

Erlei, Mathias; Schenk-Mathes, Heike

Working Paper

Bounded Rationality in Principal-Agent Relationships

TUC Working Papers in Economics, No. 06

Provided in Cooperation with:

Clausthal University of Technology, Department of Economics

Suggested Citation: Erlei, Mathias; Schenk-Mathes, Heike (2012) : Bounded Rationality in Principal-Agent Relationships, TUC Working Papers in Economics, No. 06, Technische Universität Clausthal, Abteilung für Volkswirtschaftslehre, Clausthal-Zellerfeld, <https://doi.org/10.21268/20161213-122216>

This Version is available at:

<https://hdl.handle.net/10419/107447>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Bounded Rationality in Principal-Agent Relationships

by Mathias Erlei & Heike Schenk-Mathes

Abstract: We conducted six treatments of a standard moral hazard experiment with hidden action. All treatments had identical Nash equilibria. However, the behavior in all treatments and periods was inconsistent with established agency theory (Nash equilibrium). In the early periods of the experiment, behavior differed significantly between treatments. This difference largely vanished in the final periods. We used logit equilibrium (LE) as a device to grasp boundedly rational behavior and found the following: (1) LE predictions are much closer to subjects' behavior in the laboratory; (2) LE probabilities of choosing between strategies and experimental behavior show remarkably similar patterns; and (3) profit-maximizing contract offers according to the LE are close to those derived from regressions.

JEL Classification: C72, C92, J31, L14

Keywords: experiment, logit equilibrium, moral hazard, hidden action

Version: January 5, 2012

1. Moral hazard with hidden action as an established theory in microeconomics

Agency problems have been analyzed theoretically in the last three decades under a variety of conditions, beginning with the studies of Ross (1973), Holmstrom (1979), and Shavell (1979). Standard theory postulates that the principal and the agent individually behave rationally, i.e., they have self-centered preferences and maximize expected utility. In the case of hidden action, after signing a contract, the agent takes an action or chooses an effort on behalf of the principal. The effort cannot be observed by the principal. A higher effort of the agent is associated with a higher expected outcome that accrues to the principal and results in higher effort costs for the agent. Since the outcome is exposed to risk, the principal cannot deduce the agent's effort ex post. Therefore, the asymmetric information with regard to the agent's effort induces the well-known moral hazard problem. The agent can choose an inefficiently low level of effort that the principal is unable to detect. Another version of the moral hazard problem arises when the agent's effort is not verifiable by a third party responsible for the enforcement of the contract between principal and agent. Again, the principal cannot force the agent to choose the efficient level of effort. The source of the moral hazard problem is not the stochastic outcome; rather, it is the non-enforceability of contractual agreements based on effort choice.

To mitigate or even prevent the moral hazard problem, the principal can offer a payment scheme to the agent in which payments depend on outcomes. The basic agency problem consists of determining the optimal payment scheme to maximize the principal's expected utility under two constraints. First, the incentive compatibility constraint is necessary since it is assumed that the agent chooses the effort level that maximizes his expected utility. Second, the participation constraint ensures that the agent obtains a sufficiently high payment to accept the offer.

The moral hazard problem is of practical relevance, particularly when linear payment schemes are considered (e.g., Holmstrom/Milgrom 1987, Spemann 1987). Assuming a stochastic outcome, the optimal solution is characterized by a trade-off between Pareto-efficient risk sharing and maximum motivation. The participation constraint is binding, i.e., the agent realizes only his reservation utility. If both contract parties are risk neutral, the agent will become the residual claimant. The agent bears all risk, and the principal receives a fixed payment. In this case, the first best solution is achievable. If the agent is risk averse, the optimal solution is only second-best and leads to the above-mentioned trade-off between risk sharing and motivation.

The moral hazard problem has received much attention in applied microeconomics, in topics such as labor economics, insurance economics, organizational design, and managerial accounting. Therefore, empirical studies are necessary to confirm the theoretical results. Unfortunately, field data (such as payment schemes for managers) are rarely available, and real-world contract design is much more complex than theoretical principal-agent models. Furthermore, it is impossible to control for all variables. Laboratory experiments are a more appropriate means of testing these theoretical predictions since the principals' contract offers and the agents' choices can be observed. In addition, important factors for the relationship, such as the effort costs or the choice sets of the principal and the agent, can be controlled.

In the early 1990s, experiments testing the principal-agent model with hidden action were conducted by Berg et al. (1990) and Epstein (1992), and the findings were in line with standard theory. However, their experimental design only allowed for two possible actions, and the set of feasible contracts that they used was extremely small. Later experimental results show that individuals often deviate from the predictions of standard principal-agent theory (Anderhub et al. 2002, Cochard and Willinger 2005, Fehr and Gächter 2001, Fehr et al. 2004, Güth et al. 1998, 2001, Keser and Willinger

2000, 2006). In these experiments, there are three principal findings concerning the subjects' behavior. First, the principals' contract offers are far more generous than those predicted by standard theory. Second, the agents' acceptance decisions deviate from individually rational behavior; i.e., they reject contract offers even if the offers fulfill the participation constraint. Third, the agents' effort choices are often best responses. In case of deviations, behavior can be explained by fairness norms or reciprocity.

In summary, actual behavior in experiments is sometimes better explained by fairness norms and reciprocity than by standard theory. However, these findings have several shortcomings. Similarly to the experiments conducted in the early 1990s, the agents in the more recent studies often only have two possible actions (e.g., Chernomaz 2011, Cochard and Willinger 2005, Keser and Willinger 2000, 2006). Most studies do not allow linear payment schemes and are restricted to incentive contracts consisting of the payment of a fixed wage and a fine that the agent has to pay in case of shirking.

Typically, shirking cannot be verified with certainty (e.g., Fehr and Gächter, 2001, Fehr et al. 2004)¹. Our experiment is most similar to the studies of Anderhub et al. (2002) and Güth et al. (2001). In those studies and in the other studies mentioned above, rational behavior is assumed, and deviations from standard theory are explained by introducing social preferences. Our objective is to determine whether bounded rationality provides a better explanation of the experimental data.

We conducted six treatments that included a principal-agent relationship with hidden (non-verifiable) action as a key element. We varied factors such as the procedure for role assignment or the matching of participants. During the initial periods, there were significant differences in the participants' behavior between the treatments, but these differences almost vanished in the final period. Our analysis concludes that standard theory, i.e., the subgame perfect Nash equilibrium concept, cannot satisfactorily explain the laboratory behavior. We use logit equilibrium to show that its predictions are remarkably close to the observed behavior.

¹ Their focus is not on testing the principal-agent model; rather, they test the crowding-out effect of material incentives.

The paper is organized as follows: Section 2 introduces and discusses the experimental design. Section 3 presents our experimental results. Section 4 explains the experimental data using the concept of logit equilibria, and Section 5 concludes.

2. Six experimental variations of the standard principal-agent model

2.1. Experimental design

We conducted six treatments. In each treatment we considered the following principal-agent setting:

At stage 1, the principal offers a contract including a fixed wage (*wage*) and a revenue share (*share*) in %. The revenue (*r*) depends on the agent's effort (*e*) and is given by $r(e) = 35 \cdot e$. For the *wage* and the *share*, the following restrictions must hold:

$$wage \in \{-700, -699, \dots, 699, 700\}$$

$$share \in \{0, 1, \dots, 99, 100\}.$$

In addition, the principal suggests a work effort (*suggested effort*) to the agent that is not binding on the agent.

At stage 2, the agent can accept or reject the offer. If the agent rejects the offer, both the principal and the agent earn a payoff of zero. Agents that accept the offer enter stage 3 and choose a work effort $e \in \{0, 1, \dots, 50\}$ that induces private costs of $c(e) = \frac{7}{8} e^2$.

Therefore, the payoffs of the principal, π^P , and the agent, π^A , are given by

$$\pi^P = \left(1 - \frac{share}{100}\right) \cdot 35e - wage$$

$$\pi^A = \frac{share}{100} \cdot 35e + wage - \frac{7}{8} e^2$$

Each of the six treatments comprised six periods. The participants were divided into groups of two members. One member took the role of the principal, and the other took the role of the agent. The six treatments differed in the assignment of the roles. In two of the six treatments, the roles were randomly assigned. In two other treatments, we combined the principal-agent setting with an auction to assign the role of the principal to the member giving the highest bid. In the remaining two treatments, we utilized a real-effort play to determine the principal of the group. The real-effort play will be explained in detail below.

At the beginning of a session of the *Baseline treatment*, the groups were randomly composed, and the roles were randomly assigned. In each period, the three stages of the principal-agent setting were played. The randomly composed groups and the randomly assigned roles remained the same during each session of the treatment.

The *Baseline One Shot treatment* was identical to the Baseline treatment, except that the groups were dissolved and randomly recomposed after each period.

At the beginning of a session of the *One Auction treatment*, the groups were composed. Then, the two members of each group played a second price-sealed bid auction. The member with the highest bid was given the role of the principal during the following six periods of the session and paid a price amounting to the bid of the partner, who became the agent. In case that the bids in a group were the same, the member who took the most time for bidding was given the role of the agent. The groups remained the same, and there was no role switch during a session.

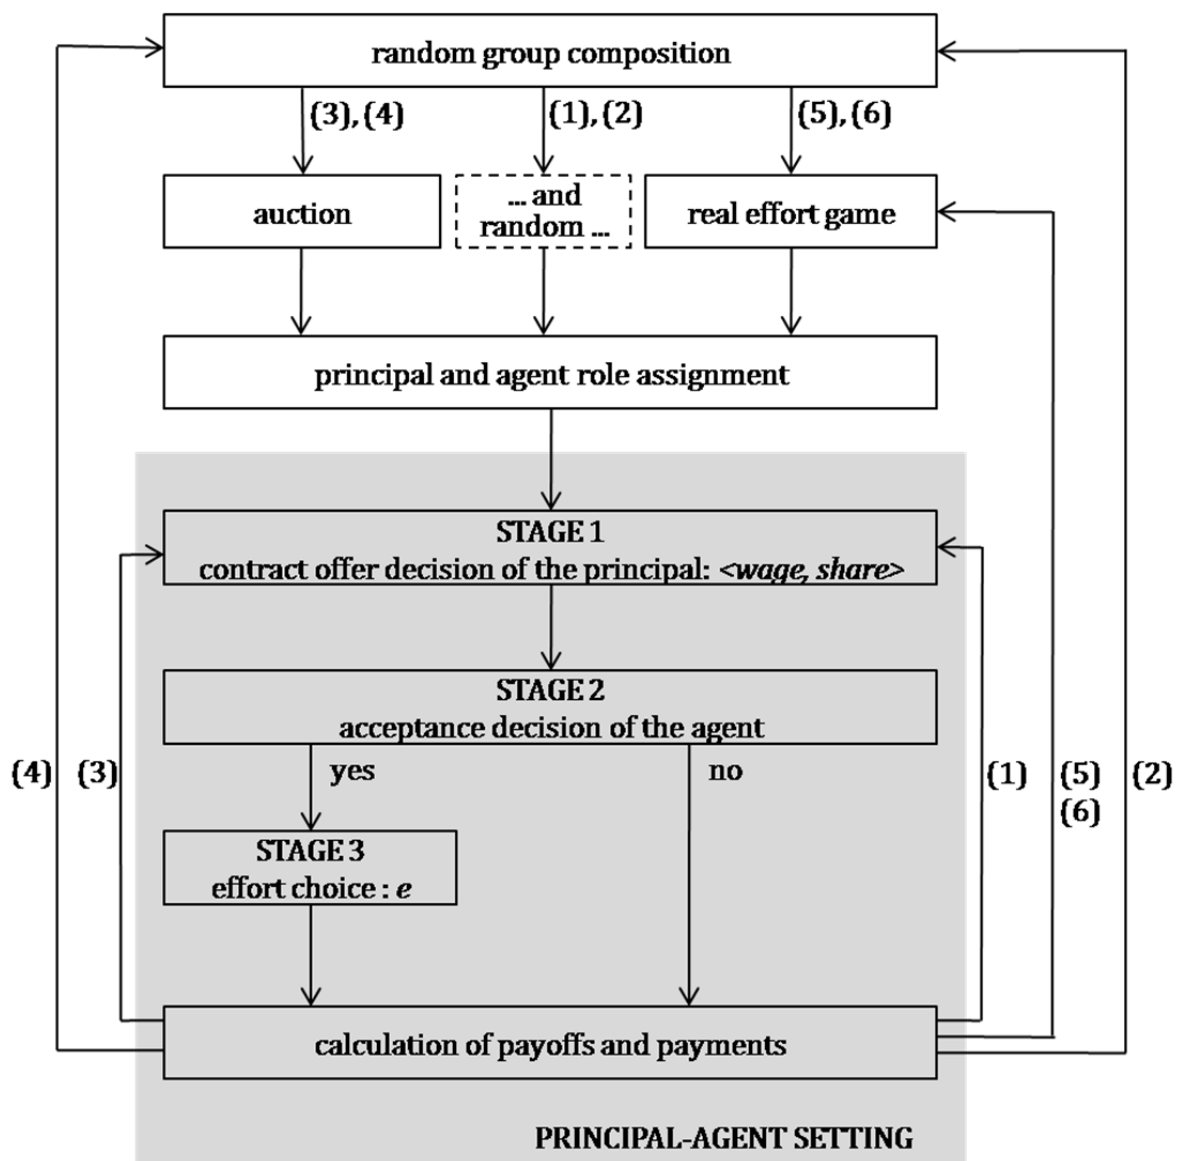
At the beginning of each period in the *Repeated Auctions treatment*, the groups were randomly composed, and roles were assigned by a second price-sealed bid auction.

At the beginning of each period of the *Real Effort treatment*, the subjects had to perform two tasks that represented graphical optimization problems (see for example van Dijk et al. (2000)). They had to find by trial and error the highest value of a function with two variables. Starting at the origin, the subjects could increase or decrease the value of the variables in discrete steps of 1. After each move, the value of the function was indicated, and the subjects had to wait 3 seconds for the next move. The time lag was introduced to ensure that no advantage was provided to experienced players of computer games. In every period, the search lasted 60 seconds. The subjects worked on two tasks of the same type, A and B, but the parameters of the function differed: the optimal values of the variables leading to the maximum of the function were not the same in the two tasks (and also changed from period to period). The subjects began work on Task A and then could switch tasks whenever they wanted. Task B was rewarded at an individual piece rate. Based on the result of Task A, the 20 subjects of a session were ranked in each period. Then, groups of two members were composed according to the following rule: The subject ranked number i , $i \in \{1, 2, \dots, 10\}$, and the subject ranked number $i+10$ formed a group. Finally, in each group, the subject with the higher score for Task A was given the role of the principal. The other subject took the role of the agent for this period. The

subjects could work on Task A to directly achieve a higher payment, or they could spend time on Task B to increase the probability of becoming a principal. Since the equilibrium payoffs of the principal are higher than the agents' payoffs, becoming a principal was valuable. In this treatment, the principal was named the employer and the agent was named the employee.

Finally, in the *Real Effort No Framing treatment*, the terms “employer” and “employee” were substituted with the neutral terms “participant X” and “participant Y.” This was the only difference between the two Real Effort treatments.

Figure 2 shows a summary of the experimental design.



(1) Baseline, (2) Baseline One Shot, (3) One Auction
(4) Repeated Auctions, (5) Real Effort, (6) Real Effort No Framing

Figure 2: The six treatments considered

2.2. Experimental procedures

The experiment was conducted at Clausthal University of Technology and consisted of 12 sessions, 2 for each treatment. In all treatments but one, 40 subjects participated. The Baseline One Shot treatment was the exception with only 36 subjects. Each subject participated in only one session. The subjects were separated from each other by blinders, and they remained anonymous.

The experiment was computerized and conducted with the help of the z-Tree experimental software developed by Fischbacher (2007). At the beginning of a session, the instructions appeared on a computer screen and were simultaneously read aloud by the experimenter. The participants had to answer a set of control questions.

Furthermore, before the 6 payoff periods began, the subjects completed two training periods to become familiar with the rules of the experiment.

The sessions lasted between 1.5 and 2 hours. The payment was 24.66 €, on average.

2.3. Discussion of the design

The experimental design is closely related to that of Anderhub et al. (2002) and Güth et al. (2001). The game of Anderhub et al. (2002) has multiple equilibria with shares between 500/7 and 100 percent resulting from a piecewise linear cost function, and the optimal effort constitutes a corner solution. We enlarged the set of feasible efforts to obtain an interior optimal solution and introduced a strictly convex cost function that leads to two subgame perfect equilibria (both with an optimal share of 100 percent). Compared with Güth et al. (2001), our sets of feasible contracts and feasible efforts are larger, but they remain discrete choice sets. In contrast to our study, Güth et al. (2001) analyze agency relationships with one principal and two agents. Their focus is not only on vertical but also on horizontal fairness. An important difference between our study and both previous studies is the embedding of the identical principal-agent relationship in different treatments (see figure 2).

As in Anderhub et al. (2002) and Güth et al. (2001), and in contrast to the other above mentioned studies, we consider a non-stochastic outcome. Risk preferences are therefore irrelevant, which allows us to concentrate on motivation devices. However, the non-stochastic outcome leads to another problem: although the effort is not observable,


it can be calculated ex post after observing the outcome. The theoretical solution to the motivation problem thus consists of a forcing contract in which the agent has to pay a fine in case of shirking. Therefore, we refer to non-verifiability of the effort, and our experimental design does not allow of forcing contracts. Hence, the reason for the moral hazard problem is not asymmetric information about the effort between the principal and the agent; rather, it is that the principal cannot punish the shirking of the agent and must motivate the agent with a linear payment scheme.

3. Experimental results

In this section, we provide an overview of the main characteristics of behavior in our six treatments. Working backwards through the game tree of our experiment, we start with the agents' effort decisions (Section 3.1), and we continue with their acceptance decision (Section 3.2) and the principals' contract offers (Section 3.3).

3.1. Effort

Effort choices according to rational behavior within the different subgames vary with the underlying contracts. Sequentially rational behavior demands effort to be chosen according to $e_i^{opt} = \frac{share_i}{5}$. In order to compare effort choices under differing contracts, we define a new variable, effort deviation, as $\Delta e = e_i - e_i^{opt}$. According to the game theoretical prediction, Δe is expected to equal zero, or it should at least converge to zero. Figure 3.1 provides a box plot of effort deviations with respect to all six treatments and all six periods. Each box plot consists of a box that is supplemented with two vertical lines (the whiskers) and (sometimes) some dots. The dots represent individual outliers. The upper (lower) vertical lines represent the range of the upper (lower) quartile of effort deviations. The box contains the inner quartile range of effort deviations. Inside the boxes, there is a horizontal line showing the median effort deviation.



Note: diff_e describes the values of Δe . Numbers 1 through 6 denote periods.

Figure 3.1: Box plot of effort deviations

Figure 3.1 clearly shows a convergence towards rational effort choices. In period 2, the median effort deviation is very close to zero. In period 6, it is exactly zero in all treatments. Furthermore, the dispersion of the effort deviation in the final period is very small. In the Baseline treatment, the remaining box representing the inner quartile range is hardly visible at all, which means that almost all choices in Period 6 correspond to the individually optimal efforts.

Table A1 given in the appendix shows the means and medians for all treatments and all periods. This table strongly confirms our interpretation of Figure 3.1. In particular, all medians in period 6 equal zero, and the means fluctuate closely around zero. Therefore, we can state our first result.

Result 1: *Effort decisions are close to the corresponding profit maximizing values.*

3.2. Contract acceptance behavior

According to the game theoretical predictions, agents will accept all contract offers that enable them to realize nonnegative profits. This condition is fulfilled whenever contract terms are such that $wage_i \geq -\frac{7}{200} share_i^2$. In order to summarize the acceptance behavior, we define classes of contract offers with respect to (fixed) wage intervals and (return) share intervals. Each wage interval is labeled according to its upper limit, and each share interval is labeled according to its upper value as a percentage. Table 3.1 shows the fraction of accepted offers in 84 classes.

Wage classes	Share classes						
	0	1...25	26...40	41...60	61...75	76...90	91...100
-750...-351	0	0		0.5			0
-350		0.50			0	0	
-349...-250		1.00	0	0.2	0		0.74
-249...-150		0.14	0	0	0.89	0.72	0.81
-149...-50		0.50	0.33	0.625	0.82	0.91	1
-49...-1		0	0.33	0.95	1		1
0		1.00	0.64	0.86	0.96	1	
1...50		0.83	1	0.96	1		
51...150	1	0.89	0.96	0.97	1		
151...250		0.95	0.93	1	1	1	1
251...350	1	1	1	0.80			1
351...750		0.85	1	0.80	1	1	

Note: Numbers for share intervals are percentages.

Table 3.1: Fractions of accepted offers

Columns 2 through 8 show a thick horizontal line demarcating the minimum fixed wage that is consistent with a rational acceptance of the contract offer. In other words, all cells above this demarcating line offer too small wages so that theory expects agents to reject the contract with a probability of one. In contrast, whenever wages are larger than zero, standard theory predicts that offers will always be accepted. Theory is somewhat confirmed by the data, because most numbers above the demarcating lines are close to zero, whereas numbers below are close to 1. The “wrong acceptances”, i.e., those cells above the demarcation line with a fraction strictly greater than zero, add up to only seven cases. None of these cases occurs in period 6. Our interpretation is therefore that the incorrect acceptances are part of the participants’ learning processes.

Result 2: *The behavior with respect to accepting or declining contract offers is largely consistent with equilibrium predictions.*

3.3. Contract offers

The experimental game considered in this paper has two subgame perfect Nash Equilibria with two rather similar contract offers. One equilibrium shows a fixed wage of -350 and a share of 100 percent. In the other equilibrium, the wage is -349, and the share again equals 100 percent. Accordingly, equilibrium predicts that principals will always choose a share of 100 percent and a wage of -349 or -350. However, in 708 cases, these equilibrium contract offers have not been proposed even once.

Table 3.2 shows the distribution of the contract offers. This distribution may be characterized by three properties: (1) There are two local peaks. One peak is characterized by shares between 90 percent and 100 percent and fixed wages between -249 and -150. The other peak corresponds to less asymmetric payoff distributions with shares between 41 percent and 60 percent and wages between -49 and +50.

(2) In each column of Table 3.2, we highlight the two cells with the highest frequencies. Combining them leads to an upward-sloping area within the table that starts with wages between 1 and 100 in the second column and ends with wages between -299 and -150. The peaks of the contract distribution lie within this range. (3) Relative frequencies out of the high-frequency range are often close to zero.

Such a distribution of contract offers cannot be explained by conventional subgame perfect Nash Equilibrium because it only contains wages of -349 and -350 in combination with a share of 100 percent, i.e., contract offers which have not been chosen at all.

Result 3: *The distribution of contract offers is inconsistent with the perfect equilibrium of the underlying principal-agent model.*

Wage	Share										Total
	10	20	30	40	50	60	70	80	90	100	
-650	0.56	0.00	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.71
-600	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-550	0.00	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.14
-500	0.14	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.14	0.42
-450	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-400	0.14	0.28	0.14	0.00	0.14	0.14	0.00	0.00	0.00	0.00	0.85
-350	0.42	0.14	0.00	0.00	0.00	0.00	0.00	0.28	0.00	0.00	0.85
-300	0.28	0.00	0.14	0.00	0.42	0.14	0.28	0.00	0.00	0.42	1.69
-250	0.00	0.00	0.28	0.00	0.14	0.00	0.00	0.00	0.00	2.82	3.25
-200	0.00	0.14	0.56	0.00	0.42	0.00	0.14	0.28	1.27	5.23	8.05
-150	0.00	0.14	0.28	0.14	0.28	0.14	0.14	1.69	1.27	3.81	7.91
-100	0.14	0.00	0.28	0.42	0.99	0.14	0.99	1.41	0.56	0.14	5.08
-50	0.00	0.14	0.14	0.00	2.12	1.27	1.13	0.42	0.14	0.00	5.37
0	0.56	0.28	0.71	<u>1.84</u>	6.50	4.52	1.98	1.84	0.28	0.14	18.64
50	1.13	0.85	1.98	2.68	4.66	2.12	0.42	0.14	0.00	0.00	13.98
100	1.55	1.27	3.39	2.82	2.40	0.85	0.28	0.00	0.00	0.00	12.57
150	0.28	0.71	1.41	1.84	0.85	0.42	0.00	0.00	0.00	0.00	5.51
200	0.14	1.55	0.42	1.13	1.84	0.00	0.14	0.14	0.14	0.14	5.65
250	0.28	0.42	0.71	0.42	0.42	0.42	0.00	0.28	0.00	0.00	2.97
300	0.14	0.28	0.28	0.28	0.42	0.14	0.00	0.00	0.00	0.14	1.69
350	0.28	0.42	0.28	0.00	0.14	0.00	0.00	0.00	0.00	0.00	1.13
400	0.42	0.28	0.00	0.00	0.28	0.00	0.00	0.00	0.00	0.00	0.99
450	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.14	0.00	0.00	0.28
500	0.28	0.00	0.00	0.00	0.14	0.00	0.00	0.00	0.14	0.00	0.56
550	0.28	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.28
600	0.14	0.14	0.00	0.00	0.00	0.14	0.14	0.00	0.00	0.00	0.56
650	0.00	0.00	0.14	0.00	0.00	0.00	0.00	0.14	0.14	0.00	0.42
700	0.00	0.00	0.14	0.00	0.14	0.00	0.00	0.14	0.00	0.00	0.42
Total	7.34	7.34	11.44	11.58	22.32	10.45	5.65	6.92	3.95	12.99	100

Notes: Wage and share classes are labeled by their upper limits. Wages are absolute numbers, shares are percentages.

Table 3.2: Distribution of contract offers.

Neglecting the different mechanisms for assigning the roles in our experiment, we can state that all treatments have identical perfect equilibria. From a theoretical point of view, we thus expect to see no major differences between treatments. Table 3.3 shows the means and medians of contract characteristics with respect to the six treatments.

Treatment	Mean (wage)	Mean (share)	Median (wage)	Median (share)
Baseline	-38.1	59.0	-30	60
One Auction	26.9	58.3	0	55
Baseline One Shot	-10.5	59.1	0	50
Real Effort	53.1	46.7	50	50
Real Effort No Framing	-54.1	47.9	-5	50
Repeated Auctions	26.1	41.2	23.5	45

Table 3.3: Means and medians of contract terms (all periods)

Differences are obviously present between the treatments. The mean fixed wages are below zero only in the treatments Baseline, Baseline One Shot, and Real Effort No Framing. The mean shares in Baseline, One Auction, and Baseline One Shot are close to 60, whereas they are well below 50 in the remaining treatments. Kruskal-Wallis equality-of-population tests for the fixed wage in period 1 reject the null hypotheses (equal populations) with p-values of 0.0042 (all treatments) and 0.0279 (all treatments except Repeated Auctions).²

It is an open question whether behavior in treatments is similar above all periods. Another open question is whether contract offers finally converge. Therefore, we also present summary statistics of contract offers in the final period.

Treatment	Mean (wage)	Mean (share)	Median (wage)	Median (share)
Baseline	-29.6	52.7	0	55.0
One Auction	-13.9	65.4	0	68.5
Baseline One Shot	-88.2	64.9	-95	55.0
Real Effort	-45.3	57.3	0	52.5
Real Effort No Framing	-83.5	60.0	0	50.0
Repeated Auctions	-27.9	44.4	5.5	47.5

Table 3.4: Means and medians of contract terms (period 6)

In period 6, the mean wages are negative in all treatments. Moreover, the mean shares of Real Effort and Real Effort No Framing have increased substantially and can hardly be

² As will be seen later, the Repeated Auctions treatment diverges somewhat from the other treatments, which is why we test for all treatments except Repeated Auctions.

distinguished from the first three treatments. The share only remains well below 50 in the Repeated Auctions treatment. Our interpretation of these data is that the participants underwent a learning process that induced a convergence of behavior. However, in the Repeated Auctions treatment, this learning process seems to have been slower than in the other treatments.

Result 4: *We find large and significant treatment effects in early periods and little evidence for treatment effects in the final period.*

Our interpretation that participants underwent a learning process is corroborated by the distribution of contracts in period 6 given in Table 3.5. Contracts outside the largely unchanged high-frequency range were chosen less frequently, and the relative frequencies of most contract terms inside the area increased. Note that we cannot observe a convergence toward the perfect equilibrium, which lies outside the high-frequency range.

The distribution of contract terms is thus very similar to the distribution for all periods, with the exception that there seems to be much less noise.

Result 5: *There is no convergence towards the perfect equilibrium.*

In summary, we find a fairly stable pattern of contract offers that is clearly inconsistent with the equilibrium prediction. The differences between treatments prove to be unstable. We interpret the treatment effects as phenomena of bounded rationality and learning that diminish with players' experience. Consequently, we have not provided an explanation for subjects' behavior in our experiment yet. In the following section, we try to provide such an explanation by referring to an equilibrium concept that explicitly accounts for bounded rationality and that may be interpreted as the endpoint of a learning process, i.e., the logit equilibrium.

Wage	Share										Total
	10	20	30	40	50	60	70	80	90	100	
-650	0.85	0.00	0.85	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.69
-600	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-550	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-500	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-450	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
-400	0.00	0.00	0.00	0.00	0.00	0.85	0.00	0.00	0.00	0.00	0.85
-350	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.85	0.00	0.00	0.85
-300	0.00	0.00	0.00	0.00	0.00	0.00	0.85	0.00	0.00	0.85	1.69
-250	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	4.24	4.24
-200	0.00	0.00	0.85	0.00	0.00	0.00	0.00	0.00	3.39	6.78	11.02
-150	0.00	0.00	0.00	0.00	0.00	0.00	0.85	0.00	1.69	5.08	7.63
-100	0.00	0.00	0.00	0.00	8.85	0.00	0.85	0.85	0.85	0.00	3.39
-50	0.00	0.00	0.00	0.00	4.24	2.54	1.69	0.00	0.00	0.00	8.47
0	1.69	0.85	2.54	1.69	10.17	5.08	4.24	1.69	0.00	0.00	27.97
50	1.69	0.85	1.69	5.93	5.93	1.69	0.85	0.00	0.00	0.00	18.64
100	0.00	0.00	3.39	0.85	0.85	0.00	0.00	0.00	0.00	0.00	5.08
150	0.00	0.00	1.69	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.69
200	0.00	0.85	0.00	1.69	0.00	0.00	0.00	0.00	0.00	0.00	2.54
250	0.00	0.00	0.00	0.00	0.85	0.85	0.00	0.00	0.00	0.00	1.69
300	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
350	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
400	0.00	0.00	0.00	0.00	0.85	0.00	0.00	0.00	0.00	0.00	0.85
450	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
500	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
550	0.85	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.85
600	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
650	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.85	0.00	0.00	0.85
700	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	5.08	2.54	11.02	10.17	23.73	11.02	9.32	4.24	5.93	16.95	100

Notes: Wage and share classes are labeled by their upper limits. Wages are absolute numbers, shares are percentages. Framed cells show the local peaks from Table 3.2.

Table 3.5: Distribution of contract offers in the final period.

4. Explaining participants' boundedly rational behavior with logit equilibria

4.1. Logit equilibrium as a device for grasping bounded rationality

In Nash equilibria, players perfectly optimize their strategy choice given the strategies of all other players. They never choose suboptimal strategies, and they never make mistakes. Logit equilibrium is a game theoretic equilibrium concept that takes erroneous behavior into account. Logit equilibrium is a particular case of quantal response equilibria, a concept developed by McKelvey and Palfrey (1995, 1998). The essence of

the logit equilibrium is to transfer the idea of probabilistic choice (Luce 1959) into the spheres of game theory.

Let $E\pi_{ik}$ denote the expected payoff (or expected utility) of player i choosing strategy k , $k \in \{1, 2, \dots, K\}$, then player i 's probability of choosing k is given by

$$pr_{ik} = \frac{e^{\frac{E\pi_{ik}}{\mu}}}{\sum_{j=1}^K e^{\frac{E\pi_{ij}}{\mu}}} \quad (3.1)$$

Here, μ denotes the error parameter. If μ converges to infinity, the player behaves completely unintelligently and chooses a strategy according to a uniform distribution. In case of $\mu \rightarrow 0$, the player is perfectly rational. For intermediate values, each strategy is chosen with strictly positive probability, and strategies providing higher expected profits are chosen with higher probabilities.

In economic games, a strategy's expected payoff is usually also dependent on the other players' strategies, i.e., $E\pi_{ik} = \sum_{s_{-i}} pr^e(s_{-i}) \cdot \pi_{ik}(k, s_{-i})$, with s_{-i} denoting the strategy combination of all players without player i and $pr^e(s_{-i})$ denoting the player i 's expected probability that the other players choose this particular strategy combination. The logit equilibrium is then defined by two requirements. (1) The probability of each player i choosing his strategy k is given by (3.1) for all players and all strategies, and (2) all players' beliefs about the distribution of the other players' strategy combinations are correct, i.e., $pr^e(s_{-i}) = pr(s_{-i})$.

For each value of the error parameter μ there is a specific logit equilibrium. In most applications, the error parameter is estimated with maximum likelihood techniques. Note that in case of $\mu \rightarrow 0$, the logit equilibrium coincides with a specific Nash equilibrium. McKelvey and Palfrey (1998) extend the concept of quantal response equilibrium in general and the concept of logit equilibrium in particular to extensive form games. The logit equilibrium then corresponds to the logit equilibrium of the agent normal form of the underlying game. In case of $\mu \rightarrow 0$ the logit equilibrium converges to a sequential equilibrium.

Anderson, Goeree, and Holt (2004) show that the logit equilibrium can be interpreted as the steady state of a noisy directional learning process in which players are subject to normal errors. In this paper, we adopt this special interpretation so that the behavior in the final period will be interpreted as the endpoint of such a learning process.

Before applying the equilibrium concept to our own experiment, it is worth mentioning that numerous successful applications of logit equilibrium to specific contexts exist (Anderson et al. 1998, 2001, 2002; Capra et al. 1999; Goeree and Holt 2000, 2001; Goeree et al. 2002). In fact, quantal response equilibrium has been so successful that Camerer et al. (2004) suggest that “Quantal response equilibrium (QRE), a statistical generalization of Nash, almost always explains the direction of deviations from Nash and should replace Nash as the static benchmark to which other models are routinely compared.”

4.2. Logit equilibrium of our experimental games

As previously mentioned, the logit equilibrium of our experimental game varies with the error parameter μ . For example, if μ converges to zero, the distribution of contract offers will converge to a probability of one that the wage will be set to -349 and the share will approach 100 percent. In contrast, if μ approaches infinity, all possible contract offers will be chosen with the same probability. Our main task, therefore, is to estimate μ . We conducted a standard maximum likelihood estimation with 118 cases for contract offers, 118 decisions whether to accept the offer, and 97 effort choices in period 6. According to our estimation, the magnitude of the error parameter is $\mu = 52.56$. This value is significantly greater than zero ($p < 0.001$).³ Consequently, we reject the hypothesis of perfectly rational behavior. Table 4.1 shows the logit-equilibrium distribution with this error parameter.

³ The log-likelihood value of our estimation is -1597.50.

Wage	Share										Total
	10	20	30	40	50	60	70	80	90	100	
-650	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	1.0
-600	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	1.0
-550	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	1.0
-500	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	1.0
-450	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.2	1.1
-400	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.2	0.3	1.3
-350	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.2	0.2	0.6	1.7
-300	0.1	0.1	0.1	0.1	0.1	0.1	0.2	0.2	0.4	1.6	3.0
-250	0.1	0.1	0.1	0.1	0.1	0.2	0.2	0.4	1.1	3.7	6.1
-200	0.1	0.1	0.1	0.1	0.2	0.2	0.4	0.9	2.9	5.2	10.2
-150	0.2	0.2	0.2	0.2	0.2	0.4	0.8	2.3	4.7	4.1	13.3
-100	0.2	0.2	0.2	0.3	0.4	0.8	2.0	3.9	4.2	2.2	14.4
-50	0.3	0.3	0.3	0.5	0.9	1.7	3.0	3.5	2.4	1.0	13.9
0	0.4	0.4	0.6	0.9	1.4	2.2	2.6	2.0	1.0	0.4	11.9
50	0.5	0.5	0.7	1.0	1.4	1.6	1.4	0.9	0.4	0.1	8.5
100	0.4	0.4	0.6	0.7	0.8	0.8	0.6	0.4	0.2	0.1	5.0
150	0.2	0.2	0.3	0.3	0.4	0.3	0.2	0.1	0.1	0.0	2.1
200	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.0	0.0	0.8
250	0.0	0.0	0.0	0.1	0.1	0.1	0.0	0.0	0.0	0.0	0.3
300	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
350	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
400	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
450	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
500	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
550	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
600	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
650	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Total	3.2	3.2	3.9	5.0	6.7	9.1	12.1	15.4	18.2	19.8	

Notes: Wage and share classes are labeled by their upper limits. Wages are absolute numbers, shares are percentages.
The framed cells show the local peaks from Table 3.2.

Table 4.1: Distribution of contract offers according to logit equilibrium
($\mu = 52.56$)⁴

Comparing Table 4.1 with Table 3.5, we find that the distribution of contract offers according to logit equilibrium has a peak at wages between –150 and -200 with shares between 91 and 100 percent. We also find a similar high-frequency range of contract offers. Most probabilities outside this range are close to zero. Therefore, logit

⁴ The fact that total numbers don't add up to 100 is due to rounding errors.

equilibrium explains the distribution of contract offers far better than standard Nash equilibrium. This finding can also be illustrated by some summary statistics given in Table 4.2. The expected value of fixed wages and the expected value of revenue shares in the logit equilibrium are much closer to the data means than the corresponding values in the Nash equilibria.

	Mean or expected values of fixed wages	Mean or expected values of revenue shares
Data (period 6)	-47.36	57.31
Nash equilibrium	-349 / -350	100
Logit equilibrium	-143.15	67.14

Table 4.2: Summary statistics of contract offers

Although logit equilibrium performs better than Nash equilibrium, there are still two major shortcomings. First, we cannot reconstruct the second peak in the distribution of empirical contract offers. This second peak located around a wage of zero and a share of 50 percent is simply absent in our logit equilibrium. We conjecture that the second peak may best be explained as a focal point. A fixed wage of zero and a share of 50 percent are both medial values within the range of feasible values. Furthermore, this contract offer realizes an equal distribution of revenues. Two other contracts lying close to the medial contract are $\langle \text{wage} = 43.75; \text{share} = 50 \rangle$ and $\langle \text{wage} = 0; \text{share} = 66.67 \rangle$. Both offers induce a symmetric distribution of profits if the agent chooses individually optimal levels of effort.

Second, the mean of the fixed wages in period 6 is still substantially greater than the expected value of wages in the logit equilibrium. We suppose that this is largely due to a combination of subjects' dislike of offering negative wages and a rather slow learning process. Table 4.3 shows the empirical means of wages over periods. The wages are obviously decreasing and they may not have reached their final values in the final periods of our sessions.

Period	1	2	3	4	5	6
Means of wages	60.9	18.4	6.0	-8.4	-25.0	-47.4

Table 4.3: Means of fixed wages over periods

Logit equilibrium also has interesting implications for normative agency theory. Subgame perfect Nash equilibrium contracts are anything but optimal. According to the logic of logit equilibrium, a contract offer of $\langle \text{wage} = -350; \text{share} = 100 \rangle$ will be rejected with a probability of 62.2 percent, and the principal's expected profit is thereby reduced to 132.19. Even the focal contract $\langle \text{wage} = 0; \text{share} = 50 \rangle$ is superior to the Nash contract. As the focal contract will be accepted with a probability of 77.3 percent, its expected profits according to logit equilibrium amount to 140.18. However, the contract maximizing the principal's profit comprises a fixed wage of -253 and a share of 100 percent. In this case, the contract will be accepted with a probability of 79.4 percent, and the expected principal's profits equal 200.76. Logit equilibrium retains the insight that a large share of 100 is efficient. However, it also shows that it is important that the agent has a sufficiently high incentive for accepting the offer. The commonly used participation constraint, which states that it is sufficient for the principal to cover the agent's monetary opportunity cost, simply does not work for laboratory behavior. Increasing the agent's fixed wage, however, generates a real incentive for the agent to enter the contractual relationship.

5. Discussion

Both the empirical distribution of contract offers and the agents' acceptance behavior show that the distribution of surplus is important for understanding subjects' behavior. It therefore seems likely that social preferences might be involved. Inequity aversion (Fehr and Schmidt 1999; Bolton and Ockenfels 2000) and social welfare preferences (Charness and Rabin 2002) are the most prominent approaches for this kind of reasoning. It shows, however, that a mere change of the utility function that leaves the concept of Nash equilibrium uncontested cannot explain the above pattern of contract distributions. Inequity aversion and social welfare preferences both leave no room for contracts with shares below 100 percent, as we assumed a non-stochastic outcome. In inequity aversion models, any fairness aspect can be resolved by paying higher fixed wages. Any decrease in the agent's share creates an inefficiency that reduces both participants' utilities. The same argument holds for social welfare preferences. Correspondingly, the high-frequency range of contracts with shares below 100 percent cannot be explained.

It is still possible, however, to combine the concept of logit equilibrium with social preferences. We calculated other logit equilibria for different utility functions that

account for such preferences, and the high-frequency range of contract offers remains close to the basic logit equilibrium without social preferences. One important difference occurred, however: The high-frequency range is shifted downward towards the contract $\langle \text{wage} = -175; \text{share} = 100 \rangle$.

We conducted new estimations for the error parameter μ of the logit models with various specifications of social preferences. In doing this, we found out that introducing social preferences simultaneously in stage 2 (accepting or rejecting the offer) and in stage 3 (effort) decreases the log-likelihood value of the estimation. This also holds for the introduction of inequity aversion in stage 3 (effort). In contrast, introducing social welfare preferences (Charness and Rabin 2002) in stage 3 slightly increased the log-likelihood. Our overall judgment is that the introduction of social preferences has only a limited impact on our ability to explain the laboratory behavior in our experiment.

Finally, we discuss the usefulness of applying stage-specific error parameters in the logit model. Up to now, it is common practice to estimate one error parameter for each game. This parameter is meant to describe the subjects' degree of rationality, which varies with each game. This is sensible because different games are differently complicated such that each game needs its own error parameter (Anderson et al. 2002, 41). However, it is also true that different stages within a game may be differently complicated so that each decision stage may require a specific error parameter to account for the different characteristics of the corresponding decision problem (cf. Goeree et al. 2005, 363).

To see whether this has an impact on our results, we estimated an error parameter for each stage using the data of period 6. We started with the error parameter in stage 3 and used the estimated value to calculate the expected profits in earlier periods. We then estimated the error parameter of stage 2 and used this value for our final estimation of the error parameter of stage 1. This backward estimation procedure provides the error parameters $\mu_1 = 65.35$, $\mu_2 = 39.43$ and $\mu_3 = 40.58$. The magnitudes of the parameters are plausible because the decisions in stages 2 and 3 are clearly less difficult than the choice of a contract offer in stage 1. The log-likelihood values also improved. However, the theoretical logit distribution of contract offers changed only marginally.

Correspondingly, we confined ourselves to a single error parameter for the experimental game.

6. Conclusion

Our experimental data indicate that a learning process occurs and that the principals' contract offers in the final period are similar to the predictions made according to logit equilibrium. In contrast, the contract offer in subgame perfect Nash equilibrium is not chosen even once. Inequity aversion or social welfare preferences may play a role but do not explain the contract distribution in our experiment if combined with the assumption of rational behavior. In particular, contracts with shares less than 100 percent cannot be explained with utility functions including social preferences as long as the Nash equilibrium concept is applied.

We conclude that logit equilibrium is much better suited to predict behavior in principal-agent settings than subgame perfect Nash equilibrium so that we corroborate the suggestion of Camerer et al. that QRE "should replace Nash as the static benchmark to which other models are routinely compared."

References

- Anderhub, V., S. Gächter, and M. Königstein, 2002, "Efficient Contracting and Fair Play in a Simple Principal-Agent Experiment", *Experimental Economics*, 5, 5-27.
- Anderson, S. P., J. K. Goeree, and C. A. Holt, 1998, "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction", *Journal of Political Economy* 106, 828-53.
- Anderson, S. P., J. K. Goeree, and C. A. Holt, 2004, Noisy Directional Learning and the Logit Equilibrium, *Scandinavian Journal of Economics*, 106, 581-602.
- Anderson, S. P., J. K. Goeree, and C. A. Holt, 2001, "Minimum-Effort Coordination Games: Stochastic Potential and Logit Equilibrium", *Games and Economic Behavior*, 34, 177-99.
- Anderson, S. P., J. K. Goeree, and C. A. Holt, 2002, "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies", *Southern Economic Journal*, 69, 21-47.
- Berg, J., L. Daley, J. Dickhaut, and J. O'Brien, 1992, "Moral Hazard and Risk Sharing", in: Isaac, M., *Research in Experimental Economics*, Vol. 5, 1-34.
- Bolton, G. E., and Ockenfels, A., 2000, "A Theory of Equity, Reciprocity and Competition", *American Economic Review*, 100, 166-193.

- Camerer, C. F., Teck-Hua Ho, and Juin Kuan Chong, 2004, "Behavioral Game Theory: Thinking, Learning, and Teaching", in: *Advances in Understanding Strategic Behavior, Essays in Honor of Werner Güth*, ed. Steffen Huck, London.
- Capra, M. C., J. K. Goeree, R. Gomez, and C. A. Holt, 1999, "Anomalous Behavior in a Traveler's Dilemma?", *American Economic Review*, 89, 678–90.
- Charness, G., and M. Rabin, 2002, "Understanding Social Preferences with Simple Tests", *Quarterly Journal of Economics*, 117, 817-869.
- Chernomaz, K., 2011, "A Model of Moral Hazard with Inequity Aversion: An Experimental Test", Discussion Paper.
- Cochard, F., and M. Willinger, 2005, "Fair Offers in a Repeated Principal-Agent Relationship with Hidden Actions", *Economica*, 72, 225-240.
- Epstein, S., 1992, "Testing Principal-Agent-Theory: Experimental Evidence", in: Isaac, M., *Research in Experimental Economics*, Vol. 5, 35-60.
- Fehr, E., and S. Gächter, 2001, "Do Incentive Contracts Crowd Out Voluntary Cooperation?", Discussion Paper.
- Fehr, E., A. Klein, and K. M. Schmidt, 2004, "Contracts, Fairness, and Incentives", CESifo Working Paper (1215).
- Fehr, E., and K. M. Schmidt, 1999, "A Theory of Fairness, Competition and Co-operation", *Quarterly Journal of Economics*, 114, 817-868.
- Fischbacher, U., 2007, "z-Tree: Zurich Toolbox for Ready-made Economic Experiments", *Experimental Economics*, 10, 171-178.
- van Dijk, F., J. Sonnemans, and F. van Winden 2000, "Incentive Systems in a Real Effort Experiment", CESifo Working Paper (272).
- Goeree, J. K., and C. A. Holt, 2000, "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games", *European Economic Review*, 44, 1079–89.
- Goeree, J. K., and C. A. Holt, 2001, "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions", *American Economic Review*, 91, 1402–22.
- Goeree, J. K., C. A. Holt, and T. R. Palfrey, 2005, "Regular Quantal Response Equilibrium", *Experimental Economics*, 8, 347–67.

- Goeree, J. K., C. A. Holt, and T. R. Palfrey, 2002, "Quantal Response Equilibrium and Overbidding in Private-Value Auctions", *Journal of Economic Theory*, 104, 247–72.
- 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.
- Güth, W., M. Königstein, J. Kovács, and E. Zala-Mező, 2001, "Fairness Within Firms: The Case of One Principal and Multiple Agents", *Schmalenbach Business Review*, 53, 82-101.
- Holmstrom, B., 1979, "Moral Hazard and Observability", *Bell Journal of Economics*, 10, 74-91.
- Holmstrom, B., and P. Milgrom, 1987, "Aggregation and Linearity in the Provision of Intertemporal Incentives", *Econometrica* 55, 231-259.
- Keser, C., and M. Willinger, 2000, "Principals' Principles and Hidden Actions: An Experimental Investigation", *International Journal of Industrial Organization*, 18, 163-185.
- Keser, C., and M. Willinger, 2007, "Theories of Behavior in Principal-Agent Relationships with Hidden Action", *European Economic Review*, 51, 1514-1533
- Luce, R. D. 1958, *Individual Choice Behavior*, New York: Wiley.
- McKelvey, R. D., and T. R. Palfrey, 1995, "Quantal Response Equilibria for Normal Form Games", *Games and Economic Behavior*, 10, 6–38.
- McKelvey, R. D., and T. R. Palfrey, 1998, "Quantal Response Equilibria for Extensive Form Games", *Experimental Economics*, 1, 9–41.
- Ross, S. A., 1973, "The Economic Theory of Agency: The Principal's Problem", *American Economic Review*, 63, 134-139.
- Shavell, S., 1979, "Risk Sharing and Incentives in the Principal and Agent Relationship", *Bell Journal of Economics*, 10, 55-73.
- Spremann, K., 1987, "Principal and Agent", in: Bamberg, G., and K. Spremann, *Agency Theory, Informations, and Incentives*, Berlin, 3-37.

Appendix A: Data tables

Table A1: Means and medians of effort deviations Δe

Means						
	Period					
Treatment	1	2	3	4	5	6
Baseline	2.8	0.1	0.7	1.2	0.5	0.6
One Auction	5.7	0.8	4.0	4.8	1.7	2.5
Baseline One Shot	2.4	3.4	2.8	0.4	0.6	-1.2
Real Effort	4.9	1.4	1.0	-0.5	-0.1	-0.2
Real Effort No Framing	8.4	3.4	1.4	0.4	6.0	-0.9
Repeated Auctions	5.2	3.1	3.5	2.6	-0.8	0.9
Medians						
	Period					
Treatment	1	2	3	4	5	6
Baseline	2	0	0	0	0	0
One Auction	5.6	0	1.5	1	0	0
Baseline One Shot	1	.5	1	0	0	0
Real Effort	2.5	0	0	0	0	0
Real Effort No Framing	4	0	1	0	1	0
Repeated Auctions	7.5	1	1	1.7	0	0