



Theories of behavior in principal–agent relationships with hidden action

Claudia Keser^a, Marc Willinger^{b,*}

^aIBM T. J. Watson Research Center and CIRANO, Montreal, France

^bBETA, Université Louis Pasteur, Strasbourg and Institut Universitaire de France, France

Received 16 November 2003; accepted 2 October 2006

Available online 17 January 2007

Abstract

In a laboratory experiment, we investigate behavior in a principal–agent situation with moral hazard. We evaluate the predictive success of two theories. One is the standard agency theory, which assumes that the agent will accept any contract offer that satisfies his participation constraint, typically requiring zero expected utility. The other is the “fair-offer” theory suggested by Keser and Willinger [2000. Principals’ principles when agents’ actions are hidden. *International Journal of Industrial Organization* 18 (1), 163–185], which requires that the principal provide full insurance against losses to the agent and leave him a share of at most 50% of the generated surplus. The treatment variable of our experiment is the cost of effort. As effort costs increase, expected net surplus of a contract decreases. We observe that fair-offer theory generally predicts observed contract offers better than standard agency theory. However, the predictive success of the fair-offer theory decreases, while the one of standard agency theory increases with decreasing expected net surplus.

© 2006 Published by Elsevier B.V.

JEL classification: C91; D82

Keywords: Principal–agent theory; Experimental economics

1. Introduction

The assumption of information asymmetry has become standard in economic theory. In particular, the design of optimal incentive schemes under moral hazard has led to fruitful

*Corresponding author.

E-mail address: willinger@lameta.univ-montp1.fr (M. Willinger).

developments in many areas, including labor economics, financial economics, public regulation, and organizational design. Nonetheless, only modest attempts have been made to test the predictive validity of the principal–agent model (see Chiappori and Salanié, 2003). A major reason is the inaccessibility of data, because firms and organizations owning such data are usually reluctant to make them available to researchers. Another reason is that real world contracts incorporate many characteristics that are not taken into account by the theory. Therefore, many factors can account for differences between observed contracts and contracts predicted by agency theory.

Laboratory experiments can overcome these obstacles by generating the particular data sets that are required for testing the main predictions of principal–agent models. Recently, several attempts have been made in this direction, leading to interesting insights about the behavior of agents and principals (Berg et al., 1992; Epstein, 1992; Anderhub et al., 2002; Güth et al., 1998; Keser and Willinger, 2000).

In this paper, we report findings based on new experimental data, gathered in order to test the predictive validity of the standard principal–agent model with moral hazard. The experiment, which is based on a design introduced in Keser and Willinger (2000), allows us to test whether observed contracts satisfy the basic assumptions of agency theory: the participation constraint and the incentive-compatibility constraint. We compare the predictive validity of these constraints to the predictive validity of behavioral principles identified in Keser and Willinger (2000). The fair-offer theory described by these principles is compatible with social preference theories.

In the experiment, a participant in the role of a principal is randomly matched with a participant in the role of an agent. They have the opportunity to conclude a contract. If the agent accepts the contract offered by the principal, he has to choose between providing either low or high effort. High effort is more costly to the agent than low effort. Each effort level generates a stochastic gain that accrues to the principal. There are two possible gains, a high gain and a low one. The high gain is more likely if the agent chooses high effort. If the agent chooses low effort, the two gains are equally likely. The principal cannot observe the agent's choice. Thus, the principal, who has to pay the agent for his performance, can make the payment dependent on the realized gain but not on the effort chosen by the agent.

The procedure of the interaction is such that the principal makes a contract offer to the agent that specifies a payment scheme. The agent can either accept the payment scheme offered by the principal and choose an effort level or reject the contract. In the latter case, the interaction between the principal and the agent immediately ends with zero profit for each party.

Assuming risk neutrality for both the principal and the agent, the game is solved by backward induction. We study a parametric version of the game, for which the subgame perfect equilibria are characterized by contract offers that induce high effort. For a particular set of parameters, we found in Keser and Willinger (2000) that most of the observed contract offers yield in both states, low and high gain, higher payments to the agent than predicted by the subgame perfect equilibrium solution. Furthermore, half of the observed payment schemes violate the incentive compatibility constraint, which should induce the agent to choose high effort. Agents tend to react as predicted by expected profit maximization. We showed that most of the observed contract offers satisfy the following three principles:

Appropriateness: The payment to the agent is larger if the high gain is realized than if the low gain is realized.

- Loss avoidance:* The payment to the agent in each of the two states covers effort costs.
Sharing power: The principal's profit is at least equal to 50% of the net surplus of the contract.

The combination of these three principles defines a subset in the contract space, called the *fair-offer area*. In Keser and Willinger (2000) we observed that more than 90% of contract offers belong to this relatively small subset. Thus, the fair-offer subset provides a good description of the experimental data.

While two of the principles defining the fair-offer area, appropriateness and sharing power, are not in conflict with standard agency theory, loss avoidance is clearly incompatible. According to agency theory, the principal can always implement low effort by offering the agent a risk-free contract in which the payment is at least equal to the cost of low effort. Thus, a profit-maximizing principal who wants to implement low effort offers a flat wage equal to the cost of low effort. A profit-maximizing principal who wants to implement high effort must offer an incentive compatible contract such that the agent incurs a net loss in the bad state and a net gain in the good state. Because the principal makes the participation constraint binding and thus keeps the agent at his reservation level, the entire expected net surplus of the contract goes to the principal.

These basic requirements of agency theory are almost always violated in the experiment by Keser and Willinger (2000). The observed contracts induce surplus sharing between the principal and the agent in both states; it is rare that the agent risks incurring a loss. A plausible reason for observing such strong differences to the predictions of agency theory is that in the experiment in Keser and Willinger (2000) the expected net surplus of a contract is quite large with respect to the cost of effort. In other words, there is a huge difference between the effort costs and the expected gain for each effort level. This might encourage principals to make generous contract offers. Furthermore, a principal who fears rejection of an unfair contract might be induced to make more generous offers.

The ultimatum component of the contracting procedure in our experiment allows us to compare our results to those from other experiments, such as experiments in ultimatum bargaining, trust, and gift-exchange games. Fairness and reciprocity theories (Bolton and Ockenfels, 2000; Fehr and Schmidt, 1999), which try to capture the general idea that participants' behavior is partly driven by social preferences, provide a comprehensive framework to organize the data collected in those experiments.

Fairness and reciprocity theories can also be applied to our setting. Our theory is that the size of the net expected surplus might affect its division between the principal and the agent. More specifically, in the case of a large net expected surplus, the principal is likely to leave a considerable share of the surplus to the agent, thereby attempting to induce high effort due to reciprocity. Alternatively, in the case of a small net expected surplus, the principal must compete with the agent to secure some profit for himself.

Standard agency theory, or the game-theoretical solution, predicts that the participation constraint is binding, regardless of the importance of the surplus. This implies that the agent only achieves his reservation payoff, whatever the size of the surplus. In other words, the game-theoretical solution ignores the distributive aspect inherent to the agency problem, which is captured by models of social preferences.

In contrast to the game-theoretical solution, the fair-offer theory predicts that the agent receives a higher payment when the net expected surplus of the contract (or the "pie") to be divided between the principal and the agent is larger. Reformulating our theory, the

smaller the net expected surplus, the better the predictive success of the game-theoretical solution.

Our experimental design generates four different levels of expected surplus by varying the agent's effort costs, from "low" to "very high." We compare the fair-offer prediction to the standard agency prediction involving (1) both a risk-neutral principal and a risk-neutral agent and (2) a risk-neutral principal but a risk-averse agent. For the latter prediction we assume a strictly increasing concave utility function for the agent. All predictions describe a subset of the contract space, which allows us to assess their respective predictive success with respect to the contract offers observed in the experiment.

We find that the fair-offer theory is a better predictor for the observed contracts than standard agency theory. However, with increasing effort costs—a shrinking expected surplus—the predictive success of the fair-offer theory decreases and the predictive success of agency theory under the assumption of a risk-averse agent increases. We show that this result can be explained by the conflict between two objectives that the principal simultaneously tries to satisfy: ensuring loss avoidance to the agent and selfish profit maximization.

Section 2 presents the experimental design, and Section 3 summarizes the theoretical predictions. Results are presented in Section 4 and discussed in Section 5. Section 6 provides the conclusion.

2. Experimental design

The experiment was run at two different sites, the University Louis Pasteur in Strasbourg (*France* thereafter), and at the University of Karlsruhe (*Germany* thereafter). At both sites observations were collected under the same procedure. Participants were randomly selected from the existing local subject pool (about 800 subjects in France and 1500 subjects in Germany). We organized eight sessions in France and six sessions in Germany. Each session involved 16 participants, eight principals and eight agents, divided into two independent player groups of four principals and four agents who interacted with each other, matched in pairs. A session consisted of 10 periods. At the beginning of each period, each of the four principals was randomly matched with one of the four agents of his group. In each group we observed 40 contracts, which altogether represent an independent observation.

Four different treatments, corresponding to situations from low to very high costs were implemented (presented in Table 1 below). Except for the low-cost treatment, we collected 160 contracts, representing four independent observations per treatment and per country. The data for the low-cost treatment had been available for the German site from a

Table 1
Experimental design

Effort level	Probability of gaining			Agent's activity costs			
	50	100		Low	Medium	High	Very high
Low (<i>L</i>)	50%	50%	C_L	13	27	34	41
High (<i>H</i>)	20%	80%	C_H	20	34	41	48

previous experiment (Keser and Willinger, 2000).¹ We collected four additional observations for this treatment at the French site.

A contract offer has two components: the payment to the agent in the case when the “bad state” occurs (i.e., a low gain of 50 for the principal) and the payment to the agent in the case when the “good state” occurs (i.e., a high gain of 100 for the principal). Gains, contract payments, and activity costs were all expressed in points.

In each period, each principal had to make a contract offer to the agent. As soon as all principals had made their offers, the offers were collected by the server of the computer network and sent in a random order to the agents in the corresponding group. Each agent, after receiving the contract offer, had to decide whether to accept or reject it. If he rejected the offer, both the principal and the agent received a zero payoff. If the agent accepted the contract offer, he had to choose between low effort (L) and high effort (H). The choice of L implied a 50-percent chance for each state, while the choice of H implied a 20-percent chance for the bad state and an 80-percent chance for the good state (see Table 1).

Points were accumulated on each participant’s account and were on permanent display on his computer screen. After each period, each participant received summary data on the proposed contract, the realized gain, the agent’s acceptance decision, and the payment transferred to the agent in each of the previous periods. Note, however, that in case of acceptance the principal was never informed about the agent’s activity choice, which remained hidden information.

Table 1 summarizes the parameters that we used for the different treatments of the experiment. Each of the four treatments corresponds to a pair (C_L, C_H) , where C_L denotes the cost of low effort and C_H the cost of high effort. With the more costly effort there is a larger probability of obtaining a large gain. Note that the cost difference between low effort and high effort remains constant at the level of seven units, across treatments. In the following sections, we identify treatments by the corresponding cost pair $(C_L - C_H)$.

3. Theoretical predictions

In this section, we provide a formal statement of the three predictions that we test on our data. Two of our predictions rely on subgame perfect equilibrium solutions, assuming selfish agents who care only about their egoistic payoff. The principal and the agent play a sequential game, in which the principal offers a contract that can be accepted or rejected by the agent. Conditionally on acceptance, the agent chooses an effort level that produces a random outcome for the principal.

If both players are expected payoff maximizers, implying that they are risk neutral, the subgame perfect equilibrium solution of this game predicts an indifference subset of contracts implementing high effort. We call this the *equilibrium under risk neutrality*.

Our second prediction, *equilibrium under risk aversion*, corresponds to the standard model of principal–agent theory, which assumes that the principal is risk neutral while the agent is strictly risk averse. Under the additional assumption that the principal knows the agent’s utility function, a unique high effort implementing contract can be defined. In the experiment, however, this is an unrealistic assumption. To account for this type of

¹Those data were generated in five sessions involving 20 participants who were divided into two independent player groups of five principals and five agents each. Thus, they contain 500 contracts, representing 10 independent observations.

incomplete information, and in order to derive a more relevant behavioral prediction, we derive the set of all possible equilibrium contract offers for the family of strictly increasing and strictly concave utility functions. Under this assumption we are able to predict a subset of potential equilibrium contracts within the space of admissible contracts. One interpretation is that the principal offers a contract based on his belief about the agent's utility function. Because the experimenter is unable to observe the principal's beliefs, we allow for a whole set of possible beliefs underlying a contract offer. Another possible interpretation is that the principal is uncertain about the agent's utility function. We model the principal's uncertainty by assuming a uniform distribution over the set of all strictly increasing and strictly concave utility functions.

Note that by allowing any strictly concave utility function, we give the best possible chances for the standard agency model to be a good predictor of our data. We simply require that the observed contract offers lie in the set of contracts predicted by the equilibrium under risk aversion.

The third prediction that we test in our experiment is the *fair-offer hypothesis* proposed in Keser and Willinger (2000). Similarly to the two subgame perfect equilibrium predictions, the fair-offer hypothesis predicts a subset of all admissible contracts. We shall thus compare the predictions on the basis of the measure of predictive success proposed by Selten and Krischker (1983). In the following subsections, we give a formal statement of each of the three predictions.

3.1. Equilibrium under risk neutrality

We model the interaction between the principal and the agent as a four-stage game. In the first stage, the principal makes a contract offer (w_{50}, w_{100}) to the agent, which specifies a payment scheme contingent on the realized gain of either 50 or 100. In the second stage, the agent decides whether to accept or reject the contract offer. A rejection terminates the game immediately, and both players earn zero profits. If the agent accepts the contract, he has to choose between low effort, L , and high effort, H , in the third stage. In the final stage, the gain is randomly drawn, according to the probabilities induced by the effort chosen by the agent. In the case of acceptance of the contract, the principal's profit is $I - w_i$, with $i \in \{50, 100\}$, and the agent's profit is $w_i - C_j$, where $j \in \{L, H\}$.

Under the assumption of risk neutrality for both players, the game-theoretical solution implies that both players maximize their expected profits. The game is solved by backward induction. The equilibrium contract offered by the principal implements high effort for any of the four cost conditions. Table 2 shows the equilibrium contracts when offers are

Table 2

Subgame perfect equilibrium contracts for a risk-neutral principal and a risk-neutral agent (the numbers in parenthesis correspond to the predicted contracts (w_{50}^*, w_{100}^*))

Effort costs (low–high)	Predicted contracts							
13–20	(0,25)							
27–34	(2,42)	(6,41)	(10,40)	(14,39)				
34–41	(1,51)	(5,50)	(9,49)	(13,48)	(17,47)	(21,46)		
41–48	(0,60)	(4,59)	(8,58)	(12,57)	(16,56)	(20,55)	(24,54)	(28,53)

restricted to be integer valued. Note that for cost situation (13–20) there is a unique integer-valued equilibrium contract.²

Prediction 1. Under the assumption of risk neutrality for both players, the subgame perfect equilibrium solutions of the game correspond to the payment schemes (w_{50}^*, w_{100}^*) shown in Table 2. For any of these contracts, the agent accepts the offer and chooses high effort.

All equilibrium contracts share the common property that the agent makes a net loss if the bad state occurs, regardless of the effort level chosen. This is a direct consequence of incentive compatibility because the agent is kept as close as possible to zero expected profit, his reservation payoff if he rejects the contract. As we shall see, most of our observed contracts do not satisfy this fundamental property of agency theory.

In the following, we derive Prediction 1 for the game without integer restrictions on the values of w_{50} and w_{100} . Using backward induction, we first determine the agent's best reply to any contract offer. Then, we take into account the agent's best reply function to identify the principal's expected profit-maximizing contract offers.

3.1.1. The agent's best reply function

The agent's best reply to a contract offer (w_{50}, w_{100}) is to *accept and choose high effort* if the participation and the incentive compatibility constraints for high effort are satisfied. The participation constraint (inequality 1) requires that the agent's expected profit with high effort is at least as high as his reservation payoff, which is assumed to be zero. The incentive compatibility constraint (inequality 2) requires that the expected profit with high effort is at least as high as the expected profit with low effort:

$$0.2(w_{50} - C_H) + 0.8(w_{100} - C_H) \geq 0 \Leftrightarrow w_{100} \geq -0.25w_{50} + (5/4)C_H, \quad (1)$$

$$\begin{aligned} 0.2(w_{50} - C_H) + 0.8(w_{100} - C_H) &\geq 0.5(w_{50} - C_L) + 0.5(w_{100} - C_L) \\ \Leftrightarrow w_{100} &\geq w_{50} + (10/3)(C_H - C_L). \end{aligned} \quad (2)$$

Similarly, the agent's best reply to a contract offer (w_{50}, w_{100}) is to *accept and choose low effort* if the participation constraint (inequality 3) and the incentive compatibility constraint (inequality 4) for low effort are satisfied:

$$0.5(w_{50} - C_L) + 0.5(w_{100} - C_L) \geq 0 \Leftrightarrow w_{100} \geq -w_{50} + 2C_L, \quad (3)$$

$$\begin{aligned} 0.5(w_{50} - C_L) + 0.5(w_{100} - C_L) &\geq 0.2(w_{50} - C_H) + 0.8(w_{100} - C_H) \\ \Leftrightarrow w_{100} &\leq w_{50} + (10/3)(C_H - C_L). \end{aligned} \quad (4)$$

If neither of the participation constraints, (1) or (3), is satisfied, the agent's best reply is to refuse the contract offer.

3.1.2. The principal's calculus

The principal takes the agent's best reply into account when making a contract offer. Let us define the principal's expected profit if the agent chooses high effort by

²In Keser and Willinger (2000) we required a strictly positive expected profit for the agent (participation constraint). As we give up this requirement here, the equilibrium contract is slightly different from the one in Keser and Willinger.

$\Pi_H(w_{50}, w_{100}) = 0.2(50-w_{50}) + 0.8(100-w_{100})$. Similarly, let $\Pi_L(w_{50}, w_{100}) = 0.5(50-w_{50}) + 0.5(100-w_{100})$ be the principal's profit if the agent chooses low effort. The principal maximizes his profit by extracting the maximum surplus from the agent, which means that he makes his contract offer such that the participation constraint is binding. If the agent chooses high effort, the principal maximizes his expected profit by offering one of the contracts that satisfies $w_{100} = -0.25w_{50} + (5/4)C_H$. His maximum expected profit is then given by $\Pi_H^*(\cdot) = 90 - C_H$. Similarly, if the agent chooses low effort, the principal maximizes his expected profit by offering one of the contracts that satisfies $w_{100} = -w_{50} + 2C_L$. His maximum expected profit is then given by $\Pi_L^*(\cdot) = 75 - C_L$. The principal implements high effort if $\Pi_H^*(\cdot) > \Pi_L^*(\cdot)$. This condition is always satisfied with the parameters of our experiment because $C_H - C_L = 7$.

It follows that the subgame perfect equilibrium solution involves the principal inducing the agent to choose high effort. The contract offers (w_{50}^*, w_{100}^*) satisfy the incentive constraint for high effort and lie on the participation constraint. In the game without integer restrictions, there are an infinite number of subgame perfect equilibrium contracts, and the principal might thus implement any one of them. Because in the experiment participants' choices were constrained to be integer numbers, we shall restrict our attention to the equilibrium contracts with integer values as summarized in Table 2.

The multiplicity of equilibria in the risk-neutral case comes from the fact that the agent's participation constraint has the same slope as the principal's iso-expected-profit lines in the (w_{50}, w_{100}) space. With the restriction to integer numbers, the number of equilibrium contracts is increasing with the cost level. Note that in the case where the agent is risk neutral, the equilibrium contracts are also Pareto-optimal contracts. The non-observability of the agent's effort affects only risk sharing but not the expected profits of the two players.

3.2. Equilibrium under risk aversion

The analysis of the game for a risk-averse agent is similar to the one presented above, except that the agent's expected profit is replaced by his expected utility for the payoffs. We assume throughout that the agent's utility function, $u(x)$, satisfies $u'(x) > 0$ and $u''(x) < 0$, for all x . If the principal wants to implement high effort, his contract offer must satisfy the participation and the incentive compatibility constraints:

$$0.2u(w_{50} - C_H) + 0.8u(w_{100} - C_H) \geq u(0), \quad (5)$$

$$0.2u(w_{50} - C_H) + 0.8u(w_{100} - C_H) \geq 0.5u(w_{50} - C_L) + 0.5u(w_{100} - C_L). \quad (6)$$

In contrast to the risk-neutral case, the equilibrium contract is not necessarily socially optimal when the agent is risk averse. More specifically, the required compensation scheme to implement high effort under non-observability incurs a larger expected wage payment than under observability of the agent's effort. This may cause a welfare loss if the principal is better off by offering the least costly contract that induces low effort.

Note that in contrast to most agency models, we do not assume that the utility of the wage payment and the disutility of effort are generated by different variables. Our approach seems reasonable in the context of our experiment, because payment and effort costs are measured in the same experimental units (points). We can therefore take the net profit (wage payment minus effort costs) as the argument of the utility function. Implicitly,

we assume that participants are able to aggregate the wage payment and the cost of effort to evaluate the net contingent profit of the contract. This assumption is theoretically justified and seems to be empirically supported by the contract offers observed in Keser and Willinger (2000). Because this assumption implies non-separability of the utility of the payment and the disutility of effort, it can be optimal for the principal, assuming that the agent is strictly risk averse, to offer a contract that fully covers the effort costs.³

Prediction 2. If the agent is strictly risk averse, i.e. $u'(\cdot) > 0$ and $u''(\cdot) < 0$, and the principal is risk-neutral, the set of contracts that implement high effort satisfy restrictions (i)–(iii) (see Appendix):

- (i) $w_{100} \leq -\frac{1}{4}w_{50} + \frac{C_L+15}{0.8}$
- (ii) $w_{100} > -\frac{1}{4}w_{50} + \frac{C_H}{0.8}$
- (iii) $w_{100} > w_{50}$

The first of these conditions states that the principal implements high effort only if he expects a larger profit than if he implemented low effort. The second condition states that the contract must satisfy the participation constraint, which implies that the contract always lies above the tangency line to the reservation indifference curve. The tangency line corresponds to the boundary case of linear (risk-neutral) utility. The third inequality follows from the monotone likelihood property: the principal offers a larger payment to the agent in the case of a high gain because the likelihood of a high gain is larger for the more costly activity. Note that if the third inequality was not satisfied, the agent would prefer to choose low effort (assuming that the participation constraint is satisfied) because this would be a stochastically dominating choice.⁴

Taken all together, conditions (i)–(iii) define an area in the space of contracts that we identify as the *equilibrium under risk aversion*.

3.3. Fair-offer prediction

In the earlier experiment presented in Keser and Willinger (2000), we found that the observed contracts for cost level (13–20) were not correctly predicted by subgame perfect equilibria, neither under the assumption of a risk-neutral agent nor under the assumption of a risk-averse agent.⁵ We showed instead that most of the observed contracts belong to a subset of contracts that satisfy the three principles outlined in the introduction: appropriateness, loss avoidance, and sharing power. Appropriateness means that the agent's payment is increasing with the principal's gain. This principle is also satisfied by the standard agency prediction, when the two effort levels satisfy the monotone likelihood

³The slope of the incentive compatibility curve for implementing activity H is given by

$$\frac{dw_{100}}{dw_{50}} = \frac{0.5u'(w_{50} - C_L) - 0.2u'(w_{50} - C_H)}{0.8u'(w_{100} - C_H) - 0.5u'(w_{100} - C_L)}$$

The sign of this expression can be positive or negative since $u'(w_{50} - C_L) < u'(w_{50} - C_H)$ by concavity of $u(\cdot)$. Recall that under the assumption of separability, dw_{100}/dw_{50} is always positive.

⁴Conditions (i)–(iii) are necessary conditions.

⁵The prediction for risk-averse agents in this earlier study was restricted to the class of utility functions with constant absolute risk aversion. In contrast to this, in the present analysis we consider the larger class of strictly increasing and strictly concave utility functions.

property. Loss avoidance, however, which means that contract offers provide the agent full insurance against losses, contradicts the standard agency prediction. Sharing power states that the principal earns at least half of the net gain from the contract.

There are several alternative ways to define the principles of loss avoidance and sharing power, depending on which cost is taken into account: C_L , C_H , or a combination of the two. For example, loss avoidance can be defined as giving at least the cost of low effort for w_{50} , and at least the cost of high effort for w_{100} (condition 2c). In total, nine different combinations of these principles are possible. Each of these combinations corresponds to a relatively small subset of the contract space, which we shall call (a variant of) the *fair-offer prediction*. The three underlying principles, with their variants, are formally defined as follows:

- | | |
|-----------------------------|--|
| (1) <i>Appropriateness:</i> | $w_{50} \leq w_{100}$ |
| (2) <i>Loss avoidance:</i> | (2a) $w_{50} \geq C_L$ and $w_{100} \geq C_L$ |
| | (2b) $w_{50} \geq C_H$ and $w_{100} \geq C_H$ |
| | (2c) $w_{50} \geq C_L$ and $w_{100} \geq C_H$ |
| (3) <i>Sharing power:</i> | (3a) $w_{50} \leq C_H + (50 - C_H)/2$ and $w_{100} \leq C_H + (100 - C_H)/2$ |
| | (3b) $w_{50} \leq C_L + (50 - C_L)/2$ and $w_{100} \leq C_L + (100 - C_L)/2$ |
| | (3c) $w_{50} \leq C_L + (50 - C_L)/2$ and $w_{100} \leq C_H + (100 - C_H)/2$ |

Prediction 3. Observed contract offers tend to belong to one of the areas of the fair-offer prediction.

Although Prediction 3 cannot be quantified, we rely on our previous findings (Keser and Willinger, 2000) showing a predictive success around 80%, which corresponds to a hit rate of 96%. In contrast to Predictions 1 and 2, Prediction 3 is not derived from a formal theory. Rather, we take it as a descriptive theory that might organize the data meaningfully. The principles, however, can be derived from a formal model based on social preferences under risk. Ideally, such a model should predict how the principal divides a contingent surplus and shares the risk of the contract, and how this affects the agent's choice of effort. Englmaier and Wambach (2005) made a first step into this direction, by incorporating inequity aversion in the agent's utility function. Inequity aversion alters the structure of optimal contracts, leading to more equitable surplus distribution. However, their model does not explicitly capture our loss aversion principle.

4. Results

For analysis of the contract offers we rely on Selten's *measure of predictive success* (Selten and Krischker, 1983; Selten, 1991). The predictive success, S , of a theory is measured by the difference $S = h - a$, where h measures the *hit rate* and a , the *area*. In our experiment, the hit rate is defined as the percentage of contract offers that fall into the predicted area. The area corresponds to the percentage of points in the contract space that belong to the predicted area. Note that the area is a measure of parsimony of a theory. More parsimonious theories predict smaller areas. The most permissive theory predicts any possible contract in the contract space and its measure of predictive success is zero. Each of the three predictions discussed in Section 3 corresponds to a specific area in the contract space.

All results and tests presented here are based on the pooled data sets collected from the two sites (France and Germany), for two reasons. First, our main interest is with respect to the treatment effect. Second, none of the statistically significant subject-pool differences appears relevant for the analysis of the treatment effect.

All tests are two-sided. Unless stated otherwise, we require significance at the 10-percent level. We use the following abbreviations: KW for the Kruskal–Wallis test, MW for the Mann–Whitney U test, and W for the Wilcoxon signed rank test.

4.1. Equilibrium under risk neutrality

To examine how well the equilibrium under risk neutrality predicts our data, we distinguish between compatible and non-compatible offers. *Compatible offers* are contract offers that are compatible with the risk-neutral prediction in that they satisfy both the incentive constraint and the participation constraint for the agent to choose high effort. We also measure how close the observed contract offers are to the prediction by taking Euclidian distances to the equilibrium solution.

4.1.1. Compatible offers

Table 3 shows, for each of the treatments, the relative frequency of compatible offers, the relative frequency of compatible offers that were accepted, and the relative frequency with which agents chose high effort after acceptance of a compatible offer. Overall, the average frequency of compatible offers is 40% and significantly below 50% (binomial test, 1% significance). When principals made a compatible offer, agents tended to accept and provide high effort. This result is in keeping with Keser and Willinger (2000).

A KW test shows no significant difference in the relative frequency of compatible offers among the four treatments. Neither does any of the pairwise comparisons (MW tests).

In each of the treatments, agents significantly tend to accept compatible offers (binomial tests, 1% significance). We observe a significant difference in the relative frequency of acceptance of compatible offers among the four treatments (KW test, 1% significance). Table 3 indicates that, on average, compatible offers are most frequently accepted in the (13–20) treatment and least frequently accepted in the (41–48) treatment. Pairwise comparisons reveal that compatible offers are more frequently accepted in the (13–20) than in the (27–34) and (41–48) treatment, and more frequently accepted in the (34–41) than in the (41–48) treatment (MW tests, 5% significance).

A possible reason for observing fewer acceptances when effort costs are very high could be lower shares of expected surplus offered to the agent. As we shall show in Section 5

Table 3

Relative frequency of compatible offers, accepted compatible offers, and choice of high effort in the case of a compatible offer

Treatment	Compatible offers	Accepted	High effort
13–20	0.44	0.94	0.64
27–34	0.37	0.85	0.56
34–41	0.38	0.88	0.61
41–48	0.37	0.81	0.55

(Table 7), though, the share of the principal's expected surplus does not significantly vary across treatments.

Overall, agents are more likely to choose high rather than low effort (W test, 1% significance), which is in keeping with the subgame perfect equilibrium solution. Although the relative frequency of accepted compatible offers that lead to high effort is not statistically different across treatments (KW and pairwise MW tests), we observe (in contrast to the overall result) in treatment (41–48) equally many groups choosing high effort with a relative frequency of at least 50% as we observe groups choosing low effort with a relative frequency of less than 50%.

We conclude that principals tend to propose contracts that deviate from the subgame perfect equilibrium. Agents, on the contrary, tend to play best reply to compatible offers. This tendency becomes weaker, though, when effort costs are very high. In the case of a non-compatible offer, which typically violates the incentive constraint, agents show no significant tendency for either effort.

4.1.2. Euclidian distances

Over all treatments, only two of the 1620 observed contract offers correspond exactly to one of the subgame perfect equilibria with a risk-neutral agent. Therefore, the corresponding measures of predictive success are all negative. However, the measure of predictive success might be too stringent, because it does not take into account the fact that contract offers might be *close* to the predicted contracts. One might assume that participants are prone to errors and that superficial evaluation of the situation could lead them to offer contracts that differ but nevertheless remain relatively close to the subgame perfect equilibrium. Furthermore, the distance from predicted contracts might vary from one treatment to another. In order to account distance from the subgame perfect equilibrium, we calculate for each cost level the average Euclidian distance between the observed contract and the *closest* predicted contract.

The average Euclidian distances for each treatment are summarized in Table 4, which also reports the respective average contract offers and the closest equilibrium contract to those averages. We observe that, as the cost of effort becomes larger and thus the net expected surplus smaller, the contract offers move “closer” to the subgame perfect equilibrium under risk neutrality, according to the Euclidian distance measure. The Euclidian distance significantly varies among the four treatments (KW test, 0% significance). Also all pairwise differences are significant at 5% (MW tests), except for the one between the two intermediate cost levels (27–34) and (34–41).

Table 4

Average contract offers and closest equilibrium offer under risk neutrality as measured by average Euclidian distance

Effort costs	Average contract offer		Average euclidean distance	Closest equilibrium contract	
	w_{50}	w_{100}		w_{50}^*	w_{100}^*
(13–20)	24.10	44.75	31.32	0	25
(27–34)	30.83	51.04	23.34	14	39
(34–41)	35.38	54.36	22.24	21	46
(41–48)	37.37	60.45	17.61	28	53

We conclude that as effort costs increase, principals tend to offer contracts that are closer to the subgame perfect equilibrium.

4.2. The equilibrium under risk aversion versus the fair-offer prediction

In this subsection, we compare the predictive success measures of the equilibrium under risk aversion and the fair-offer theory. For this comparison we take into account all contract offers, whether or not they are accepted, since our aim is to evaluate the predictive value of principal–agent theory with respect to contract offers. Each of the theories predicts a specific area in the contract space. Recall that in Keser and Willinger (2000) we found that contract offers for treatment (13–20) were more accurately predicted by the fair-offer hypothesis than by the subgame perfect equilibrium solution with either a risk-neutral or a risk-averse agent.

The fair-offer hypothesis combines the three principles: appropriateness, loss avoidance and sharing power. We compute the predictive success of all possible combinations of principles (fair-offer sets). Similarly to Keser and Willinger (2000), two of these combinations give clearly better results than all other possible combinations, combinations (1–2a–3a) and (1–2c–3a). The following analysis is based on the fair-offer subset (1–2c–3a), which gives a slightly better measure of predictive success (reported in the last column of Table 5). In this subset the agent receives at least the low effort cost in the bad state, at least the high effort cost in the good state, and less than half of the net surplus assuming high cost in both states.

Table 5 reports the average measures of predictive success separately for each principle of the selected combination. Appropriateness and sharing power have on average better measures of predictive success than loss avoidance. All measures are significantly positive (Binomial test, 5% significance), except for the loss avoidance principle in treatment (41–48), where two of the eight measures are negative (no significance), and in treatment (13–20), where one of the eight measures is negative (10% significance). As can be seen in Table 5, loss avoidance appears to be the weakest of the three principles in terms of predictive success.

There are no significant differences in the measures of predictive success for appropriateness and loss avoidance across treatments (KW tests). However, there is a significant difference for sharing power (KW test, 5% significance). Specifically, the predictive success for treatment (41–48) is significantly lower than for any of the other treatments (MW tests, 10% significance).

Table 5

Measures of predictive success^a for the fair-offer set (1–2c–3a), defined by $w_{100} \geq w_{50}$, $w_{50} \geq C_L$ and $w_{100} \geq C_H$, $w_{50} \leq (50 + C_H)/2$ and $w_{100} \leq (100 + C_H)/2$

	Appropriateness (1)	Loss avoidance (2c)	Sharing power (3a)	(1–2c–3a)
13–20	0.49	0.24	0.74	0.82
27–34	0.47	0.31	0.66	0.71
34–41	0.45	0.22	0.62	0.61
41–48	0.49	0.24	0.57	0.49

^aNote that the measure of predictive success cannot be larger than 1 but might be negative.

Table 6

Average measures of predictive success for the equilibrium under risk aversion and the fair-offer prediction

Effort costs	Risk aversion	Fair offer
13–20	0.05	0.82
27–34	0.15	0.71
34–41	0.18	0.61
41–48	0.26	0.49

Table 6 reports the average measures of predictive success of the fair-offer set and the risk-aversion hypothesis. The predictive success of the risk-aversion hypothesis clearly increases with the level of effort costs, while the opposite tendency can be observed under the fair-offer hypothesis.

The predictive success for the risk-aversion hypothesis is significantly lower for the lowest cost level (13–20) than for any of the higher cost levels (MW tests, 5% significance). None of the other comparisons shows a significant difference.

Although statistically not significant, the risk aversion hypothesis shows the highest measure of predictive success in treatment (41–48). This measure is 0.26, which is significantly below the measure of 0.49 of the fair-offer prediction in this treatment (W test, 10% significance). Obviously, also in each of the other treatments the fair offer theory is more successful than risk aversion in predicting actual contract offers (W tests, 1% significance).

5. Discussion

The main result of our experimental investigation is that although fair-offer prediction is always more successful than standard agency theory in predicting observed contract offers, the difference in the predictive success between the two theories becomes smaller as effort costs increase. In other words, the standard agency theory becomes a better predictor with a decreasing net expected surplus, while the predictive success of the fair-offer theory deteriorates with a decreasing net expected surplus.⁶

There are several plausible explanations for this result: low incentives leading to anomalous behavior, principals requiring a minimum level of expected profit, or principals requiring larger shares as costs becomes larger.

A possible reason for observing behavioral differences between the very high cost treatment (41–48) and the other treatments might be the extremely low incentives in the very high cost treatment. Gneezy and Rustichini (2000) report lab and field experiments showing that low incentives can affect participants in a contradictory way because of a possible conflict between intrinsic motivation and financial reward. Furthermore, Camerer and Hogarth (1999) conclude from a vast survey of experiments that in some cases incentives improve performance, and in other cases they have no effect, or even worse, hurt performance.

In the (41–48) treatment the net surplus in the bad state is almost zero if the agent chooses high effort. It could be that under this extreme condition, principals tend more

⁶Similarly, the area rate of standard agency theory increases, while the area rate of the fair-offer theory decreases with a decreasing net expected surplus.

than in other treatments to offer contracts involving a loss for the agent in the bad state—a behavior that is compatible with standard theory. On the other hand, when the net surplus becomes larger under lower effort costs, principals might tend to behave more cooperatively, as observed in many other experiments, such as ultimatum bargaining experiments. They offer a substantial share of the net surplus to the agent—a behavior that is compatible with social preference theories.

Such behavioral differences cannot be due to differences in monetary incentives, though, because our experimental design keeps the financial incentives constant across treatments. For instance, for French participants we implemented the following incentive scheme: cash payment = FF 80 (*French Francs*) + 0.15*[number of points–*c*], where *c* is a constant whose value was announced only at the end of the session.⁷ With this payoff scheme, participants earned FF 80, to which 15 cents was added or subtracted for each extra point above or below the constant. For profit-maximizing participants the objective is simply to maximize the number of points. Because participants were randomly assigned to treatments, there is no reason to assume that the proportion of profit-maximizing participants differs from one treatment to another.

An alternative explanation of our main result is that principals require a minimum level of expected profit, independent of the effort costs. In other words, they disregard the effort cost of the agent. This would contradict the sharing power hypothesis that takes the effort costs into account in defining an upper threshold level for contract offers.

According to this explanation, the principals have a psychological threshold for the range of expected profits. This threshold typically differs from one principal to another. As the effort costs increase, more and more principals have to take a larger proportion of the expected surplus, in order to secure their threshold. By requiring a large share of the expected surplus, the offers get closer to the contracts predicted by subgame perfect equilibrium under risk aversion, leading also to a reduction of the average Euclidian distance to the closest equilibrium contract. However, this line of reasoning is not compatible with our data as shown below.

Let

$$S_j = \frac{\pi_j(50 - w_{50}) + (1 - \pi_j)(100 - w_{100})}{\pi_j 50 + (1 - \pi_j)100 - C_j}$$

be the principal's share of the net expected surplus for a choice of effort *j* by the agent, where π_L is the probability of the bad state if the agent chooses low effort, and π_H is the corresponding probability for high effort. Note that according to subgame perfection, S_j should be equal or very close to 100% for the equilibrium contract.

While principals take significantly less than 100% in each treatment, they take a larger share of the expected surplus than the agents, both with respect to low effort and with respect to high effort (Table 7). Table 7 also reveals that, as the effort costs increase, principals take on average a slightly larger proportion of the expected surplus. Pairwise comparisons show that in treatment (13–20) the principal's share is lower than in any other treatment. For none of the other comparisons can we reject the null hypothesis of equal

⁷An equivalent scheme was used for German participants. Each session involved two independent groups of participants. Half of the members of a group acted as principals and the other half as agents. For each group, *c* was set equal to the average number of points of the other group. Participants knew that the value of the constant was different for principal participants and for agent participants.

Table 7
Principals' average share of the net expected surplus (all contract offers)

Treatment	Low effort cost	High effort cost
13–20	0.655	0.704
27–34	0.710	0.768
34–41	0.735	0.805
41–48	0.767	0.813

Table 8
Number (percentage) of contract offers with $w_{50} < C_L$, and relative frequency of those accepted

Effort costs	Number of offers	Accepted
13–20	30 (5%)	0.50
27–34	58 (18%)	0.53
34–41	94 (29%)	0.61
41–48	138 (43%)	0.72

surplus shares (MW tests). Thus, our conclusion is that principals are more likely to require a constant share of the net expected surplus rather than a constant level of expected profit independent of effort costs.

A detailed analysis of the principal's state conditional surplus share⁸ shows that principals always require a larger conditional share for the bad state than for the good state. This provides an incentive for the agent to exert high effort because when the outcome is favorable, the principal rewards the agent by giving him a larger share of the net conditional surplus. Note that in contrast to principal–agent theory, the agent tends to receive a positive surplus share in each state (loss-avoidance principle) although, as shown in Table 8, the percentage of contract offers implying a loss for the agent in the bad state increases from 5% in treatment (13–20) to 43% in treatment (41–48).

Interestingly, as the effort costs increase, the principal requires a relatively larger conditional surplus share in the bad state with respect to the conditional share in the good state. This implies that for higher effort costs principals offer contracts that secure them both an overall larger expected surplus share (see Table 7) and a relatively larger conditional share in the bad state. In other words, this implies that the incentives for the agent become stronger when the effort costs increase, providing him with a larger relative share in the good state. For very high costs, such an incentive scheme can be achieved only by contract offers implying a loss for the agent in the bad state.

Note that the intersection of the fair-offer set and the area predicted under risk aversion is non-empty, and its relative size is about 1% of the contract space, whatever the treatment. The common area satisfies both the profit maximizing condition (for implementing high effort) and loss avoidance. We observe that as effort costs increase, more and more contracts fall into this common area, and satisfy the profit maximizing

⁸We define the principal's conditional surplus share for the bad state as $50 - w_{50}/50 - C_j$, where C_j is the cost of the effort chosen by the agent. Similarly, for the good state the principal's surplus share is $100 - w_{100}/100 - C_j$.

condition. At the same time we observe that overall fewer contracts satisfy loss avoidance. Our interpretation is that principals try to satisfy two apparently conflicting objectives: avoiding losses for the agent and trying to maximize their own profits.

With higher effort costs the second objective becomes dominant, and the frequency of contracts inducing a loss in the bad state increases (see Table 8). This tendency is particularly clear if we compare treatment (13–20) with treatment (41–48). The percentage of contract offers that induce a loss for the agent in the bad state is at 5% in treatment (13–20). In treatment (41–48) this percentage is as high as 43%.

Simultaneously, we observe that at higher cost levels agents are more likely to accept contracts that induce a loss in the bad state than at lower cost levels. This is consistent with the contract offers made by principals when the effort costs increase. When the expected net surplus is large, principals make more generous offers, and the very few contracts involving a loss are frequently rejected by the agents. Alternatively, as effort costs increase, the principals' offers become less generous, and simultaneously the contracts that involve a potential loss are more likely to be accepted by the agents.

It is as if the conflict between the profit maximizing objective and the loss avoidance objective would be solved clearly in favor of loss avoidance at low cost but less clearly at high cost, where at the same time profit maximization plays a more important role than at low cost. It also appears as if the principal and the agent tacitly agree on the implicit hierarchy of objectives with respect to the cost level. Nevertheless, as we showed above, the agent receives a conditional surplus share in the good state, and the relative size of this share increases with the cost level.

6. Conclusion

In the experiment reported in this paper we test a simple version of the principal–agent model with hidden action. The treatment variable is the cost of effort. According to the standard agency prediction, the principal designs the incentive compatible contracts in such a way as to appropriate the entire net expected surplus generated by the agent's effort. In other words, the agent receives only his reservation payoff, whatever the cost of effort. Our results do not support this prediction, in particular not when effort costs are low. In that case, a large net surplus is generated by the contractual relationship. Similar to experiments on, for example, ultimatum bargaining, we observe a more equitable sharing of the surplus—in contrast to what standard agency theory predicts. The incorporation of inequity aversion in the agent's utility function (Englmaier and Wambach, 2005) leads to predictions about surplus sharing that are compatible with our findings. However, when effort costs increase and the generated net surplus decreases until it becomes negligible, equity considerations lose some of their impact and leave more room for egoistic profit considerations. In such a situation, agency theory under the assumption of a risk-averse agent has a relatively high predictive success for actual human behavior—although its success measure is still below that of the fair-offer theory.

This raises an interesting question about social preference theories. According to our observations, social preference theories provide a good predictor for participants' behavior when the amount of resources to be divided is large enough. However, under resource scarcity and stronger competition for surplus appropriation, principals are more reluctant to cover the agents' risk and tend to take a larger share of the expected net surplus of the contract. Of course, this tentative conclusion needs to be made more precise.

More experimental research is required to identify clearly how the size of the expected surplus affects the principals' contract offers and the agents' acceptance rate of contracts involving losses. Furthermore, it would be of interest to develop a behavioral principal–agent model taking into account both inequity aversion and risk-sharing. Most models of social preferences are deterministic and therefore can neither explain how a random pie will be divided, nor how such a division might be affected by moral hazard. Although we need careful definitions of the concepts of inequity aversion in a stochastic environment and of “fair risk-sharing,” we conjecture that risk-sharing motivations might, to some extent, go against fairness considerations.

Acknowledgments

We would like to thank Kene Boun My for his support in running the experiment. Helpful discussions with François Cochard, John Dickhaut, Bentley MacLeod, Isabelle Maret, Phu Nguyen Van, Anne Rozan, François Salanié and Hubert Stahn are gratefully acknowledged. The experiment was run while Claudia Keser was working at the University of Karlsruhe. Financial support by the Sonderforschungsbereich 504 at the University of Mannheim is gratefully acknowledged.

Appendix

For simplification we use the following notations in this appendix: w_1 is the payment to the agent in the bad state, and w_2 is the payment to the agent in the good state.

Let $u(x)$ be the agent's utility function and assume that for all x , $u'(x) > 0$ and $u''(x) < 0$. We show that the contracts, for which the principal implements high effort, must satisfy the three restrictions (i)–(iii)⁹ below:

- (i) $w_2 \leq -\frac{1}{4}w_1 + \frac{C_L+15}{0.8}$
- (ii) $w_2 > -\frac{1}{4}w_1 + \frac{C_H}{0.8}$
- (iii) $w_2 > w_1$

Step 1: The principal can never implement high effort by offering a flat contract $w_2 = w_1$. For such a contract the agent maximizes his expected utility by choosing the least costly effort because $C_L < C_H$ and $u(w-C_L) > u(w-C_H)$. Therefore, the principal can implement low effort with the riskless contract (C_L, C_L) , which is the profit-maximizing contract for implementing low effort.

Step 2: Restriction (i) means that the principal implements high effort only if the expected profit from high effort is larger than the expected profit from the implementation

⁹More generally inequalities (ii) and (iii) are, respectively:

$$(ii) \quad w_2 > -\frac{\pi_H}{1-\pi_H}w_1 + \frac{C_H}{1-\pi_H}$$

$$(iii) \quad w_2 \leq -\frac{\pi_H}{1-\pi_H}w_1 + \frac{C_L + (\pi_L - \pi_H)(R_2 - R_1)}{1-\pi_H}$$

where π_L is the probability of state 1 if the agent chooses low effort, and π_H is the corresponding probability for high effort. R_i is the principal's profit in state i . The following inequalities are assumed: $R_1 < R_2$, $C_L < C_H$, and $\pi_H < \pi_L$.

of low effort. Because for low effort the principal maximizes his profit with the contract (C_L, C_L) , the following inequality holds for implementing high effort: $0.2(50 - w_1) + 0.8(100 - w_2) \geq 75 - C_L$. This is equivalent to inequality (i).

Step 3: In order to implement high effort, the principal must satisfy the agent's participation constraint: $0.2u(w_1 - C_H) + 0.8u(w_2 - C_H) \geq u(0)$. Without loss of generality, we assume that $u(0) = 0$. The slope of the participation constraint for high effort is given by

$$\frac{dw_2}{dw_1} = -\frac{1 u'(w_1 - C_H)}{4 u'(w_2 - C_H)}.$$

Since $u'(x) > 0$, the participation constraint curve is strictly decreasing and convex ($\frac{\partial}{\partial w_1} \frac{dw_2}{dw_1} > 0$), with slope $(-1/4)$ at the point $w_2 = w_1 = C_H$. For the participation constraint to be satisfied, contract offers must be such that $w_2 \geq -\frac{1}{4}w_1 + \frac{C_H}{0.8}$, where $w_2 = -\frac{1}{4}w_1 + \frac{C_H}{0.8}$ is the equation of the tangency curve to the participation constraint at the point (C_H, C_H) .

Step 4: We show that the incentive compatibility constraint for implementing high effort is never satisfied for contracts such that $w_1 > w_2$. To show this, assume that inequalities (A.1) and (A.2) below are simultaneously satisfied.

$$0.2u(w_1 - C_H) + 0.8u(w_2 - C_H) \geq 0.5u(w_1 - C_L) + 0.5u(w_2 - C_L), \quad (\text{A.1})$$

$$w_1 > w_2. \quad (\text{A.2})$$

We show that this assumption leads to a contradiction. Inequality (A.1) can be rewritten as

$$0.8u(w_2 - C_H) - 0.5u(w_2 - C_L) \geq 0.5u(w_1 - C_L) - 0.2u(w_1 - C_H). \quad (\text{A.3})$$

Some additional rewriting of (A.3) leads to

$$\begin{aligned} & -0.3(u(w_1 - C_L) - u(w_2 - C_H)) - 0.5(u(w_2 - C_L) - u(w_2 - C_H)) \\ & \geq 0.2(u(w_1 - C_L) - u(w_1 - C_H)). \end{aligned} \quad (\text{A.4})$$

Because $C_H > C_L$ and $w_1 > w_2$, and because $u(\cdot)$ is strictly increasing, all utility differences in (A.4) are strictly positive; hence, the contradiction. We conclude that the incentive compatibility constraint can be satisfied only for contracts such that $w_2 > w_1$.

References

- Anderhub, V., Gächter, S., Königstein, M., 2002. Efficient contracting and fair play in a simple principal-agent experiment. *Experimental Economics*, 5–27.
- Berg, J.E., Daley, L.A., Dickhaut, J.W., O'Brien, J., 1992. Moral hazard and risk sharing: Experimental evidence. *Research in Experimental Economics* 5, 1–34.
- Bolton, G., Ockenfels, A., 2000. A theory of equity, reciprocity and competition. *American Economic Review* 90 (1), 166–193.
- Camerer, C.F., Hogarth, R., 1999. The effect of financial incentives in economic experiments: A review and capital-labour-production framework. *Journal of Risk and Uncertainty* 1, 7–42.
- Chiappori, P.A., Salanié, B., 2003. Testing contract theory: A survey of some recent work. In: Dewatripont, M., et al. (Eds.), *Advances in Economics and Econometrics*, vol. 1. Cambridge University Press, Cambridge.
- Englmaier, F., Wambach, A., 2005. Optimal incentive contracts under inequity aversion. Discussion Paper No. 1643, IZA.
- Epstein, S., 1992. Testing principal-agent theory, experimental evidence. *Research in Experimental Economics* 5, 35–60.

- Fehr, E., Schmidt, K., 1999. A theory of fairness, competition and co-operation. *Quarterly Journal of Economics* 114, 817–868.
- Gneezy, U., Rustichini, A., 2000. Pay enough or don't pay at all. *Quarterly Journal of Economics*, 791–810.
- Güth, W., Klose, W., Königstein, M., Schwalbach, J., 1998. An experimental study of a dynamic principal–agent relationship. *Managerial and Decision Economics* 27, 327–341.
- Keser, C., Willinger, M., 2000. Principals' principles when agents' actions are hidden. *International Journal of Industrial Organization* 18 (1), 163–185.
- Selten, R., 1991. Properties of a measure of predictive success. *Mathematical Social Sciences* 21, 153–167.
- Selten, R., Krischker, S., 1983. Comparison of two theories for characteristic function experiments. In: Tietz, R. (Ed.), *Aspiration Levels in Bargaining and Economic Decision Making*. Springer, Berlin, pp. 259–264.