 87(3), 314–341.
- National Research Council, 1991. *Pay for Performance. Evaluating Performance Appraisal and Merit Pay*. National Academy Press.
- Rabin, M., 1993. Incorporating fairness into game theory and economics. *The American Economic Review*, 83(5), 1281–1302.
- Rossi, A., 1999. Incentives in managerial compensation: A survey of experimental research. Discussion Paper, University of Trento.
- Roth, A.E., 1995. Bargaining experiments. In J. Kagel and A. E. Roth (Eds.), *The Handbook of Experimental Economics*. Princeton: Princeton University Press.