Efficient Contracting and Fair Play in a Simple Principal-Agent Experiment

Vital Anderhub, Humboldt University at Berlin
Simon Gächter, University of Zürich
Manfred Königstein, Humboldt University at Berlin*

Abstract

Modern 'principal-agent theory' has made a lot of progress in proposing theoretical solutions to agency problems. This paper contributes to a better understanding of *behavior* in agency situations. In particular, we provide experimental evidence on offered contracts and effort choices in a simple agency game. In line with principal-agent theory we find that in our experiments many contracts proposed by principals are 'incentive compatible' and most agents behave optimally given the terms of the contract. However, in contrast with economic predictions, we find that agents (i) reject 'unfair contracts' and that (ii) given acceptance, their effort choices are to some extent driven by reciprocity. It seems that contract design has to regard an equity constraint that has so far been neglected by contract theory. In fact, most contract offers observed in the experiment aim at fair surplus sharing.

Keywords: Principal-agent theory, contract theory, fair sharing, incentive contracts, reciprocity, experiments.

JEL-Codes: J33, J41, L14, C91.

*Postal addresses: Anderhub and Königstein: Department of Economics, Humboldt-University at Berlin, Spandauer Strasse 1, D-10178 Berlin, Germany. Gächter: University of Zurich, Institute for Empirical Economic Research, Blümlisalpstrasse 10, CH-8006 Zurich, Switzerland. E-mail addresses: anderhub@wiwi.hu-berlin.de; gaechter@iew.unizh.ch; mkoenig@wiwi.hu-berlin.de.

1. Introduction

Agency problems abound in economic life. An agency problem is likely to occur whenever one individual (usually called the 'agent') takes actions on behalf of some other person (the 'principal'). Hence, as it has been pointed out very early on, almost all contractual relationships have some agency element. This is of particular importance for organizations, since they can be understood as a system of "contracting relationships among individuals" (Jensen and Meckling 1976, p. 310). Examples are the relationships between employers and employees, managers and subordinates, and managers and owners. Agency problems are of course also prevalent in business relations between firms and between individuals, between landlords and tenants, voters and the government, and so on. Hart and Holmström (1987, p. 75) make a particularly general argument about the importance of agency relationships by pointing out that agency problems are likely to arise wherever there are gains to specialization.

A crucial aspect of an agency relationship is that the agent does not necessarily take actions that are in the best interest of the principal. This problem is particularly severe if the principal's preferred action cannot be contractually enforced. The reason for this is that there is some asymmetric information either between the principal and the agent or between them and a third party like the courts that prevents the conclusion of enforceable contracts.

Agency problems have been formalized as principal-agent (PA) models and have been a major subject of economic theorizing especially within the last two decades (see, e.g., Ross 1973;Holmström 1979; Shavell 1979; Grossman and Hart 1983; Hart and Holmström 1987; Holmström and Milgrom 1987; and Salanié 1997 for a recent textbook account). In these models, if contracts can not be enforced completely, the principal has to design a work contract that obeys *'incentive compatibility constraints'* and a *'participation constraint'*. The former are necessary to align interests between the principal and the agent because of the presumed opportunism of the agent. The latter constraint is necessary to guarantee that the agent agrees to the contract in the first place. The principal takes optimizing, selfish behavior of the agent for granted and derives the optimal contract which (i) is incentive compatible and (ii) leaves the agent just indifferent between accepting or rejecting the contract. In other words, the principal extracts the whole surplus that is created in the exchange. The agent (iii) accepts the contract and (iv) responds optimally to the incentives set in the contract.

In this paper we study the *behavioral validity* of these four predictions within a simple PA-game. Within this game, the principal first designs a linear contract; i.e., he may offer the agent a fixed wage and a return share. The fixed wage may be positive or negative (within some upper and lower bounds) and the return share can vary between 0 and 100 percent. A negative fixed wage is tantamount to a payment from the agent to the principal; a return share of 100 percent is tantamount to the agent possessing the whole return. Second, the agent decides on whether or not to accept the contract. Thereafter, she chooses an effort level which generates a return and causes a cost (to be beared by the agent). These three stages comprise a single base game. In the experiment the subjects play two repeated games each consisting of six base games.

Within the given restrictions any combination of fixed wage and return share constitutes a feasible contract in our experiment. If, however, the agent is selfishly motivated, a *fixed wage contract* (with a zero return share) does not at all align interests between the principal and the agent; it leads to full shirking of the agent. Incentive compatibility requires the *return share* offered to the agent to be sufficiently high. In our model, the unique 'trembling hand' perfect equilibrium of the game calls for a return share of 100 percent, i.e., the agent owns the whole return for which he pays a price (i.e., a negative fixed wage) that amounts to the generated surplus. However, this solution as well as all subgame perfect equilibria of the game are questionable from the viewpoint of fairness because the principal receives all surplus.

The benchmark for our analysis are some empirical observations made in the last decade of experimental research on people's social motivations in bargaining contexts and in cooperation games. For our purposes two sets of results are particularly relevant. First, by now there is a lot of evidence that people act less often opportunistically than is assumed by standard economic theory and in most agency models. For example, Fehr, Gächter and Kirchsteiger (1997) have shown that in a simple agency game that was designed such that selfish agents had an incentive to shirk, much less shirking actually occurred than theoretically predicted. Instead of shirking, agents made their effort choices dependent on the generosity of the work contract offered to them by the principal. They responded with 'high' effort choices if the principal offered them 'generous' wages and they put forward 'low' effort levels if the terms of the contract were greedy. Put differently, agents, on average, behaved *reciprocally.* Under certain conditions reciprocity turned out to be a very powerful contract enforcement device. This result has by now been replicated several times in different games and under different institutions (see Güth, Klose, Königstein, and Schwalbach 1998; and Fehr and Gächter 1998 for an overview). The observation of reciprocity in agency relationships is not restricted to the laboratory. To give just one example, Bewley (1997) provides extensive field evidence for the importance of reciprocal behavior in employment relations.

A second set of results that is relevant here comes from ultimatum games.[1] In the subgame perfect equilibrium of that game the proposer offers the smallest amount of money and the responder just accepts it. As we have just pointed out, this is also the prediction about the principal's contract offer and the agent's acceptance decision. The experimental results in the ultimatum game are not at all in line with this prediction. Subjects reject allocations that give them only a fraction of the pie. Accordingly, proposers usually offer between 40 and 50 percent of the pie. Again, this result has been replicated many times (see Camerer and Thaler 1995, Güth 1995, and Roth 1995 for overviews).

While the first set of results suggests that people act less often opportunistically than assumed by standard economic theory, the latter set of results may be informative about the behaviorally relevant participation constraint. These observations raise several interesting questions that we want to investigate in this paper. How do agents respond to the incentives set in the contract? Given that people are often reciprocally motivated, how does this influence agents' behavior under incentive contracts? In turn, if agents are less opportunistic than assumed by standard theory, what kind of contracts do principals design? To what extent do they rely on incentive contracts? Does the result that people reject unfair offers carry over to more complex contractual arrangements? In other words, what is the behaviorally relevant participation constraint in an agency relationship?

Our most important results, presented in detail in sections 3 and 4, are as follows. In line with economic predictions, principals offer incentive compatible return shares and ask for negative fixed wages in many cases. Agents choose conditionally rational effort levels in many cases, i.e., we find optimal effort decisions on and off the equilibrium path. Deviations from the individually optimal effort level can be explained by reciprocity. Agents who receive generous contracts are more likely to increase their effort above the conditionally optimal level than agents who receive greedy offers. However, agents are also prepared to reject 'unfair' contracts. Thus, the observations from the ultimatum game carry over to our principal-agent experiment. The behaviorally relevant participation constraint is

[1] In this game a proposer can divide a fixed sum of money between him and a responder who can either accept the proposal, or reject it. In case the responder rejects both get nothing, whereas an acceptance leads to the implementation of the proposed allocation.

one that involves fair sharing between the principal and the agent.

Our paper contributes to the behavioral investigation of agency problems. In this respect we complement empirical investigations that use naturally occurring data on issues of incentive contracting (see Prendergast 1999 for a comprehensive overview of this literature). However, there are by now not many investigations that study the *behavioral elements* that make up a contractual solution of an agency problem. The papers that are closest to ours in this respect are the experimental studies of Berg, Daley, Dickhaut, and O'Brien (1992); Epstein (1992); Güth, Klose, Königstein, and Schwalbach (1998); and Keser and Willinger (1997) who all study contract design and agents' behavior in *moral hazard* situations. Güth, Königstein, Kovács, and Zala (1999) study a PA–game with multiple agents.[2]

In contrast to most of these papers, in our design principals can choose among a very large number of contracts - including pure fixed wage contracts, pure incentive contracts and any mixture of incentive and fixed wage contracts. Agents can choose among a large set of effort levels. These design features allow us to investigate the relevance of fine-tuned incentive contracts, 'fair sharing', as well as the principles which are behind effort choices.

In section 2 we present the game underlying our study, its theoretical solution and the applied experimental procedures. Thereafter we present our experimental data on contract design (section 3) and on the behavior of agents (section 4). Section 5 concludes.

2. The Experimental Design

2.1. The game

In the experiments we implemented a finitely repeated game that consisted of six repetitions of the following base game between two players, a principal P and an agent A. First, the principal designs a work contract and makes a 'take it or leave it'–offer to the agent. Secondly, the agent either accepts or rejects the

[2] Other PA–studies which are, however, less closely related to ours are: Bull, Schotter and Weigelt (1987) who test the incentive effects of piece rate and tournament payment schemes; Nalbantian and Schotter (1997), who investigate various contracts to provide group incentives; Hackett (1993), who studies incomplete contracting, and Chaudhuri (1998), who investigates ratchet effects in a principal-agent relationship. Plott and Wilde (1982); DeJong, Forsythe, and Lundholm (1985); and DeJong, Forsythe, Lundholm and Uecker (1985) study moral hazard problems with multiple buyers and sellers.

contract and, thirdly, chooses her work effort. In detail, the three stages and the parameters of the base game in period t are as follows:

STAGE 1: P chooses a work contract $w_t = (f_t, s_t)$ that consists of two components, a fixed wage f_t and a return share s_t. P also states a 'suggested work effort' $\tilde{e}_t$ which is not binding for A's choice later on. The work contract and the suggested work effort have to obey the following restrictions:

$$\begin{aligned} f_t &\in \{-700, -699, ..., +700\} \\ s_t &\in \{0, 0.01, ..., 1\} \\ \tilde{e}_t &\in \{0, 1, ..., 20\}. \end{aligned}$$

In period $t = 1$ the principal is allowed to design up to two work contracts (and to suggest a work effort for each contract).

STAGE 2: A may either accept ($\delta_t = 1$) or reject ($\delta_t = 0$) the contract offered by P (in $t = 1$ she may accept one of the two offered contracts or reject both). This decision determines the 'implemented contract'. For $\delta_t = 1$ the implemented contract is equal to the offered contract (one of the offered contracts in $t = 1$). In case of $\delta_t = 0$ the implemented contract is equal to $\bar{w} = (0, 0)$, the 'status quo contract'.

STAGE 3: A chooses work effort e_t with

$$e_t \in \{0, 1, ..., 20\}.$$

The agent has complete effort discretion, i.e., she is not restricted by $\tilde{e}_t$ suggested by the principal.[3]

The repeated game proceeds by P getting informed about δ_t and e_t after A's effort choice. Then, either the next period follows (if $t < 6$) or the game ends (if $t = 6$).

A's work effort e_t determines the return according to the return schedule $r_t = r(e_t) = 35e_t$. The players' repeated game payoffs are the sum of the following base game payoffs:

$$\pi_t^P = (1 - s_t) \cdot r_t - f_t \tag{2.1}$$

$$\pi_t^A = s_t \cdot r_t + f_t - c_t \tag{2.2}$$

[3] Technically, A also has to make an effort choice if she has declined $P's$ contract offer.

with s_t, and f_t according to the implemented contract, the return r_t, and cost of work effort $c_t = c(e_t)$. The cost function $c(e_t)$ is piecewise–linear, increasing in e_t, and convex (for a tabular representation of this cost function as it was presented to the subjects, see the instructions in the Appendix):

$$c(e_t) = \begin{cases} 5e_t & \forall e_t = 0, ..., 4 \\ -20 + 10e_t & \forall e_t = 5, ..., 8 \\ -60 + 15e_t & \forall e_t = 9, ..., 12 \\ -120 + 20e_t & \forall e_t = 13, ..., 16 \\ -200 + 25e_t & \forall e_t = 17, ..., 20. \end{cases}$$

For our purposes, this type of cost function has some useful properties which will be discussed later on. In the following we turn to the game-theoretic solution of this game, and then discuss some important features of our design.

2.2. Game-theoretic solution

The game-theoretic solution we now derive assumes rationality and selfishness of all players. We start by characterizing the set of subgame perfect equilibria of the base game in $t = 6$. For a given implemented contract w_6 the agent maximizes (2.2) by choice of e_6. Her best reply effort function $\hat{e}_6(s_6)$ satisfies:

$$\hat{e}_6 \equiv \arg\max_{e_6} \pi_6^A \tag{2.3}$$

In contract theory this condition is known as the 'incentive compatibility constraint'. We refer to $\hat{e}_t$ as the 'conditionally rational effort choice'. It is the effort level that equates A's marginal return and her marginal cost. Since marginal return is constant and marginal cost is a step function, one can easily derive the conditionally rational effort choices as a function of the return share s_6 (see figure 1 and table 1). Thus, the piecewise–linear specification of the cost function allows for only six different effort levels that might be conditionally rational. In the data analysis below, this feature will make it easy to determine whether agents choose conditionally rational effort levels.

Insert figure 1

Return Share (s_6)	Conditionally Rational Effort ($\hat{e}_6$)
$0 \leq s_6 < \frac{1}{7}$	0
$\frac{1}{7} \leq s_6 < \frac{2}{7}$	4
$\frac{2}{7} \leq s_6 < \frac{3}{7}$	8
$\frac{3}{7} \leq s_6 < \frac{4}{7}$	12
$\frac{4}{7} \leq s_6 < \frac{5}{7}$	16
$\frac{5}{7} \leq s_6 \leq 1$	20

TABLE 1: Conditionally rational effort choices. Note that figure 1 suggests that $\hat{e}_6$ is not unique for $s_6 \in \left\{\frac{1}{7}, \frac{2}{7}, \frac{3}{7}, \frac{4}{7}, \frac{5}{7}\right\}$. However, within the experiment s_6 could be chosen only in increments of 0.01 such that non–uniqueness was impossible.

Before deciding upon an effort level, A has to choose whether to accept or to reject the offered work contract. A rejection $\delta_6 = 0$ means that the status quo contract $\bar{w} = (0, 0)$ becomes effective and that by her subsequent choice of $\hat{e}_6 = 0$ she can guarantee herself a payoff of at least zero. Accordingly, $\delta_6 = 1$ if and only if

$$\pi_6^A(\hat{e}_6) \geq 0 \tag{2.4}$$

which will be referred to as the 'participation constraint'. Thus, the principal's problem is described as follows:

$$\max_{\hat{e}_6, f_6, s_6} \pi_6^P$$

$$\text{subject to (2.3) and (2.4).}$$

Given the participation constraint is satisfied, (2.3) says that effort solely depends on s_6. Furthermore, note that multiple return shares may lead to the same effort and that optimal behavior requires for P to induce a given effort at minimal cost (see, e.g., Grossman and Hart 1983). Therefore, choosing $\hat{f}_6$ such that (2.4) holds with equality implies that every $(\hat{f}_6, s_6)$ induces $e_6 = \hat{e}_6(s_6)$ at minimal cost. However, this also implies that the principal captures all the surplus. Thus, maximizing π_6^P implies surplus maximization and one easily finds that this is the case for all $s_6 \geq \frac{5}{7}$. In summary, the paths of all subgame perfect equilibria (SPE)

(f_6^*, s_6^*, e_6^*) are:

$$\begin{aligned} f_6^* &= 300 - 700 s_6^* \\ s_6^* &\in \left[\tfrac{5}{7}, 1\right] \\ e_6^* &= 20. \end{aligned}$$

All SPE are payoff equivalent leading to $\pi_6^P = 400$ and $\pi_6^A = 0$. Invoking equilibrium refinement one can show that the unique path of all perfect equilibria (see Selten 1975) is characterized by $f_6^* = -400, s_6^* = 1, e_6^* = 20$; i.e., the principal 'sells the return to the agent' at a price that equals the return that is generated by $e_6^* = 20$.[4]

This is the solution for period 6. By backward induction it follows that the solutions for periods 1 to 5 are equivalent: Note that in $t = 5$ any long-term agreement that specifies punishment (in $t = 6$) in case of deviations (in $t = 5$) is not subgame perfect. So, the solution is as in $t = 6$. This reasoning goes through all periods.

2.3. Discussion of the design

The present design reflects our main goal to investigate the principals' choice of contracts as well as the agents' behavior in the most simple but yet general enough environment that allows for a test of main elements of proposed solutions of agency problems. In particular the design gives us the opportunity to check to what extent contract design takes incentive compatibility and a participation constraint into account. It also allows us to observe the agents' (incentive compatible) behavior when confronted with a particular contract. However, our game has some specific features that distinguishes it from general PA–games and we briefly want to discuss some of these aspects.

First, the space of feasible contracts among which the principal can choose is restricted to the space of linear contracts. This is a restriction compared to the standard PA–framework (see, e.g., Grossman and Hart 1983). However, within the space of linear contracts, incentive compatible contracts that induce efficiency are feasible (and optimal).

Second, in contrast to most PA-models, effort in our model is observable. In this case economic theory suggests that the optimal contract imposes high punishments on A for effort choices that are suboptimal from P's perspective. However,

[4] For all $s_6^* < 1$, a 'trembling hand' effort choice of the agent reduces the principal's expected payoff.

with limited liability, or if effort is not verifiable to a court, such contracts are not enforceable. Our restriction of the set of feasible contracts essentially plays an equivalent role. Contracts that impose high punishments for deviating behavior are infeasible. Thus, despite observable effort, we preserve some general features of PA–games with non–observable effort: The agent can deviate from the required effort without being punished; a pure fixed wage policy induces a rational and selfish agent to shirk; and when designing a work contract the principal has to take incentive effects into account.

Last, our model is non–stochastic. Our design has the advantage that we have modelled an agency problem in which we can investigate behavior without having to deal with risk preferences and the cognitive complexities that go along with the introduction of stochastic outcomes.[5]

Linear contracts that are akin to the ones investigated in this paper (i.e., that use some combination of fixed payments and return shares) can be found in reality. For example, they are used in sharecropping (see, e.g., Alston and Higgs 1982; Allen and Lueck 1993), in franchising (see e.g., Mathewson and Winter 1985); between cab owners and cab drivers (e.g., Camerer, Babcock, Loewenstein and Thaler 1997); and between authors and publishers or between actors and producers (for an example from the motion pictures industry see Chisholm 1997). Prendergast (1999) provides a comprehensive survey on empirical studies that investigated (linear) incentive contracts. A theoretical justification of linear contracts is given by Holmström and Milgrom (1987) and, in the context of sharing contracts, by Bhattacharyya and Lafontaine (1995).

Linear contracts are simple. This is an advantage both in reality as well as for the purposes of our experimental investigations. Linear contracts can be easily understood by experimental subjects, which increases control. It should be noted, however, that the space of contracts we allow for is actually quite large compared to other experimental studies. For instance, in Berg et al. (1992) and Epstein (1992) contracts had to be chosen among three alternatives. In our experiment

[5]Theoretically, under the assumption of von Neumann-Morgenstern preferences, there are experimental techniques that allow for the control of risk preferences (see Berg et al. 1986 for a general statement). However, there is a lot of experimental evidence that subjects are not expected utility maximizers (see Camerer 1995 for a survey of results). Therefore, some researchers have doubts on the validity of these techniques. For example, Selten, Sadrieh, and Abbink (1995) present experimental test results that show that these techniques do not at all work in a variety of economic decision situations. Loomes (1998) points in a similar direction. Keser and Willinger (1997) discuss these arguments in the context of their principal-agent experiment.

the principal can design contracts with every mixture of fixed wages f_t and return shares s_t. For example, a contract of the form $w_t = (s_t = 0, f_t > 0)$ means that the principal just offers a fixed wage to the agent and no return share at all. At the 'opposite extreme' a contract of the form $w_t = (s_t = 1, f_t < 0)$ is tantamount to 'selling the return' to the agent at a selling price of f_t. Furthermore, all 'mixed contracts' $w_t = (0 < s_t < 1, f_t \gtreqless 0)$ are possible (within the given limits set for f_t). This great flexibility in contract design allows us to study the importance of incentive compatibility *simultaneously* with issues of 'fair sharing' and reciprocity in one experimental design. In this respect our study is among the first to extend isolated previous investigations on bargaining and reciprocity to the realm of agency problems.

2.4. Experimental procedures

The experiment comprised four sessions that were conducted at the University of Zürich in June 1997. The participants were 94 volunteer students with various backgrounds (except economics). They were all recruited from a large data base. After subjects arrived at the laboratory, they were randomly allocated to their roles as principals or agents, respectively. In the experiments the principals were called 'participant X' and the agents were called 'participants Y'. Then subjects had to read the instructions (a translation can be found in the Appendix). Except for some role-specific wordings, the instructions were exactly the same for both the principal and the agent. Subjects also had to answer a set of control questions to test for their understanding of payoff calculations. The experiments did not start before all subjects had answered all questions correctly. Each subject played two repeated games (referred to in the sequel as RG1 and RG2) with no role switch. It was explained to the subjects that all decisions will be taken anonymously and that identities will never be revealed. At the beginning of RG1 it was publicly announced that subjects will play a repeated game with six periods with the same opponent. At the end of RG1 subjects were informed that they will play a second repeated game (RG2) of six periods with a different opponent and again with no role switch.

The experiments were computerized and conducted with the help of the experimental software 'z-Tree' developed by Fischbacher (1998). Subjects were separated from each other by blinders and matched anonymously via a computer network. They never learned the identity of their opponent players. Each session took about one hour. The average payoff per subject was about 40 swiss francs (roughly \$ 33 at the time of the experiment and including a show up fee of 15 Swiss francs).

3. Contract design

We present the results of our experiments as follows. First we discuss the features of contract design; i.e., offered fixed wages and return shares as well as suggested efforts. Furthermore, we investigate surplus sharing as suggested by principals. Basically, we do that by looking at the raw data without providing statistical tests. Then we turn to the analysis of the agents' behavior, their decisions to reject or accept an offer as well as their effort choices. Here, we will state three hypotheses for which we provide rigorous statistical tests. We will summarize all our findings as 'Observations'.

3.1. Distributions of fixed wages, return shares and suggested effort

Figures 2a to 2c display the empirical distributions of fixed wages and return shares as offered by the principals, as well as their suggested work efforts.

Insert figures 2a to 2c

We see that the majority of principals (about 70 percent) offered negative fixed wages.[6] About 70 percent of all contracts exhibited return shares of $s \geq 72$ percent. Remember that this is what incentive compatibility suggests. Looking at the suggested work efforts we find overwhelming evidence for efficiency; i.e., by far most principals (more than 82 percent) aimed at inducing an efficient effort choice. In the following we will refer to these contracts as 'efficiency contracts'. Comparing suggested work efforts and return shares one observes that the fraction of efficiency contracts is higher than the fraction of incentive compatible return shares. This indicates that at least some principals aimed at efficiency without providing proper monetary incentives. We summarize these findings in the following observations.

Observation 1: *In most cases the offered return share is in the range predicted by the subgame perfect equilibria.*

[6] Note that we report the contracts for all base games. Since each of 47 principals played 12 base games this makes 564 contracts.

Observation 2: *Principals aim at inducing efficient (maximal) effort.*

Observation 3: *Most contracts exhibit a negative fixed wage.*

These results suggest that the principals clearly recognized the incentive problem. In the following, we take a closer look at the *types* of contracts the principals actually designed.

3.2. Types of contracts

Remember that principals could offer contracts that contained any feasible combination of fixed wages and return shares. Table 2 reports which types of contracts were offered in RG1, respectively RG2. Pure fixed wage contracts were very rare in RG1 and were never chosen in RG2. The majority of offered contracts is of a 'mixed' type. Interestingly, a stable share of at least 24 percent of the contracts combined positive shares with non-negative fixed wages (which, under the standard assumption do not have any incentive effect). The perfect equilibrium predicts 'sell the firm' which describes up to thirty percent of all contracts. We keep these results as

Observation 4: *'Selling the firm' occurs in about 30 percent of the cases. Roughly 70 percent of all contracts are of a 'mixed' type. Of those, at least a quarter contains positive fixed wages.*

A comparison of RG1 and RG2 reveals, furthermore, that contract design was relatively stable across periods and repeated games.

CONTRACT TYPE			RG1	RG2
'pure fixed wage'	$s_t = 0$	$f_t > 0$	1.5 %	0 %
'mixed'	$1 > s_t > 0$	$f_t \geq 0$	30.1 %	24.6 %
	$1 > s_t > 0$	$f_t < 0$	41.7 %	41.7 %
'sell the firm'	$s_t = 1$	$f_t < 0$	24.6 %	30.4 %
other			2.1 %	3.3 %
$\sum$			100 %	100 %

TABLE 2: Relative frequency of different contract offers in both repeated games.

A necessary condition for the agent to accept the contract is that the agent's 'participation constraint' (i.e., 2.4) is satisfied. To check this, we determined the maximal payoff that the agent could get given the offered contract. Thus, we calculated the agent's payoff by assuming a conditionally rational effort choice. Violations of the participation constraint (i.e., cases in which the maximal payoff the agent could achieve was negative) occurred only in 13 out of 564 cases (i.e., in 2.3 percent of the contracts).[7] Hence, these results support

Observation 5: *Almost all offered contracts satisfy the participation constraint.*

According to the theoretical predition, the participation constraint can be satisfied even with a very unequal distribution of the surplus. In the following subsection we present the results of the actually offered surplus sharing by the principal.

3.3. Suggested surplus sharing

The game theoretical solution under the assumption of rationality and selfishness suggests the following linear relation between fixed wages and return shares (see section 2.2):

$$f_t = 300 - 700 s_t \quad \forall s_t \geq \tfrac{5}{7}. \tag{3.1}$$

This implies that the principal 'sells his asset' and demands the whole surplus that is created by $A's$ rational effort choice as the price for selling his asset. Figure 3 is a scatter plot of all observed combinations of return shares and fixed wages. It clearly suggests a negative tradeoff between the two contract components, at least for $s_t \geq 72$ percent $> \frac{5}{7}$ (note that in the experiment s_t had to be an integer number between 0 and 100).

Insert figure 3

Since the observations 1 to 3 are qualitatively in line with economic theory, one might ask whether principals indeed tried to extract the whole surplus from the agent. However, in general, this is not the case. The solid line that is drawn in figure 3 displays equation (3.1). It represents the combinations of fixed wage and

[7] In two cases the agent nevertheless accepted the contract. Both cases occurred in the first period of the first game (!), which indicates decision errors due to inexperience rather than anything systematic.

return share at which the agent receives zero income given he chooses efficient effort. Obviously, most observations lie *above* this line; i.e., for a given return share the fixed wage offered to the agent was higher than the rational fixed wage level. This implies that firms do not fully extract the created surplus. It holds in almost all cases.

The broken line in figure 3 represents contracts at which the efficient surplus is split 50–50 between both players. Thus, most contracts offer the agent a share between 0 and 50 percent of the efficient surplus. This is also shown in figure 4 where we look at the surplus distribution from yet another angle. It depicts the 'implied surplus distribution' between the principal and the agent for all efficient contracts. Specifically, the figure displays the surplus share of the agent as it is implied by the offered work contract.[8]

Insert figure 4

The mode of the distribution is at the equal split and almost all shares are between zero and 0.5. Thus, in the majority of cases, the principals offer a 'fair' split of the pie in the following sense: The payoff distribution that results if the agent follows the suggestion of the principal and chooses efficient effort, is less asymmetric than the rather uneven payoff distribution that is theoretically predicted. Note that this does not preclude that the suggestion of the principal favors himself relative to the agent. These results closely resemble the above mentioned findings in ultimatum bargaining games and provide support for

Observation 6: *Surplus sharing suggested by principals is less asymmetric — i.e. more fair — than predicted by the theoretical solution.*

In summary, the principals' contract offers are largely in the *direction* predicted by standard economic theory. In particular, in the overwhelming number of cases, the offered contracts satisfy the participation constraint and contracts contain a strictly positive return share. However, quantitatively, there are some important deviations from standard predictions. Many contracts contain return shares that are 'too low', as well as positive fixed wages that, under the maintained assumptions of rationality and selfishness, do not have any incentive effects. Moreover, most offered surplus shares are much more equal than predicted.

[8]The agent's surplus share is evaluated based on the terms of the offered contract and assuming that the agent indeed chooses efficient effort (as suggested).

In the next section we turn to the behavior of the agents. We first investigate their acceptance decisions and then present results on the agents' actual effort choice dependent on the offered contract.

4. Behavior of Agents

4.1. Acceptance of contracts

Remember that about 97 percent of the offered contracts satisfied the 'participation constraint'. Despite this, 112 contracts out of 551 contracts that satisfied the participation constraint, were *rejected* by the agents. This is in clear violation of standard economic theory that predicts a rejection if and only if the participation constraint is violated.

An important question is, therefore, why agents rejected the offered contract in 20 percent of the cases. Results from numerous studies on ultimatum games offer important hints. For example, Slonim and Roth (1998) have shown that in the ultimatum game the acceptance rate of the responder was positively correlated with the relative payoff she received. In the context of our principal-agent game this result suggests that acceptance behavior is related to the offered surplus share. We formulate this intuition in the following

Fairness–Hypothesis: *The influence of the agent's surplus share on her acceptance rate is positive.*

To test this hypothesis, we estimated a logit regression model with the agent's acceptance decision (1 = acceptance, 0 = rejection) as dependent variable and the agent's surplus share as explanatory variable. Only contracts that satisfy the participation constraint enter the regression. The agent's surplus share is determined as her payoff relative to the surplus, given that she chooses a conditionally rational effort level. Thus, this model proposes that the agent looks at the split of income between the principal and herself that will result if she chooses her income maximizing effort level. The null hypothesis, derived from standard arguments, is that there is no correlation between the surplus share and the acceptance rate. The estimation results are contained in table 3. Figure 5 illustrates the acceptance rate as a function of the offered surplus share.

Variable	RG1 Coeff.	RG1 S.E.	RG1 P	RG2 Coeff.	RG2 S.E.	RG2 P
Constant	-1.08	0.46	0.000	-1.41	0.51	0.005
Surplus Share	6.49	1.27	0.019	7.54	1.41	0.000
Model						
$-2LL$	231			241		
$-2LL$ (restr.)	267			278		
χ^2	36			37		
N	263			275		

TABLE 3: Logit regression of the agent's acceptance decision on her surplus share. (Note that 13 observations are missing. These are cases in which the surplus is zero, so that the surplus share is not defined).

Insert figure 5

We estimated the model separately for both repeated games. Figure 5 shows that there is virtually no difference between RG1 and RG2. The influence of the surplus share is significantly positive and therefore supports the fairness–hypothesis. The prediction of a zero correlation is clearly rejected by the data.

The observed effect is statistically significant in both repeated games. However, statistical testing assumes that errors are uncorrelated, which is questionable here, since each subject was observed several times (repeated measurement). Furthermore, tests based on RG2 data are especially problematic due to the strategic interaction of subjects in RG1. To take care of the repeated mesurement problem we applied the procedure suggested by Slonim and Roth (1998) to analyze RG1 data. We used a subject's average acceptance rate in periods other than the current one as an additional explanatory variable (for details see their article). However, this procedure changed the regression results only marginally. In additon, we ran a very conservative test of the fairness–hypothesis based on average data for RG1. We took a subject's mean surplus share and its acceptance rate across the six base games of RG1 and then determined the Spearman rank correlation coefficient between these variables. The correlation was positive (0.33) and significant on the 5 percent level ($p = 0.011$, $N = 47$, one–tailed). Thus, even this very conservative test procedure supports the fairness–hypothesis. These results, therefore, give credit to

Observation 7: *The agents' acceptance decisions support the 'Fairness-Hypothesis'.*

We conclude that whether or not the agent accepts the contract offered by the principal does not only depend on the implied absolute payoff but also on the implied relative payoff. We interpret this as an indication of the agent's care for fairness. It suggests that the behaviorally relevant participation constraint in PA–games is not characterized by the reservation payoff which the agent may receive outside the relationship (e.g., in the market), but is given by some payoff that captures 'fair sharing'.

How did agents who accepted a contract, react to the incentives set by the contract? To answer this question, we turn in the following subsections to the agents' actual effort behavior.

4.2. Rational effort choices

Figure 6 displays the distribution of actual (instead of suggested) effort decisions for all base games in which an agreement was reached. The data are shown separately for both repeated games, RG1 and RG2.[9] We observe a striking similarity between the two repeated games. Obviously, experience does not affect the effort distribution very much. 60 percent of all base game effort decisions in RG1 are efficient (67 percent in RG2).

Insert figure 6

According to the game-theoretical analysis (see table 1), only the effort levels $\hat{e}_t \in \{0, 4, 8, 12, 16, 20\}$ might be rational for an agent with selfish preferences. Notice that this feature of our design allows for a strong test of incentive compatible behavior. To determine whether or not an effort choice might be considered (sequentially) rational, we look, in figure 7, at the *conditional effort distributions* for each given theoretical prediction $\hat{e}_t$. In 264 base games rationality calls for $\hat{e}_t = 20$, and in 230 of these cases (87 percent) agents did indeed choose the conditionally rational effort level. Moreover, these effort choices are also efficient. Looking at the conditional effort distribution for $\hat{e}_t = 16$, one again finds a substantial proportion of rational choices: 20 out of 81 cases (25 percent) ar conditionally rational. There is an even larger number of efficient effort choices

[9] Remember that a principal's suggested effort was not binding for the agent's actual choice.

(35 cases, i.e., 43 percent). Thus, more than two thirds of the decisions are either rational, or efficient. These two kinds of effort decisions (i.e., rational and efficient effort choices) are also taken more often than anything else in case $\hat{e}_t = 12$. For the cases $\hat{e}_t \in \{0, 4, 8\}$ rational effort is always the modal decision, but we observe no efficient effort any longer (despite some dispersion).[10]

Insert figure 7

Inspecting these distributions clearly suggests the following:

Rational–Effort–Hypothesis: *Effort choices are to a large degree conditionally optimal.*

Table 4 shows, furthermore, that the percentage of conditionally rational effort choices increases with experience.

	RG1		RG2	
$t =$	$1,2,3$	$4,5,6$	$1,2,3$	$4,5,6$
Rational Effort	53%	61%	66%	76%
Other	47%	39%	34%	24%
$\sum$	100%	100%	100%	100%

TABLE 4: Percentage of rational effort choices for different levels of experience.

In order to test the Rational–Effort–Hypothesis we look at the behavior of experienced subjects in RG1. Specifically, in the last period ($t = 6$) of RG1 26 out of 39 agents (who had accepted the offered contract) chose conditionally rational effort while 13 chose some other effort. Thus the fraction of rational choices is $\frac{2}{3}$. In principle, behavior might be driven by rationality as well as other forces; i.e., a behavioral theory might comprise various concepts. For example one might think that individuals are driven by several motivating forces at the same time and choices represent these mixed motives. Or one may think that some individuals are rational while others are not, which produces a distribution of choices

[10]For $\hat{e}_t = 4$ rationality is only weakly modal, but this conditional distribution is based on just four observations.

at the aggregate level. Either way, one may ask the following question: What is the lowest probability weight that such a behavioral theory may put on rational effort? More specifically: Which is the lower bound for the probability of rational choices that is compatible with the observed frequency of rational effort, for some desired degree of confidence? Since we code effort choices binary (rational or non–rational), this means essentially that we are looking for a lower confidence bound for $prob(e = e^*)$ (the probability for rational choice) based on the binomial statistic. If we fix the confidence level at 99 percent, the binomial statistic for observing 26 out of 39 rational choices gives a lower confidence bound for $prob(e = e^*)$ at 0.54. Accordingly, this procedure allows us to reject the null hypothesis $H_0 : prob(e = e^*) \leq 0.54$ in favor of the alternative hypothesis $H_1 : prob(e = e^*) > 0.54$ at the 1% significance level ($N = 39$, one–tailed).

This evidence is strong support for the rational–effort–hypothesis. Remember that in each case the conditionally rational choice is given by a single effort level out of 21 possible levels. By classifying the decisions into 'rational' versus 'other' decisions our test therefore rejects a large class of behavioral theories. Every theory that induces a probability weight less than 0.54 for the rational decision can be rejected. For example, we can reject a theory that combines rationality and 'noise', if the noise level reduces $prob(e = e^*)$ below the above threshold. The test does not use information regarding the distribution over non–rational decisions. It makes use only of the probability weight that a theory puts on the rational decision versus a non–rational decision. Since the exact distribution over non–rational decisions is irrelevant for the result of the test, this is a rather general support of the Rational–Effort–Hypothesis.[11] We conclude this subsection by stating

Observation 8: *Agents' effort choices support the 'Rational-Effort-Hypothesis'.*

Figure 7 shows that, although a large number of effort choices can be explained by Observation 8, a non-negligible number of effort decisions deviate from the conditionally rational effort level. In the following we investigate whether *reciprocal behavior* can account for them.

[11] The same test procedure could have been applied to the RG2 data. We abstained from statistical testing because the strategic interaction of RG1 might cause correlated errors in RG2.

4.3. Reciprocal effort choices

As mentioned in the introduction, other studies on games that involve a PA-structure have found that effort choices can be characterized as reciprocal: contracts that are favorable for the agent trigger effort choices that are above the individually optimal effort level (positive reciprocity) whereas unfavorable contracts trigger agents' willingness to punish the principal by choosing a low effort level (negative reciprocity). In particular, in some of these experiments principals were even able to achieve high effort levels with fixed wage contracts only. This observation suggests that, behaviorally, even fixed wages can have a positive influence on effort levels.

How does the support for the Rational-Effort-Hypothesis square with the robust observation of reciprocal behavior made in many other experiments? Figure 7 suggests that the deviations from rational effort do not appear as random; rather, agents deviated toward efficiency in many cases. Thus, we propose that the deviations depend on the agent's payoff relative to that of the principal. In our present setup, how favorable a contract for the agent is, can be measured by the surplus share the principal offers to the agent. Hence, the results from previous experiments suggest the following

Reciprocal–Effort–Hypothesis: *Deviations from the conditionally rational effort level are positively correlated with the agents' surplus share.*

Similar to the above analysis of rejection behavior, we determined the mean surplus share of the agent in RG1, respectively RG2, which she could have achieved by choosing conditionally rational effort; i.e., we took the mean across those periods in which she did not choose rational effort. Figure 8 displays the correlation between the agent's mean surplus share (classified in five ranges) and her average deviation from the conditionally rational effort level. The figure indicates, at least for RG1, a positive correlation as proposed by the Reciprocal–Effort–Hypothesis. The Spearman rank correlation coefficient is 0.35 in RG1. It is qualitatively as predicted and statistically significant ($p = 0.024$, $N = 33$, one–tailed). In RG2 the coefficient is positive as well (0.22). Due to the strategic interaction of the subjects, a statistical test in RG2 is problematic, however.

Insert figure 8

Hence, the results support

Observation 9: *The deviations from the conditionally rational effort levels are compatible with the 'Reciprocal–Effort–Hypothesis'.*

In summary, agents' behavior can be characterized as follows. Agents reject 'unfair contracts' (i.e., contracts that give them only a relatively small share of the surplus). The agents clearly react to the incentives set in the contract. Higher return shares clearly lead to higher effort choices. The agents, in particular, make conditionally rational effort choices in a large number of cases. Deviations from conditionally optimal levels can be explained with reciprocity.

5. Summary

In this paper we have studied behavior within a simple principal–agent experiment. Our design allowed for a large class of linear contracts which, so far, has rarely been investigated experimentally. Principals could offer contracts with any feasible combination of positive or negative fixed wages and a return share between 0 and 100 percent. This great flexibility in contract design allowed us to study agency considerations *simultaneously* with issues of 'fair sharing' and reciprocity - found to be important in previous studies - within one experimental design.

Under standard economic assumptions, the optimal contract induces efficient effort. To achieve this, the contract has to provide proper monetary incentives requiring a return share that is sufficiently high. Furthermore, the theoretical model predicts agreement despite a rather asymmetric split of the surplus.

Several important features of our empirical data are, by and large, in line with orthodox economic theory. The principal designs a work contract that offers an incentive compatible return share (**Observation 1**), aims at implementing efficiency (**Observation 2**), and asks for a negative fixed wage (**Observation 3**). 'Selling the firm' occurs in about 30 percent the contracts (**Observation 4**). Almost all contracts satisfy the participation constraint (**Observation 5**). However, their contract offers are more fair than standard economic arguments would suggest (**Observation 6**). The latter two results have also been observed by Keser and Willinger (1997). Thus, the broad message from Observations 1 to 5 is that principals clearly recognize the agency problem and react accordingly but take fair sharing into account.

The agents' behavior can be characterized by three features. First, they reject contracts that leave them with a low share of the surplus (**Observation**

7). This suggests that the behaviorally relevant 'participation constraint' is one that regards an 'equity constraint'. Second, agents respond, to a very large degree, optimally to the incentives set by the contracts (**Observation 8**). Whereas empirical analyses with field data often have difficulties to find a clear 'pay–for–performance' relation (see, e.g., Prendergast 1999), we received strong support for it in the laboratory: the higher the return share, the higher has been the effort level. Third, deviations from the optimal effort level can be explained by reciprocity which often contributed to further efficiency gains (**Observation 9**).

Our study of a simple PA–game has shown that people are to a non-negligible degree guided by individual rationality as proposed by standard economic theory. However, our results also demonstrate again that fairness - expressed in the willingness to reject 'unequal' offers - and reciprocity are other important behavioral patterns that are present at the same time. These behavioral principles have been repeatedly observed in many simple controlled laboratory experiments that were designed to study these motives in isolation. Our experiments demonstrate that they significantly influence behavior also in more complex agency relations. Moreover, and even more importantly, these motives do not contradict optimizing behavior: Most agents were at the same time prepared to choose at least the optimal effort level, or to even go beyond it for reciprocal reasons, and to reject very unfair offers. Hence, besides assuming optimizing behavior, agency theory should also take fairness and reciprocity as important determinants of behavior into account.

Acknowledgments: This paper is part of the EU-TMR Research Network ENDEAR (FMRX-CT98-0238). We also acknowledge support by the Swiss National Science Foundation under project number 1214-051000.97 (Gächter) and by the Deutsche Forschungsgemeinschaft DFG through SonderForschungsBereich 373 (Anderhub and Königstein). We have also benefitted from comments by Armin Falk, Ernst Fehr, Urs Fischbacher, Werner Güth and Peter Skott and by seminar participants at the Universities of Aarhus, Karlsruhe and Zurich, at the ESA-meeting in Mannheim, the Workshop for Corporate Governance in Berlin, the IAREP Workshop in Frankfurt/Oder and the Annual Meeting of the Swiss Society for Statistics and Economics in Fribourg. We are particularly grateful to Urs Fischbacher for writing the computer program and for his support during the experiments.

References

[1] Allen, D. and D. Lueck (1993): Transaction costs and the design of cropshare contracts, *RAND Journal of Economics* 24, 78-100.

[2] Alston, L. and R. Higgs (1982): Contractual Mix in Southern Agriculture since the Civil War: Facts, Hypotheses, and Tests, *Journal of Economic History* XLII, Vol. 2, 327-352.

[3] Berg, J., L. Daley, J. Dickhaut, J. O'Brien (1986): Controlling Preferences for Lotteries on Units of Experimental Exchange, *Quarterly Journal of Economics* 101, 281-306.

[4] Berg, J., L. Daley, J. Dickhaut, J. O'Brien (1992): Moral Hazard and Risk Sharing: Experimental Evidence, in: Mark Isaac (ed.), *Research in Experimental Economics,* Vol. 5, 1-34, Greenwich: JAI Press Inc.

[5] Bewley, T. (1997): *Why Not Listen to Business? A Study of Wage Rigidity.* Unpublished book manuscript, Yale University.

[6] Bhattacharyya, S. and F. Lafontaine (1995): Double-sided moral hazard and the nature of share contracts, *RAND Journal of Economics* 26, 761-781.

[7] Bull, C., A. Schotter, and K. Weigelt (1987): Tournaments and Piece Rates: An Experimental Study, *Journal of Political Economy* 95, 1-33.

[8] Camerer, C. (1995): Individual Decision Making, in: J. Kagel and A. Roth (eds.): *The Handbook of Experimental Economics.* Princeton: Princeton University Press.

[9] Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler (1997): Labor Supply of New York City Cab Drivers: One Day at a Time, *Quarterly Journal of Economics* CXII, 407-441.

[10] Camerer, C. and R. Thaler (1995): Anomalies: Ultimatums, Dictators, and Manners, *Journal of Economic Perspectives* 9, 209-219.

[11] Chaudhuri, A. (1998): The ratchet principle in a principal agent game with unknown costs: an experimental analysis, *Journal of Economic Behavior and Organization* 37, 291-304.

[12] Chisholm, D. (1997): Profit-Sharing Versus Fixed-Payment Contracts: Evidence From the Motion Pictures Industry, *Journal of Law, Economics, and Organization* 13, 169-201.

[13] DeJong, D., R. Forsythe, and R. Lundholm (1985): Ripoffs, Lemons, and Reputation Formation in Agency Relationships: A Laboratory Market Study, *Journal of Finance* 40, 809-823.

[14] DeJong, D., R. Forsythe, R. Lundholm, and W. Uecker (1985): A Laboratory Investigation of the Moral Hazard Problem in an Agency Relationship", *Journal of Accounting Research* 23 (Supplement) 81-120.

[15] Epstein, S. (1992): Testing Principal-Agent Theory, in: Mark Isaac (ed.), *Research in Experimental Economics,* Vol. 5, 35-60, Greenwich: JAI Press Inc.

[16] Fehr, E. and S. Gächter (1998): Reciprocity and economics: The economic implications of *Homo Reciprocans, European Economic Review* 42, 845-859.

[17] Fehr, E., S. Gächter, and G. Kirchsteiger (1997): Reciprocity as a Contract Enforcement Device - Experimental Evidence, *Econometrica* 65, 833-860.

[18] Fischbacher, U. (1998): z-Tree: Zurich Toolbox for Readymade Economic Experiments. *Mimeo,* University of Zurich.

[19] Grossman, S. and O. Hart (1983), An Analysis of the Principal–Agent Problem, *Econometrica*, Vol. 51, 7–45.

[20] Güth, W. (1995): On ultimatum bargaining experiments - a personal review, *Journal of Economic Behavior and Organization* 27, 329-344.

[21] Güth, W., W. Klose, M. Königstein and J. Schwalbach (1998): An Experimental Study of a Dynamic Principal-Agent Relationship, *Managerial and Decision Economics* 19, 327-341.

[22] Güth, W., M. Königstein, J. Kovács, and E. Zala (1999): Fairness within firms: The case of one principal and multiple agents, *Working paper,* Humboldt University Berlin.

[23] Hackett, S. (1993): Incomplete Contracting: A Laboratory Experimental Analysis, *Economic Inquiry* XXI, 274-297.

[24] Hart, O. and B. Holmström (1987): The theory of contracts. in: T. Bewley (ed.), *Advances in Economic Theory. Fifth World Congress.* Cambridge: Cambridge University Press.

[25] Holmström, B. (1979), Moral Hazard and Observability, *Bell Journal of Economics*, Vol. 10, 74–91.

[26] Holmström, B. and P. Milgrom (1987): Aggregation and Linearity in the Provision of Intertemporal Incentives, *Econometrica* 55, 231-259.

[27] Jensen, M. and W. Meckling (1976): Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure, *Journal of Financial Economics* 3, 305-360.

[28] Keser, C. and M. Willinger, (1997): Principal-agent relations with hidden actions: an experimental investigation. Forthcoming in: *International Journal of Industrial Organization.*

[29] Loomes, G. (1998): Probabilities vs Money: A Test of Some Fundamental Assumptions About Rational Decision Making, *Economic Journal* 108, 477-489.

[30] Mathewson, G. F. and R. Winter (1985): The Economics of Franchise Contracts, *Journal of Law & Economics* XXVIII, 503-526.

[31] Nalbantian, H. and A. Schotter (1997): Productivity Under Group Incentives: An Experimental Study, *American Economic Review* 87, 314-341.

[32] Plott, C. and L. Wilde (1982): Professional Diagnosis vs. Self-Diagnosis: An Experimental Examination of Some Special Features of Markets with Uncertainty, in: V. Smith (ed.), *Research in Experimental Economics,* Vol. 2, 63-11, Greenwich: JAI Press Inc.

[33] Prendergast, C. (1999): The Provision of Incentives in Firms, *Journal of Economic Literature* XXXVII, 7-63.

[34] Ross, S. (1973): The economic theory of agency: the principal's problem, *American Economic Review* 63(2), 134-139.

[35] Roth, A. (1995): Bargaining Experiments, in: J. Kagel and A. Roth (eds.): *Handbook of Experimental Economics.* Princeton: Princeton University Press.

[36] Salanié, B. (1997): *The Economics of Contracts. A Primer.* The MIT Press.

[37] Selten, R.(1975): Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, *International Journal of Game Theory*, Vol. 4, 25-55.

[38] Selten, R., A. Sadrieh, and K. Abbink (1995): Money Does Not Induce Risk Neutral Behavior, but Binary Lotteries Do even Worse, *Discussion paper No. B-343,* University of Bonn.

[39] Shavell, S. (1979): Risk sharing and incentives in the principal and agent relationship, *Bell Journal of Economics* 10, 55-73.

[40] Slonim, R. and A. Roth (1998): Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic, *Econometrica* 66, 569-596.

Figures

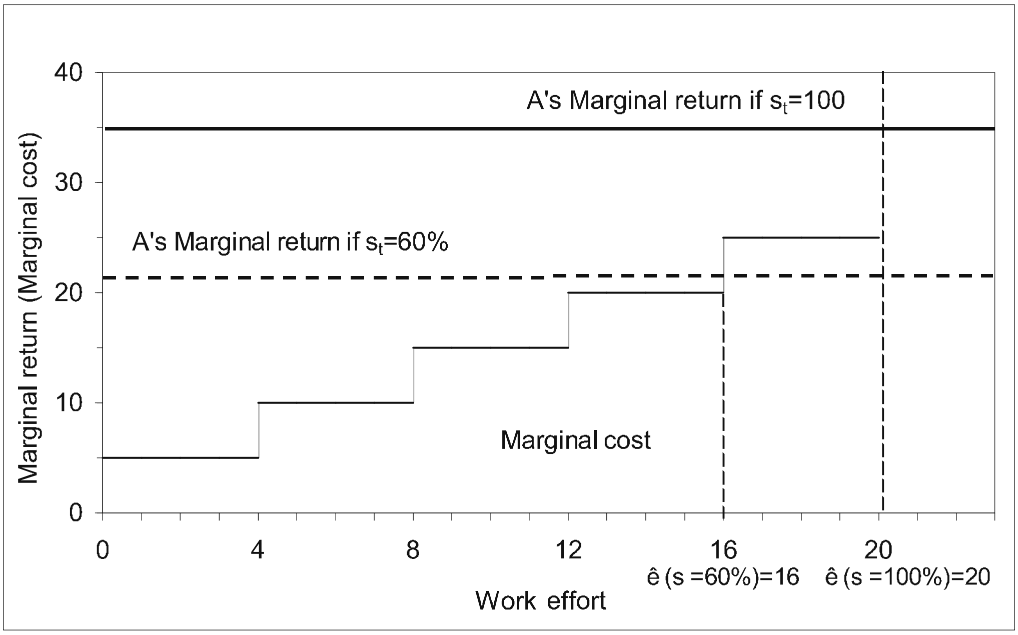

Figure 1: Marginal return function and marginal cost function as well as conditionally rational effort choices for $s = 60\%$, respectively $s = 100\%$.

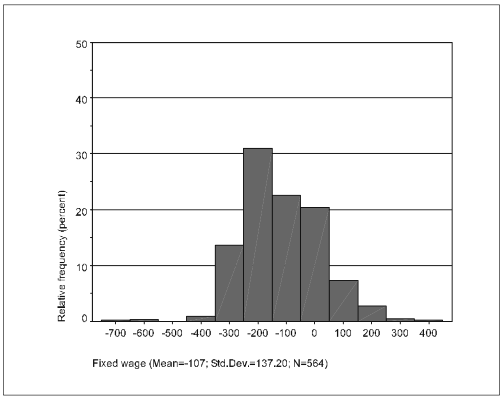

Figure 2a: Distribution of offered fixed wages

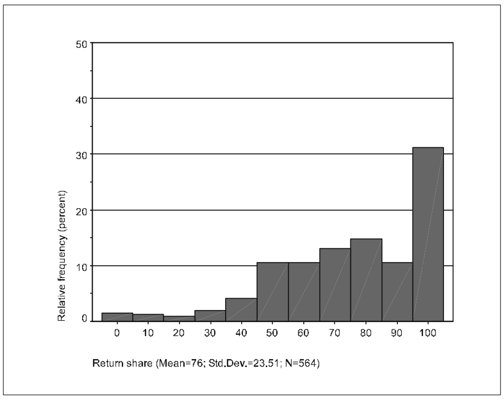

Figure 2b: Distribution of offered return shares

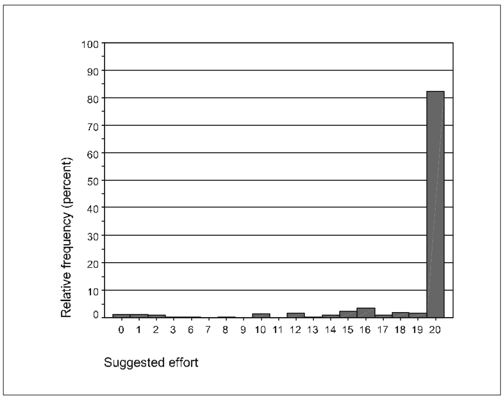

Figure 2c: Distribution of suggested effort

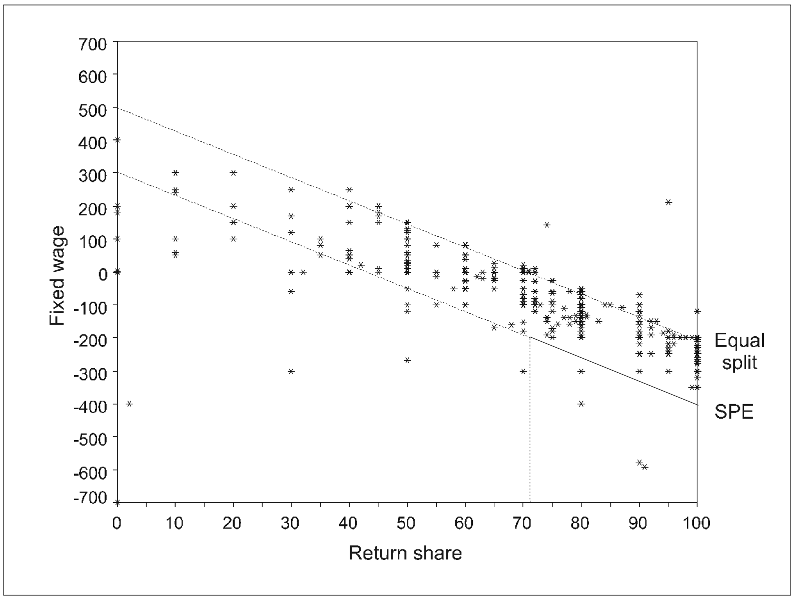

Figure 3: Combinations of return share and fixed wage as offered by the principals

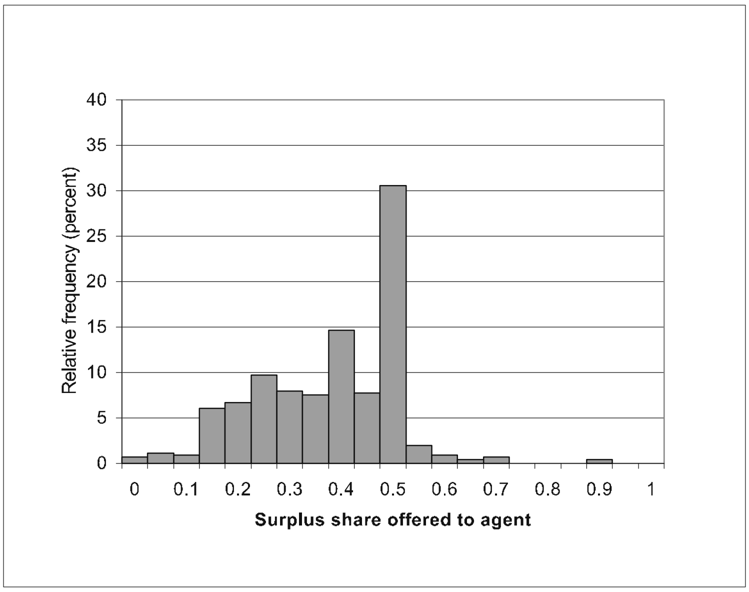

Figure 4: Surplus share offered to the agent (for efficient contract offers)

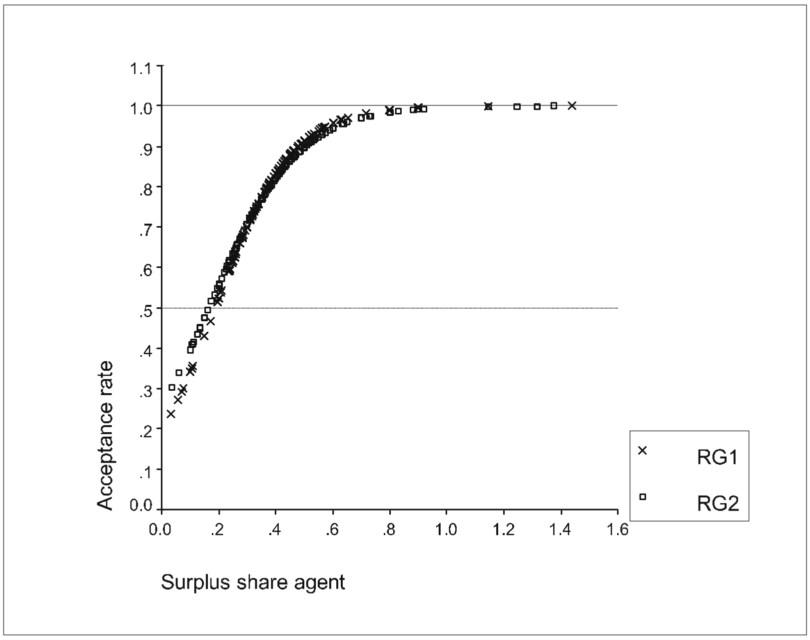

Figure 5: Predicted acceptance rate (probability of $\delta = 1$) conditional on the surplus share offered to the agent according to logit regressions for RG1, respectively RG2

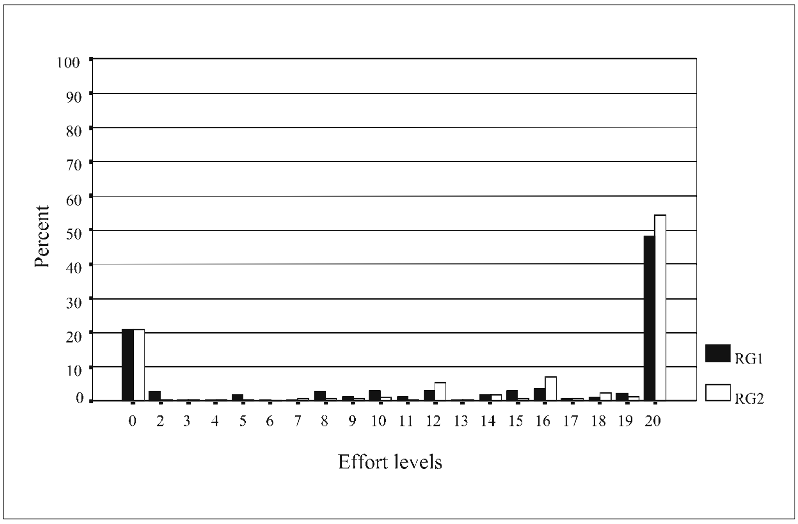

Figure 6: Frequency distribution of agents' actual effort choices in both sequences of the repeated games (RG1 and RG2)

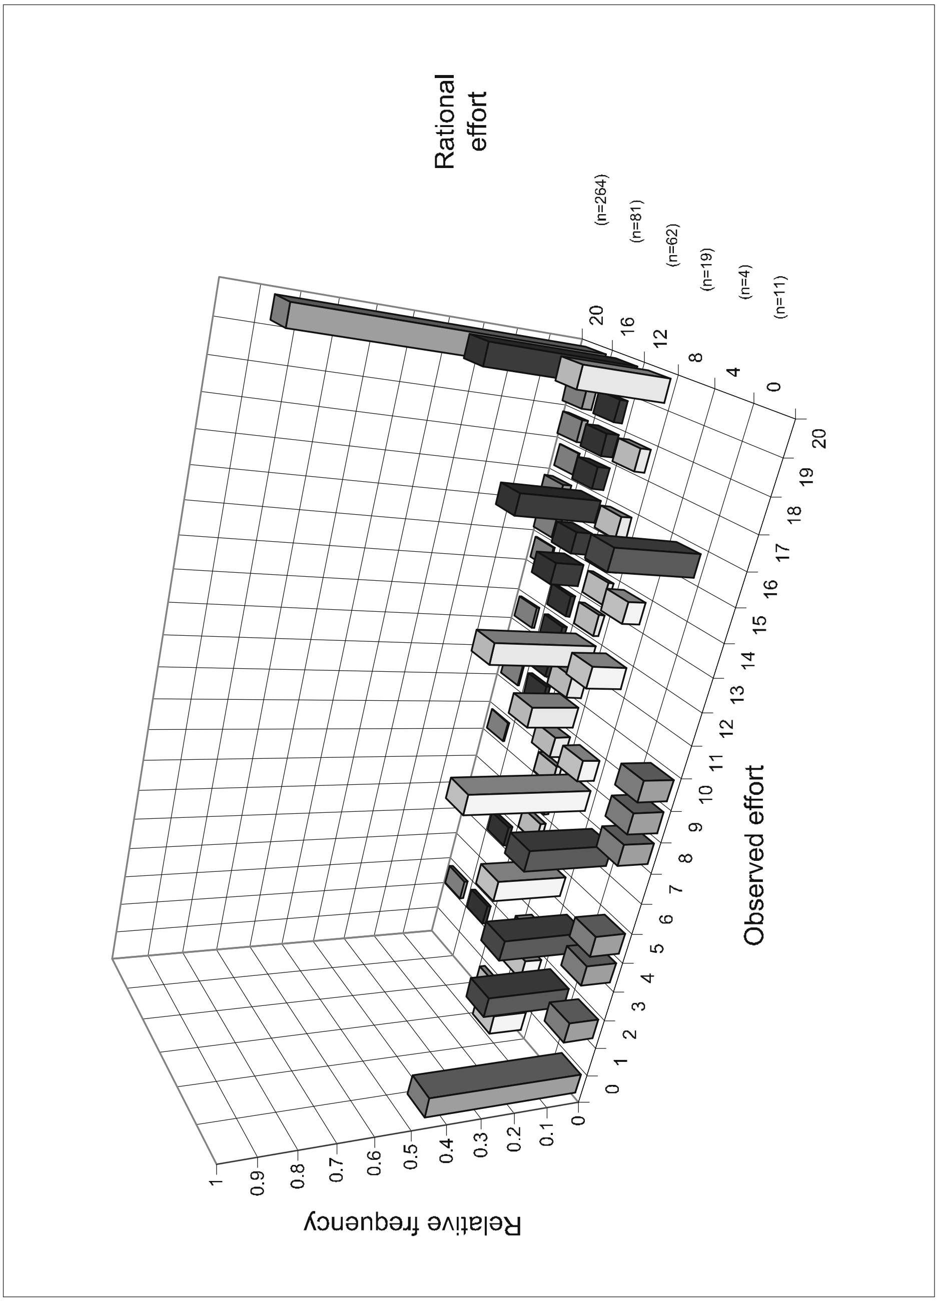

Figure 7: Conditionally rational effort choices

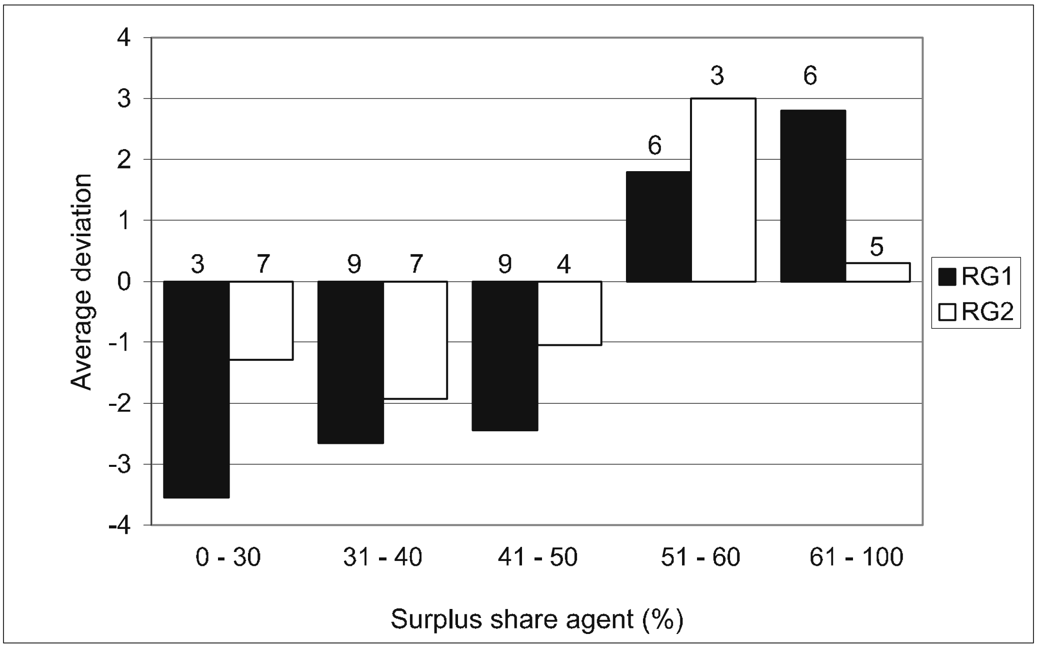

Figure 8: The number of individuals who deviated from rational effort $(e - \widehat{e})$ for a given surplus share is indicated above the respective bar