

Reputation and Contract Design*

Björn Bartling^{a)}
University of Zürich

Ernst Fehr^{b)}
University of Zurich, CESifo and CEPR

Klaus M. Schmidt^{c)}
University of Munich, CESifo and CEPR

July 17, 2008

Abstract

When should a principal control an agent? Control limits shirking but reduces the agent's productivity. We analyze experimentally how an imperfect signal about the agent's past performance affects optimal contract design. Without signal all principals control and pay low wages. If agents can acquire a reputation, many principals do not control and pay higher wages to high reputation agents. However, some agents get stuck with a poor reputation and low wage control contracts. We decompose the total effect of reputation in a sorting and an incentive effect. Competition for agents and jobs fosters the incentive effect and further increases efficiency.

[100 words]

Keywords: contract design, reputation, reciprocity, trust, control, dual labor markets.

JEL: C91, D86

*) We would like to thank Theo Eicher, Florian Englmaier, Georg Gebhardt, Hannah Hörisch, Sandra Ludwig, Christina Strassmair, Joachim Winter and seminar participants at Amsterdam, Bocconi, Edinburgh, Mannheim, Munich, Tel Aviv, Vienna, Zurich and Hebrew University for many helpful comments and suggestions. Financial support by Deutsche Forschungsgemeinschaft through SFB-TR 15 and the Excellence Initiative of the German government is gratefully acknowledged. Ernst Fehr also gratefully acknowledges support from the Swiss National Science Foundation (project number 101312-103898/1) and the Research Priority Program on the "Foundations of Human Social Behavior" at the University of Zurich.

^{a)} Björn Bartling, Institute for Empirical Research in Economics, University of Zurich, Bluemlisalpstrasse 10, CH-8006 Zurich, Switzerland, email: bartling@iew.uzh.ch

^{b)} Ernst Fehr, Institute for Empirical Research in Economics, University of Zurich, Bluemlisalpstrasse 10, CH-8006 Zurich, Switzerland, email: efehr@iew.uzh.ch and Collegium Helveticum, Schmelzbergstrasse 25, CH-8092 Zurich, Switzerland

^{c)} Klaus M. Schmidt, Department of Economics, University of Munich, Ludwigstrasse 28, D-80539 Muenchen, Germany, email: klaus.schmidt@LMU.de (corresponding author)

1 Introduction

This famous quote attributed to Lenin refers to a fundamental question in any principal-agent relationship: To what extent should the principal leave discretion to the agent? An agent who is trusted may shirk. Control limits shirking, but it also restricts the flexibility of the agent, it reduces his productivity, and it may adversely affect his motivation. In the real world we observe very different structures across and within firms. Some employees are tightly monitored and controlled, they have to obey strict workplace attendance rules, they have to follow detailed procedures how to perform a task, and they are not allowed to mix private and professional activities. Other employees enjoy a lot of discretion when and how to do their jobs. It would be possible to control at least some of their actions at a relatively small cost. However, their employers choose not to do so.

In this paper we explore experimentally how an agent’s record, i.e. his performance with other principals in the past, affects the actual and optimal design of contracts in one-shot interactions. In our base treatment the principal has no information about the agent’s past performance. She can offer a fixed wage to the agent and decide whether to control or to trust him. If she exercises control, the agent is forced to spend at least a minimum amount of effort that is higher than his cost minimizing effort level.¹ However, control comes at a cost because it reduces the agent’s productivity. Each principal interacts with 15 different agents in a row. In the beginning many principals trust their agents and try to induce high effort levels by offering high wages. However, in the long run Lenin is right: trust is rarely honored and eventually almost all principals go for control contracts with low wages - despite the efficiency advantage of trust contracts.

¹ The enforceable minimum effort requirement is smaller than the first best effort level. Otherwise the problem would be trivial because no effort enforcement problem existed. The enforceable minimum effort level can also be interpreted as a reduced form representation of an explicit incentive that ensures that agents always provide at least this effort level.

This picture changes dramatically if principals get the opportunity to observe an imperfect signal about the agent's past performance with previous principals. An agent's track record is informative about future performance. Most principals understand this and offer trust contracts with high wages to agents with a good reputation and control contracts with low wages to agents with a poor reputation. However, a significant minority of principals never trust. For the agents it is profitable to invest in a good reputation and to work hard. Most agents do so if they are offered a high wage, but they don't if the wage offer is low. But there is also a significant minority of agents who always shirk.

The interaction of heterogeneous agents and principals gives rise to a segmentation of the labor market. Some agents work hard, acquire a good reputation, and are offered trust contracts with high wages most of the time. Other agents shirk, are left with a poor reputation, and get repeatedly control contracts with low wages which further discourages them to provide effort. Overall, about 40 percent of all contracts are trust contracts.

In a third treatment, we analyze the effects of competition. Competition drives up wages for agents with a good reputation. This makes it more profitable to acquire a good reputation and it fosters learning. Principals realize that they do not get agents with the best past track records by offering control contracts with low wages. Agents realize that they are left behind if they do not have a good reputation. We observe that the fraction of trust contracts increases to more than 75 percent, wages are significantly higher, agents work much harder, and total efficiency reaches almost the efficient level. Our results show that reputation and competition are complements that reinforce each other.

Our paper is related to three branches of the experimental literature. First, there is a large literature on gift exchange games. Fehr, Riedl and Kirchsteiger (1993) and others have shown that higher wages induce more effort on average. More recently, however, Fehr, Klein and Schmidt (2007) demonstrated that the effect is typically too small to be profitably exploited by the principal. This is confirmed by our results in the base treatment where agents

cannot build a reputation. However, the reputation treatment shows that paying high wages can be profitable if the agent can acquire a reputation for good performance.

Second, there are a few recent papers (e.g. Fehr and Rockenbach (2003), Falk and Kosfeld (2006)) on “hidden cost of control” showing that agents may withdraw voluntary effort provision if the principal chooses to control them. Falk and Kosfeld (2006) consider a dictator game in which the receiver can control the proposer by forcing him to give a minimum amount.² They find that the receivers are offered more on average if they do not control the proposer. They also consider a gift exchange game that is similar to our base treatment. Again, they find (small) hidden costs of control, but they do not discuss whether it is profitable for the principal to abstain from controlling. In our base treatment control is profitable even though it yields an efficiency loss. In the reputation and the competition treatment control has no significant impact on the agent’s effort, and trust contracts are chosen because of their efficiency advantage. Thus, even if there are hidden costs of control they are too small to affect contractual choices.

Finally, there are several experimental papers that discuss the impact of reputation on cooperation. Brown, Falk and Fehr (2004) consider an experimental set-up where a principal and an agent can choose to interact repeatedly (relational contracting). They show that parties stay much longer together if the agent works hard. In contrast, we look at principals and agents who are involved in one-shot interactions in order to separate the effects of reputation and of repeated interaction. We show that reputation alone is sufficient to support cooperation. Keser (2002) and Bolton, Katok and Ockenfels (2004) consider trust games with one-shot interaction and show that the provision of feedback about the agents’ past behavior has strong positive effects on efficiency. Huck, Ruchala and Tyran (2006, 2007) consider a market for experience goods. Similar to our results they show that reputation does have a

² Fehr and Rockenbach (2003) also document an adverse effect of control. They consider a trust game in which the investor can announce that he will impose a fine on the trustee if the trustee does not return at least a minimum amount. They show that trustees return more if the investor chooses not to impose the fine.

strong positive effect that is further increased by competition.³ However, none of these papers considers the question of optimal contract design nor do they consider the principal's option to control the agent.

There is also a large theoretical literature starting with the seminal paper by Kreps, Milgrom, Roberts and Wilson (1982) showing that in finitely repeated games with incomplete information the possibility of reputation building can induce parties to cooperate in a prisoners' dilemma game in all but the last finite number of periods. In our set up a similar equilibrium in which principals pay high wages and agents work hard in all but the last periods exists. Our experimental results are largely consistent with this equilibrium.

The remainder of the paper is organized as follows. Section 2 outlines the experimental design and procedural details. In Section 3 we discuss the behavioral predictions of the self-interest model, of models of fairness, and of the notion of 'strong reciprocity.' Section 4 presents the experimental results of the base and the reputation treatment. Here we show that the total effect of reputation can be decomposed in a sorting effect and an incentive effect and that these two effects are strong complements. We also discuss the segmentation of the labor market that arises in the reputation treatment. Section 5 studies the effects of competition on effort provision and contract choice. Section 6 concludes.

2 Experimental Design and Procedures

Consider a principal who hires an agent to carry out production. The agent generates a monetary gross profit $b \cdot e$ if he expends effort e . The parameter b reflects the agent's productivity. Gross profits accrue directly to the principal, while the agent incurs private

³ In Huck, Ruchala and Tyran (2006) there is no competition in prices but only in reputations. In this paper competition has a strong positive effect. In Huck, Ruchala and Tyran (2007) there is price competition between agents for principals. This reduces the agents' profits and thereby reduces their incentives to acquire a good reputation. Thus, in their model efficiency is improved if competition is reduced by regulating prices.

effort costs $c(e) = e$, measured in monetary terms. Thus, the principal prefers the agent to choose high effort levels, but the agent prefers low effort.

The principal can offer an employment contract to the agent that specifies a fixed wage w and a desired, non-binding effort level \tilde{e} . The wage has to cover at least the costs of the desired effort. The contract cannot be conditional on effort, nor on effort costs, nor on gross profits. These variables are observable by both parties, but they cannot be verified to the courts. If the agent rejects the contract offer, no wage is paid, no effort is exerted, and both parties receive their reservation utilities of 0. If the agent accepts, the principal must pay the offered wage - irrespective of the actual effort chosen by the agent. Payoffs are given by $\Pi = b \cdot e - w$ for the principal and $U = w - e$ for the agent.

There are two types of contracts that the principal can offer: a *trust contract (TC)* and a *control contract (CC)*. These contract types differ in two dimensions:

1. *Minimum effort level*: In a TC the agent can choose an effort level between 1 and 10, whereas in a CC he must choose an effort level of at least 3, given he accepts the contract.
2. *Efficiency*: In a TC the efficiency of the relationship is given by $b = 5$, whereas in a CC the efficiency parameter is only $b = 4$.

Table 1 summarizes the differences between TCs and CCs, and the principal's and agent's payoff functions.

TABLE 1—CONTRACTS AND PAYOFF FUNCTIONS IN EACH PERIOD

	Trust Contract (TC)	Control Contract (CC)
feasible effort levels	$e \in \{1, \dots, 10\}$	$e \in \{3, \dots, 10\}$
efficiency parameter	$b = 5$	$b = 4$
payoff if contract is accepted	$\Pi = 5 \cdot e - w$ $U = w - e$	$\Pi = 4 \cdot e - w$ $U = w - e$
payoff if contract is rejected	$\Pi = U = 0$	$\Pi = U = 0$

This design captures a fundamental trade-off between control and efficiency. Control forces the agent to obey some minimum standards, but it also restricts his ability to react in a flexible and efficient way to a changing environment. For example, the principal can regulate working hours by using time cards to monitor attendance, impose reporting obligations to better assess performance, or establish strict production procedures to govern the agent's action directly. However, regulated working hours force the agent to work when he might not be most productive, reporting obligations absorb the agent's time and attention, and strict production procedures forfeit other possibly more efficient practices. The harder the agent works, the more costly it is to restrict his actions. This is reflected by the reduction of the productivity parameter b .

We started out with two treatments, the *base treatment* and the *reputation treatment*. Each treatment lasted for 15 periods and involved 18 principals and 18 agents per session. In each period, a principal was randomly matched with a new agent to eliminate repeated game effects. In the base treatment, a principal did not receive any information about his current agent. In the reputation treatment, a principal was informed about his current agent's effort choices in the last three periods.⁴ Note that a principal did neither observe the types of contract, nor the wage offers, nor the desired effort levels that his current agent faced in the last three periods. Agents knew that future principals would be able to observe their current effort choice. Apart from the information that was given to the principals in the reputation treatment, the two treatments were identical.

The reputation treatment reflects the fact that a principal sometimes has the opportunity to receive information about an agent's past performance before hiring him. For example, the principal may see letters of reference, he may have talked to a previous principal about the agent, or he may have observed the agent directly in his previous position. However, this information is incomplete. Even if the principal receives an accurate signal

⁴ If the agent did not choose an effort level because he rejected a contract, the principal received this information. In periods 1-3 a principal could only be informed about the effort levels that were available by then.

about the agent’s previous performance, he does not observe which contract induced the observed behavior and how well the agent was treated. This is reflected in the experimental design where the principal observes the agent’s actions but not the contracts he was offered.

Figure 1 summarizes the sequence of events in each period.

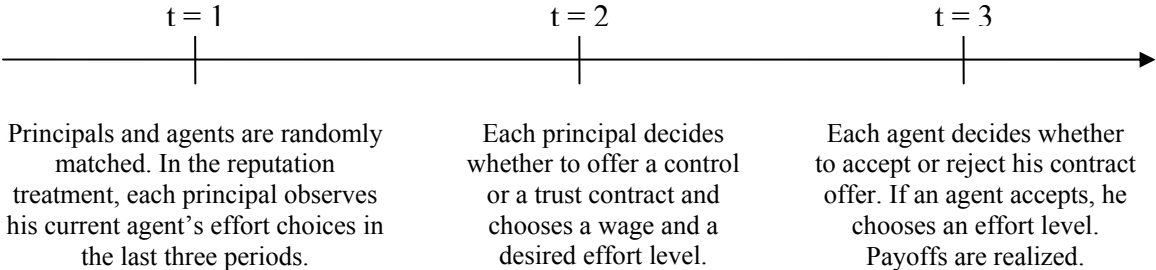


FIGURE 1—SEQUENCE OF EVENTS IN BASE AND REPUTATION TREATMENT

We conducted three sessions with the base treatment and three sessions with the reputation treatment with 36 participants in each session.⁵ Upon arrival at the lab, half of the subjects were randomly and anonymously assigned the role of a principal, the other half the role of an agent. The experiment was framed as an employment relationship. Principals were called ‘employers’ and agents ‘employees.’ Value laden terms like control, trust, or efficiency were not used. For example, the control (trust) contract was called “employment contract with (without) limitation of possible effort choices.”⁶ We also conducted two sessions with 32 participants each on a control treatment in which there was competition between principals for

⁵ The experiments were computerized with the software z-Tree (Fischbacher 2007). The recruitment was done with the software ORSEE (Greiner 2004).

⁶ After the subjects had played the treatment for 15 periods we conducted the second treatment with the same agents. While playing the first treatment, subjects did not know that a second treatment was going to be played. After the first treatment was completed, they were informed that the experiment continued with a second treatment and that the session would definitely end thereafter. Thus, each subject actually participated in both treatments. Subjects that were a principal (agent) in the first treatment remained a principal (agent) in the second treatment (no role reversal). It turned out that there are small but significant order effects. This is why we do not use the results of the second treatment in the main part of the paper. However, in Section 4.4 we use the information on each agent in order to characterize the agents’ types. The results of the treatments played in the second round are reported in the Appendix where it can be seen that they are qualitatively very similar to the results of the treatments played in the first round.

agents and between agents for jobs. This competition treatment is described in more detail in Section 5.⁷

Sessions lasted about 2½ hours and took place at the Institute for Empirical Research in Economics at the University of Zurich. Subjects were students from the University of Zurich and the Swiss Federal Institute of Technology. On average, subjects earned about CHF 46 (about \$37), which includes a show-up fee of CHF 15 (about \$12).⁸

3 Behavioral Predictions

The central question addressed by the experiments is what kind of contract the principal should use in order to induce the agent to work hard, depending on what he observes about the agent's past performance. Different behavioral approaches suggest different answers to this question.

3.1. Self-interest model

The standard neoclassical approach assumes that all people are fully rational and only interested in maximizing their own material payoff. In this case the (second best) optimal contract is straightforward. In the base treatment, the agent always chooses the effort level that minimizes his cost, which is $e = 1$ in a TC and $e = 3$ in a CC. Furthermore, he accepts all contract offers that yield a non-negative payoff. Therefore, the principal offers a wage that holds the agent down to his reservation payoff of 0 . The contract that maximizes the principal's profit is thus a CC with a wage of $w = 3$. This yields a profit of $\Pi = 4 \cdot 3 - 3 = 9$. Offering a TC with a wage of $w = 1$ yields a profit of only $\Pi = 5 \cdot 1 - 1 = 4$. This prediction is not affected by the possibility of reputation building. In the last period of the reputation

⁷ In these sessions only the competition treatment was played.

⁸ The experimental currency were 'points', and 10 points were converted to CHF1.25 (about \$1) in the base and reputation treatment and to CHF 2.50 in the competition treatment (while the participants played only 15 rounds, the competition sessions lasted as long as the base/reputation sessions with altogether 30 rounds).

treatment, agents have no reputation to lose and will thus choose the minimum effort level. Principals anticipate this and offer a CC with a wage of $w = 3$. By backward induction, this outcome is the unique prediction also for all previous periods.

***Hypothesis 1 (Self-interest Model):** If agents are paid fixed wages and have a finite horizon they always choose minimum effort levels. Therefore, in both treatments, principals should offer control contracts with a wage of 3.*

3.2 Social preferences

Models of social preferences (e.g. Rabin 1993, Fehr and Schmidt 1999, Dufwenberg and Kirchsteiger 2004) predict that some agents are “fair” and reciprocate to high wages with high effort levels, while other agents are mainly self-interested. These models also predict that controlling an agent does not reduce his effort as long as he is offered a fair wage.⁹ If the principal cannot observe the agent’s past record, her optimal contract offer depends on the share of “fair” agents in the population. If we assume with Fehr and Schmidt (1999) that about 60 percent of the population are selfish and 40 percent are fair, then CCs with low wages are optimal in the base treatment.¹⁰ However, wages have to be sufficiently high to induce agents to accept them. A control contract with a wage of 6 splits the surplus equally if the agent chooses the minimum effort of 3. Thus wages above 6 should always be accepted.

How is this prediction affected by the possibility to build a reputation? With reputation there exists an equilibrium along the lines of Kreps, Milgrom, Roberts and Wilson (1982):¹¹ All principals offer generous trust contracts in all but the last few periods to agents with a high

⁹ The reason is that in all of these models fairness is evaluated by payoff consequences only. Thus, if the wage is fair, controlling the agent has no impact on the perceived fairness of the situation. In the Fehr and Schmidt (1999) model a control contract would even increase effort of the fair-minded agents. The reason is that because of the smaller productivity parameter b fair agents will work harder to equalize payoffs.

¹⁰ The assumption of 60 percent selfish and 40 percent fair players is consistent with Fehr and Schmidt (1999) and seems to be a roughly accurate description of the behavior observed in similar experiments (see e.g. Fehr, Klein and Schmidt, 2007).

¹¹ Our game differs in several respects from KMRW (1982). First, in our game each agent interacts with each principal only once. Second, the principal does not observe the wages that previous principals offered. Third, there is not a small probability of a “commitment type” but rather a distribution of types whose concerns for fairness differ. Nevertheless, the construction of the equilibrium follows the same lines as KMRW (1982).

reputation, and control contracts with low wages to agents with a low reputation. Fair agents with a high reputation (or no reputation) accept generous trust contracts and work hard for them in all periods. They reject control contracts and trust contracts with low wages. Selfish agents mimic fair agents in all but the last few periods where they start to randomize between spending a high effort of 10 and a low effort of 1. Once they have lost their good reputation they shirk forever after.

***Hypothesis 2 (Fairness Models):** In the base treatment control contracts with low wages are optimal which induce agents to choose the minimum effort of $e=3$. In the reputation treatment principals should offer trust contracts with high wages in all but the last period if and only if the agent has a good reputation. Otherwise they should offer control contracts with low wages.*

3.3 Strong reciprocity

Fehr and Rockenbach (2003) and Falk and Kosfeld (2006) argue that control crowds out voluntary effort provision. This had been called ‘strong reciprocity’ (Gintis, 2000, Fehr and Rockenbach, 2003). [Ernst: Ist die Verwendung des Begriffs “strong reciprocity” hier o.k.?)

Agents motivated by strong reciprocity will choose lower effort levels under a CC, because controlling will be perceived as a signal of distrust—especially since control is costly to the principal and reduces efficiency. Thus, the notion of strong reciprocity implies that, for given wage levels, reciprocal agents choose higher effort levels under a TC than under a CC. These hidden costs of control could be amplified in the reputation treatment. If the principal controls an agent who worked hard in the past, the agent knows that the principal knows that he is a trustworthy agent. Such an agent might be especially offended by being controlled.

***Hypothesis 3 (Strong Reciprocity):** Control contracts crowd out voluntary effort provision. This hidden cost of control is stronger in the reputation treatment than*

in the base treatment. If the minimal enforceable effort level is not too high principals should thus use trust contracts in both treatments.

4 Reputation

In this section we discuss the effects of reputation by comparing the experimental results in the base treatment, where there is no scope for reputation effects, to the reputation treatment, where the principal observes his agent's past behavior.

4.1 The Base Treatment

In the first period, 52 percent of the principals choose a TC. However, this fraction falls quickly to less than 20 percent in the last four periods (see Figure 8 below). The large majority of contracts (more than 70 percent) are control contracts. Most principals experiment and choose a trust contract at least once, but almost all of them eventually turn to control contracts.

***Result 1 (Dominance of Control Contracts):** If agents cannot build a reputation, the large majority (71 percent) of contract offers are control contracts.*

To understand why control contracts are so popular, consider the effort choices of the agents.

***Result 2 (Effort Provision without Reputation):** Average effort increases with wage, but the effect is small. For any given wage, control contracts induce agents to spend more effort than trust contracts.*

Figure 2 displays the average effort chosen by the agents for different wage intervals. It shows that on average effort increases with wage and that agents work harder under a CC than under a TC for any wage interval.¹²

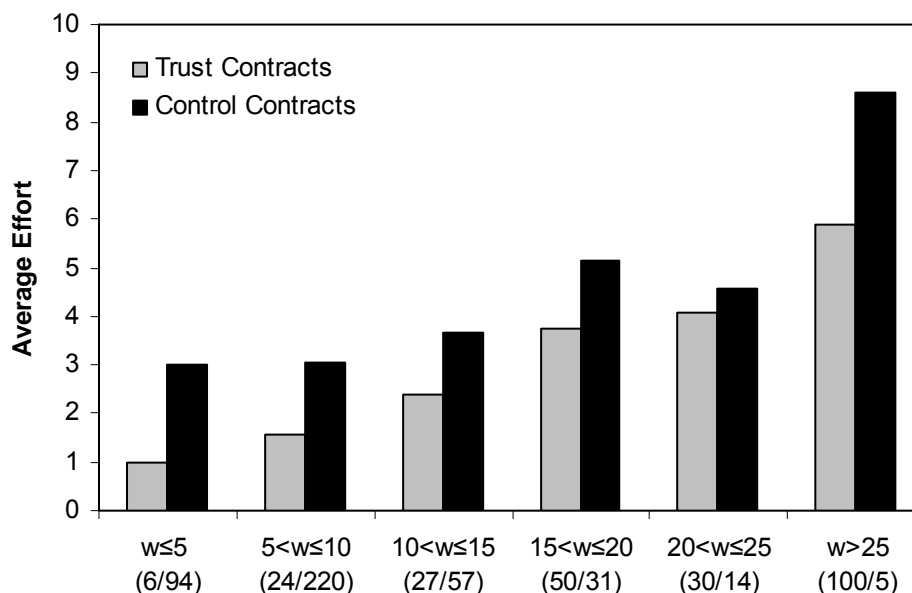


FIGURE 2—AVERAGE EFFORT CHOICE FOR DIFFERENT WAGE INTERVALS IN BASE TREATMENT

This is confirmed by a simple OLS regression of effort as a function of wage and contract type. Table 2 shows that the positive slope of the wage-effort relation is significantly reduced if a control contract is offered (0.14 versus 0.20). This may be interpreted as a crowding out of voluntary effort provision.¹³ However, the CC also raises the intercept by 2.4 units, which reflects the fact that the CC forces the agents to work at least 3 rather than 1. On average, this second effect dominates the effect on the slope of the wage-effort relation for all reasonable wage levels (for all $w \leq 40$). Thus, in our experimental set-up the ‘hidden costs of control’ observed by Falk and Kosfeld (2006) do not render control contracts ineffective.¹⁴

¹² The figure considers accepted contracts only. If a contract was rejected, no effort had to be chosen.

¹³ The difference in slope disappears in a Tobit regression.

¹⁴ Falk and Kosfeld (2006, p. 1625ff) consider a control treatment with a gift exchange game that is similar to our experiment in the base treatment. In their experiment the principal was restricted to offer wages of 10, 30, 60 or 120. Average efforts are 10.7, 12.6, 19.1 and 31 if the principal controls and 5.0, 10.1, 20.9 and 32.6 if the principal does not control (p. 1626). Thus, average effort is lower without control for wages of 10 and 30, but a

TABLE 2—DETERMINANTS OF EFFORT AND PRINCIPALS’ PROFITS IN BASE TREATMENT

	OLS (robust)	OLS (robust)	OLS (robust)	OLS (robust)
	effort	effort	effort	profit
wage	0.20***	0.20***	0.20***	-0.01
control*wage	-0.06**	-0.06**	-0.06**	-0.26**
control	2.39***	2.39***	2.39***	6.54***
desired effort	—	0.01	0.01	—
period dummies	—	—	insig. ¹⁵	—
constant	-0.28	-0.30	-0.66	-1.31
observations	658	658	658	810
R ²	0.29	0.29	0.30	0.06

Notes: In the effort regressions we consider accepted contracts only because there is no effort choice if a contract is rejected. In the profit regression we included all contracts. A low wage contract may be rejected, but this has a clearcut effect on profits.

*** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

Result 3 (Optimal and Actual Contract Choice): *Even though control reduces the efficiency parameter b , it is optimal for the principal to control the agent. Paying high wages is not profitable. However, if a principal pays a wage lower than 7, the probability of rejection is very high. Most principals choose control contracts with low wages which is profit maximizing.*

The regression results of Table 2 also show that principals’ expected profits are maximized by choosing CCs and paying low wages. When a principal offers a TC, he must increase the wage by five units to receive one additional unit of effort. Thus, with $b = 5$ the principal is just indifferent whether or not to increase the wage. If he offers a CC, he has to raise the wage by seven units to increase the agent’s effort by one unit while the marginal

little higher for wages of 60 and 120. Falk and Kosfeld do not report whether these differences are significant or whether it pays for the principal to relinquish control.

¹⁵ Only period 3 is significant at the 10 percent level and period 7 is significant at the 5 percent level.

benefit of effort is only $b = 4$. Thus, increasing the wage in a CC reduces the principal's profit.

But the principal has to ensure that his contract offer gets accepted. While almost all TCs got accepted (2 rejections out of 239), 26 percent of all CCs got rejected (150 out of 571). The acceptance probability of CCs falls sharply for wages below 7.¹⁶ The profit maximizing strategy of the principal would have been to offer a CC with a wage that just ensures acceptance. Principals seem to understand this. The average wage offer in a CC was 8.2 which yielded an average profit of 3.1. On the other hand, those principals who chose a TC offered significantly higher wages (on average 22.2) and made an average loss of 1.2.

The experimental results of the base treatment show that control contracts with low wages are optimal if agents cannot build a reputation, that the large majority of principals chooses contracts that are close to optimal, and that for any given wage control does not induce the agents to provide less effort. This is consistent with the predictions of the self-interest model (Hypothesis 1) and models of fairness (Hypothesis 2), but it does not support the notion of strong reciprocity (Hypothesis 3). The self-interest model predicts in addition that paying higher wages has no impact on effort and that the principal should offer a wage that holds the agent down to his reservation utility. However, we observe that an increase of wages has (on average) a significantly positive (but small) effect on effort, and that principals have to pay wages that give agents significantly more than their reservation utility in order to make them accept the contract. This is consistent with models of fairness.

¹⁶ While almost no CC with a wage of 8 or higher got rejected, the acceptance rates for wages of 7, 6, 5, 4, and 3 were 87, 70, 61, 17 and 12 percent, respectively.

4.2 The Reputation Treatment

In the reputation treatment the principal observes his agent's effort choices of the last three periods (if available). To aggregate this information we construct a 'reputation index' that is the average of the observed effort choices. We speak of a low reputation if the index < 3.5 , of a medium reputation if $3.5 \leq \text{index} < 6.5$, and of a high reputation if the index ≥ 6.5 .

***Result 4 (Reputation and Contract Choice):** In the reputation treatment most principals condition their contract choices on agents' reputation, offering trust contracts to agents with medium and high past effort choices much more often than to agents with low past effort choices. On average, this leads to a considerably higher share of trust contracts in the reputation treatment than in the base treatment.*

The average fraction of trust contracts is 46 percent and fairly stable over time (except for the last period where it drops to 30 percent, see also Figure 8 below). Figure 3 shows a histogram of fractions of TCs, both in the base treatment and in the reputation treatment. On average, the fraction of TCs is more than twice as high in the reputation treatment as in the base treatment. Moreover, in the reputation treatment, the fraction of TCs increases with the reputation index. Principals offer TCs almost three times more often to agents with high reputation than to agents with low reputation. Thus, principals make their contract offers contingent on agents' reputations. (See also Table 3 below).

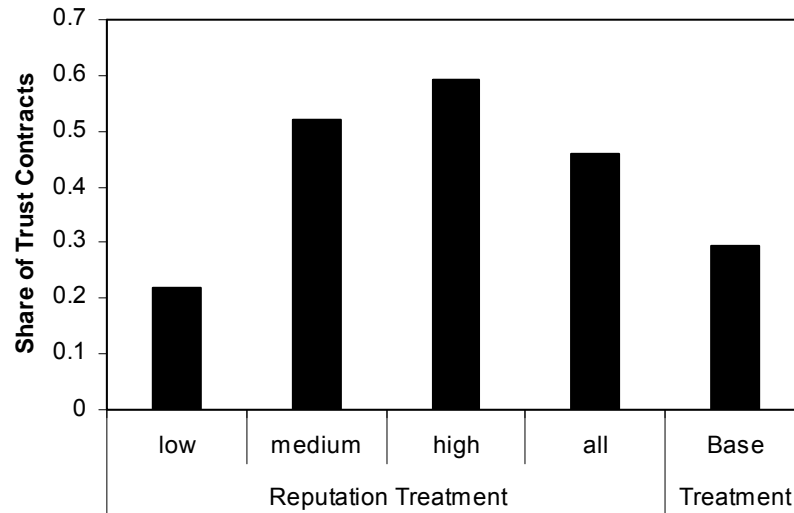


FIGURE 3—AVERAGE SHARE OF TRUST CONTRACTS

***Result 5 (Reputation and Wage Offers):** In the reputation treatment principals condition their wage offers on the agents' record, offering higher wages to agents with a medium or high reputation. On average, this leads to much higher wage offers than in the base treatment.*

On average, wages are 21 percent higher in the reputation treatment than in the base treatment (14.9 vs. 12.3). Moreover, principals offer wages that are almost twice as high on average to agents with high reputation (19.1) than to agents with low reputation (10.3).

Why are TCs more popular than CCs for medium and high-reputation agents? There are two possible explanations. First, it could be the case that control contracts crowd out voluntary effort provision and that agents work harder if they are offered a trust contract (hidden cost of control). Second, it could be the case that the trust contract does not elicit more effort than the control contract but that it is still more profitable because of the agent's higher productivity.

Result 6 (Effort Provision with Reputation): For almost all wage levels agents work more under a CC than under a TC. Thus, control does not adversely affect the agents' effort choices. The higher profitability of TCs is entirely due to higher productivity.

Evidence for Result 6 is provided by Figure 4 which displays the average effort chosen by the agents for different wage intervals. It shows that average effort increases with wages for both, CCs and TCs. However, for almost all wage levels agents work more under a CC than under a TC. Only for wages larger than 20 TCs induce more effort. However, we have only 7 CCs with wages between 20 and 25 and only one CC with a wage larger than 25 (while there are 57 and 103 TCs, respectively).¹⁷




FIGURE 4—AVERAGE EFFORT CHOICES IN THE REPUTATION TREATMENT

This is confirmed by the OLS regressions of effort reported in Table 3. Increasing the wage by one unit increases average effort by 0.20 under a TC and by 0.07 under a CC. However, offering a CC increases the intercept by 2.8. This is very similar to the effort regression in the base treatment (Table 2).

¹⁷ A Wilcoxon-Mann-Whitney rank-sum test shows that the difference in effort levels is not significant in the wage class $20 < w \leq 25$ ($p=0.91$). For the wage class $w < 25$ the difference is significant ($p=0.09$), but this should be interpreted with caution as we have only a single observation for control contracts in that wage class.

The dummy variables for medium and high reputation are not significant in the OLS regressions on effort, but they are highly significant when they are interacted with wage. Thus, medium and high reputation agents do not work harder per se, but they are prepared to spend more effort if they are offered higher wages. We also interacted the reputation dummies with wage and control. If agents with medium or high reputation would feel offended by a control contract the slope of their effort-wage relation should be smaller and the coefficients should be negative. However, both coefficients are insignificant.

TABLE 3—DETERMINANTS OF EFFORT AND PRINCIPALS' PROFITS IN REPUTATION TREATMENT

	OLS (robust)	OLS (robust)	OLS (robust)	OLS (robust)	OLS (robust)
	effort	effort	effort	effort	profit
wage	0.22***	0.22***	0.21***	0.21***	0.20*
control*wage	-0.13***	-0.12***	-0.13***	-0.12***	-0.56***
control	2.84***	2.78***	2.82***	2.74***	8.02***
medium-reputation*wage	0.08***	0.08**	0.08**	0.08**	0.34***
high-reputation*wage	0.11***	0.12***	0.12***	0.12***	0.47***
last period	-1.81***	-1.81***	-1.80***	-1.73***	-7.13***
medium-reputation	—	-0.06	-0.08	-0.03	—
high-reputation	—	-0.19	-0.24	-0.24	—
desired effort	—	—	0.03	0.03	—
other period dummies	—	—	—	insig.	—
constant	-0.23	-0.14	-0.24	-0.28	-2.02
observations	655	655	655	655	745
R ²	0.62	0.62	0.62	0.62	0.29

Notes: We consider only observations for which a reputation index is available, i.e. observations with at least one previous effort choice. In the effort regressions in addition we consider accepted contracts only as no effort is chosen in rejected contracts. In the profit regression we consider both accepted and rejected contracts. Here, a low wage may induce the agent to reject the contract which has a direct effect on the principal's profit. 'Control' is a dummy variable that takes on value 1 if the CC is offered. 'Medium-reputation' and 'high-reputation' are dummy variables that take on value 1 if the reputation index is in [3.5,6.5) or [6.5,10], respectively. 'Last period' is a dummy variable that takes on value 1 for observations in period 15.

*** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

The OLS regressions of Table 3 also show that the agent's reputation is an informative signal about his future behavior. It does not pay to offer a TC with high wages to an agent

with a low reputation, but it is profitable to do so for a medium or high reputation agent, except for the last period:

- If the agent's reputation is low, a wage increase by one unit increases effort by about 0.20 under a TC and by about 0.07 under a CC. Thus, with an efficiency factor of $b=5$ ($b=4$, respectively), paying high wages does not pay off. Furthermore, introducing control increases the agents' effort by about 2.8 units, very similar to what we observed in the base treatment.¹⁸ Thus, a CC with a low wage just high enough to be accepted by the agent is optimal.
- If the agent's reputation is high (medium, respectively), increasing the wage by one unit increases effort by 0.33 (0.30) under a TC and by 0.20 (0.17) under a CC. Thus, the principal should offer either a TC with a high wage or a CC with a low wage that just gets accepted. The profit regression shows that offering a high wage of 30 with a TC yields an expected profit of 17.7 (14.4) while a CC with a wage of 8 yields a payoff of only 6.9 (6.0). Thus, a TC with a high wage is indeed optimal, except for the last period where effort drops by 1.83.

Even though trust contracts with high wages are optimal for high and medium reputation agents, not all principals offer such contracts. There is a gap of 25 percentage points between the fraction of agents with medium or high reputation (70 percent) and the fraction of trust contracts (45 percent). This is because a sizable fraction of principals does not respond to the information about the agents' past behavior: Almost 17 percent of principals (9 out of 54) always chose the CC in the reputation treatment. Another 4 percent (2 out of 54) chooses the TC only once.¹⁹

¹⁸ As in the base treatment, we conducted separate Tobit regressions for TCs and CCs. While the wage-effort relation is flatter for CCs for agents with low reputation also in the Tobit regressions, this does not hold for agents with medium or high reputation. The details of these regressions are reported in the appendix.

¹⁹ A closer look at the data shows that all of these non-responsive principals faced an agent with a reputation index of 7 or higher at least once, and 7 out of the 11 principals even faced an agent with the maximum reputation index of 10 at least once. Furthermore, the principals that always chose the CC had agents with an average reputation index of 5.25 while the overall average of the reputation index was 5.24. Hence, the non-

Result 7 (Optimal and Actual Contract Choices): *In the reputation treatment it is optimal to offer trust contracts with high wages to agents with a medium or high reputation and control contracts with low wages to agents with a low reputation. Most principals do so. However, there is a significant minority of about 20 percent of the principals who never trust and always offer control contracts.*

The experimental results of the reputation treatment clearly reject the prediction of the self-interest model that all principals should offer CCs with low wages (Hypothesis 1). There is a large fraction of principals who offer TCs with high wages. This is indeed the optimal contract choice with agents of medium and high reputation. The experimental results also reject the predictions of the notion of ‘strong reciprocity’ (Hypothesis 3). We do not find that agents work less under CCs. For almost all wage intervals, the average effort under CCs exceeds the average effort under TCs.

The predictions of models of fairness (Hypothesis 2) are, however, partially confirmed by the experimental results. In the reputation treatment it is optimal to offer TCs with high wages to agents with a medium or high reputation. However, not all principals understand this and not all agents acquire a high reputation. This is why only half of the contracts that are offered are trust contracts. In the last period there is a large and highly significant drop of effort both in the reputation but not in the base treatment (see Figure 8 and Tables 2 and 3) as predicted by models of fairness with heterogeneous agents.

4.3 Incentive and Sorting Effects of Reputation

To better understand the effects of reputation Figure 5 compares the agents’ effort choices in the base treatment to the reputation treatment. It shows that agents work harder in the

responsive principals did not face a worse distribution of agents as the principals that conditioned their contract choices on agents’ track records.

reputation treatment no matter what wages they were offered. This effect is stronger for TCs than for CCs.

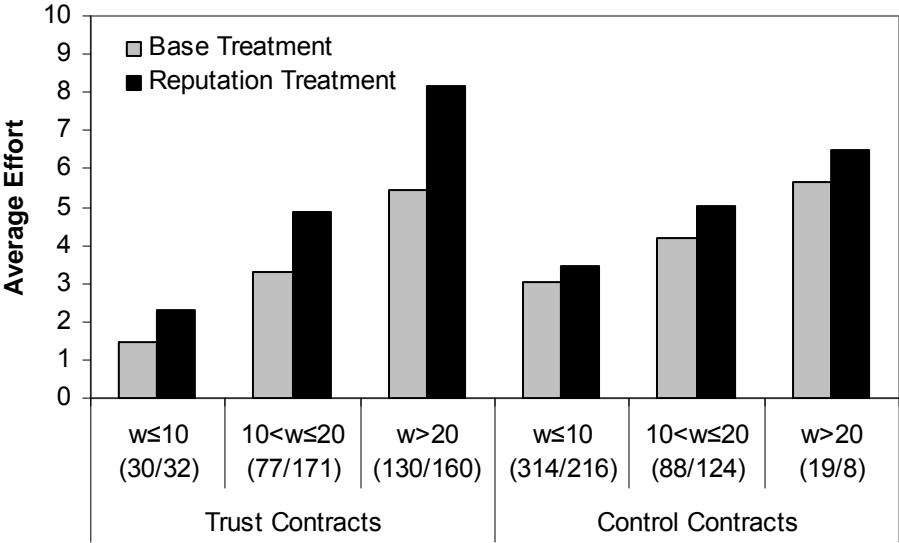


FIGURE 5—AVERAGE EFFORT CHOICES IN BASE AND REPUTATION TREATMENTS

The average effort in the reputation treatment is 5.1 as compared to 3.7 in the base treatment, an increase of 38 percent. Furthermore, there is an additional effect because more trust contracts are offered in the reputation treatment which is more efficient. If we compare the average total surplus in the reputation treatment to the base treatment we find:

Result 8 (Effort and Total Surplus with and without Reputation): *Agents work harder in the reputation treatment than in the base treatment for all wage intervals both with trust and with control contracts, and they are offered more trust contracts. On average total surplus is 58 percent higher in the reputation treatment than in the base treatment.*

There are two effects that explain this increase in efficiency by more than 50 percent:

1. *Sorting (or composition) effect:* If the principal observes the agent’s record he can use this information for sorting, that is, for offering different contracts to different types of agents. By offering a TC with high wages to agents with a good reputation and CCs with low wages to agents with a poor reputation, principals can induce high effort

from the trustworthy agents with limited risk of being exploited by an agent who is not trustworthy.

2. *Incentive Effect:* If agents know that principals observe their past record and that a good reputation will be rewarded with more generous wages (and less control), they have an additional incentive to provide more effort in order to improve their reputation.

In this section we analyze how much of the efficiency increase can be attributed to each of the two effects in isolation.

***Pure sorting effect:** Suppose principals can observe the agents' reputation and can offer different contracts to agents with different records, but agents do not know this. Therefore, they do not consider the impact of a higher reputation on future income. How large is the remaining pure sorting effect?*

We could address this question with a corresponding control experiment, but this would require misleading the participants. A more elegant way is to use our existing data and conduct a thought experiment. We consider the contracts that were offered in the reputation treatments for every given reputation index. Then we consider the data of the base treatment and ask what effort an agent with a given reputation index would have chosen if he was offered the contract that principals offered to this type of agent in the reputation treatment. This can be done by estimating effort as a function of wage, contract type and reputation index. Then we compute the hypothetical surplus in this situation. The difference between this hypothetical surplus and the actual surplus in the base treatment is the pure sorting effect.²⁰

***Pure incentive effect:** Suppose agents believe that principals observe their reputation but in fact they do not. In this case the incentive effect is present, but principals cannot use the agents' records for sorting. By how much would the average effort increase in this case?*

²⁰ The details of this calculation and the calculation of the pure sorting effect can be found in the appendix.

Again, we answer this question through a thought experiment. We consider the reputation treatment and estimate how the agents react to different contracts. However, the contracts that were offered in the reputation treatment are affected by the sorting effect. Therefore, we use the contracts that have been offered in the base treatment and compute how the agents would have reacted had they been randomly offered these contracts in the reputation treatment. Again, we compute the hypothetical surplus in this situation. The difference between this hypothetical surplus and the surplus in the base treatment is the pure incentive effect.

TABLE 4—SORTING AND INCENTIVE EFFECTS (SURPLUS)

average effort with	no sorting	sorting
no reputation incentives	10.25 (surplus in base treatment)	10.98 (hypothetical surplus with pure sorting effect)
reputation incentives	13.02 (hypothetical surplus with pure incentive effect)	16.17 (surplus in reputation treatment)

Table 4 shows the average actual and hypothetical surplus levels with and without sorting and incentive effects. The pure sorting effect is positive but small. If agents are offered the contracts of the reputation treatment in the base treatment, average surplus increases by 0.73 (about 7 percent). The pure incentive effect increases surplus by 2.77 (about 27 percent). However, the sorting and the incentive effects combined raise surplus by 5.67, almost twice as much as the sum of the pure incentive and pure sorting effects. The explanation for this difference must be that the two effects are not additive but complements. If there is no incentive effect, it is optimal to offer control contracts with low wages to everybody. Therefore, all agents will spend little effort. Hence, the reputation index contains little information about the agents and the principals cannot sort the agents. On the other hand, if

there is no sorting effect, all principals offer control contracts with low wages to the agents. Thus, agents do not benefit in the future from working harder. Therefore, without sorting the incentive effect is small.

***Result 9 (Complementarity of Sorting and Incentives):** Sorting and incentives are strong complements. Without the incentive effect all agents work very little and sorting is ineffective. Without sorting, most principals offer CCs with low wages and there are little incentives to provide high effort.*

An important implication of this result is that if the principals observe the agent's past records they should not hide this fact from the agents. Sorting without incentives is ineffective.²¹

4.4 Segmentation of the Labor Market

Why do some agents acquire a good reputation while others don't? Figure 6 shows that agents with a higher average reputation get (on average) a higher monetary income. Nevertheless, a significant fraction of the agents always choose low effort levels in the reputation treatment even when they are offered high wages.

²¹ Milinski and Rockenbach (2007) discuss the question whether an observer would want to see 'unobserved' behavior (the observed person does not know that she is observed). While this permits observing a person's 'true' behavior, more altruistic or reciprocal behavior might be triggered if the person knows that she is observed. Interestingly, humans differ from most animals in that we have a large white sclera of either side of the iris that openly signals where one watches. Milinski and Rockenbach argue that this shows that there has been a net selective advantage of signalling the directions of our gazes. This is consistent with the results of our experiments showing that principals would—if they had the choice—prefer the agents to know that their track record is observed.

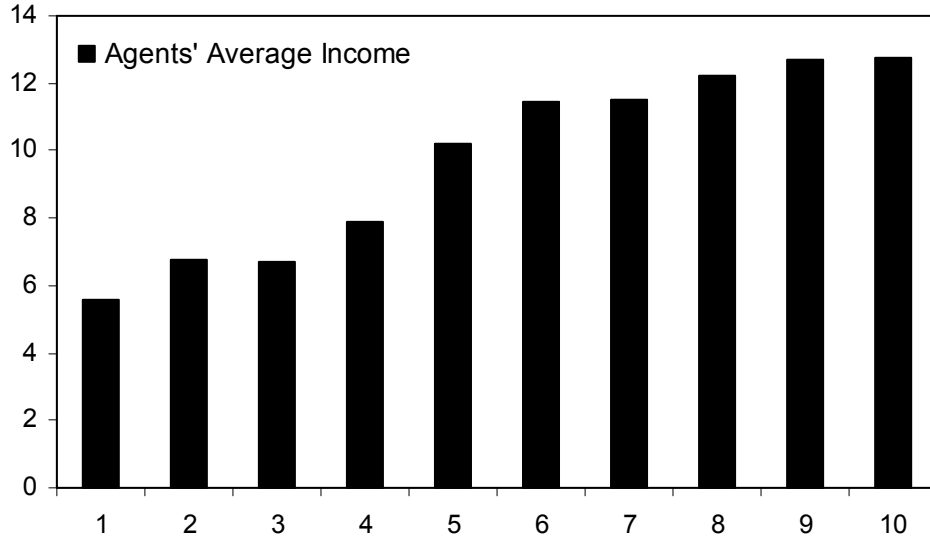


FIGURE 6—AGENTS’ AVERAGE INCOME AS FUNCTION OF THEIR REPUTATION INDEX

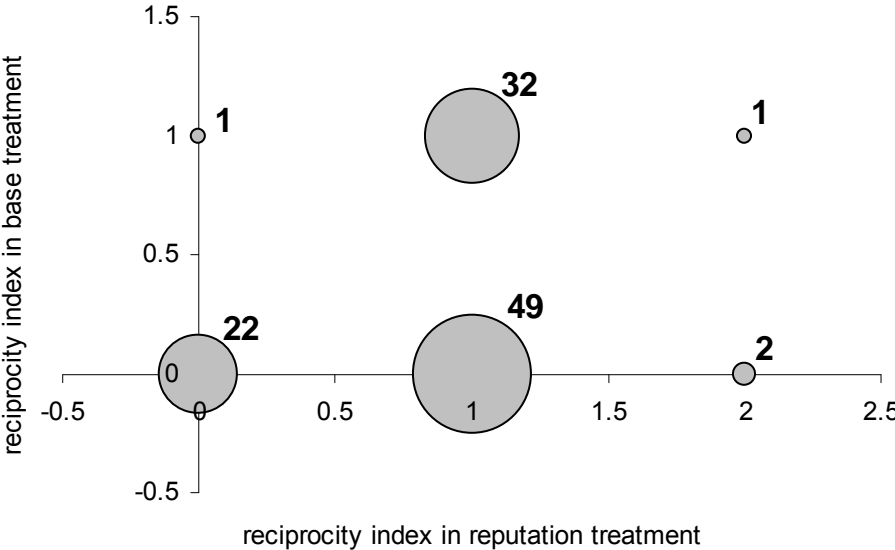
In order to classify the agents’ behavior, we use the fact that each agent participated in both treatments. So far we did not use the data of the second treatment because of small order effects (see Footnote 6 above). However, here this data can be used to classify the behavior of each agent in a simple and intuitive way. To do so we construct for each agent the following reciprocity index in the base and in the reputation treatment:

$$\alpha_i = \frac{\sum_{t=1}^{15} (\text{actual effort}^t - e_0^t)}{\sum_{t=1}^{15} (\text{fair effort}^t - e_0^t)}$$

Here e_0^t is the minimum effort agent i could choose in period t (i.e. 1 if a trust contract was offered and 3 if a control contract was offered). The “fair effort” in period t is the effort level that equalizes the payoffs of the agent and the principal given the contract offered by the principal, i.e. $b \cdot \text{fair effort} - \text{wage} = \text{wage} - \text{fair effort}$. We only consider those periods where the wage was high enough that the fair effort exceeded the minimum effort. However, for each agent we have at least one such observation. Thus, an agent who always chooses an effort level that equalizes payoffs has a reputation index of 1, while an agent who always chooses the minimum effort has a reputation index of 0. For each agent we have a reputation

index in the base treatment and in the reputation treatment. In Figure 7 these indices are rounded (to natural numbers) and plotted against each other. Three clusters of different types of agents arise:²²

- *Narrowly self-interested types:* About 20 percent of the agents have a reputation index close to zero in both, the reputation and the base treatment. These agents do not reciprocate to high wages with high effort under any condition even though this would be profitable in the reputation treatment. In the reputation treatment they choose an average effort of only 3.3. These agents get stuck with a low reputation, they are offered low wages (11.0 on average) and control contracts (77.3 percent of all contracts), and their average income is only 7.2.
- *Reciprocal types:* About 30 percent of the agents have a reputation index close to one in both treatments. These agents always reciprocate to high wages with high effort. In the reputation treatment they choose an average effort of 6.0. They acquire a medium or high reputation, they are offered considerably higher wages (15.7 on average), they are offered more trust contracts (46 percent), and their average income is 9.5.



²² Note that there is only one agent with an index of one in the base treatment and of zero in the reputation treatment. Three agents have a reputation index close to 2 in the reputation treatment, indicating that they spend considerably more than the fair effort to acquire a good reputation.

FIGURE 7—CLASSIFICATION OF AGENTS' TYPES

- *Strategic types*: A little less than 50 percent of the agents have a reputation index close to zero in the base treatment and close to one in the reputation treatment. These agents act strategically and reciprocate if their track record is observed, but do not reciprocate if low effort goes undetected by future principals. In the reputation treatment they are very similar to the reciprocal types: They choose an average effort of 5.8, they acquire a medium or high reputation, they are offered relatively high wages (16.6 on average) and mostly trust contracts (52.7 percent), and their average income is 10.8.

Result 10 (Endogenous Segmentation of the Labor Market): *Narrowly self-interested agents are predominantly offered control contracts with low wages in the reputation treatment, while agents who reciprocate (for strategic or intrinsic reasons) are much more likely to be offered trust contracts with more generous wages.*

This result suggests that people who only consider their narrow, short-term self-interest are more likely to end up with jobs that are tightly controlled and leave no rents on the table, while people who behave reciprocally are more likely to get jobs that leave more discretion and offer higher rents. This is reminiscent of the literature on dual labor markets (Doeringer and Piore 1971, Piore 1980) that provides a stylized description of actual labor markets in terms of a primary and a secondary market. In the primary market, employees enjoy higher wages and job security while in the secondary sector low wages, high turnover, and low job security prevail. Furthermore, jobs in the primary market tend to give much more discretion to employees than jobs in the secondary sector that are often tightly controlled and monitored. Piore (1980) stressed that dual labor markets often arise *within* firms. Bulow and

Summers (1986) and Saint-Paul (1997) link the description of dual labor markets with efficiency wage theories that are based on differences in monitoring costs or employment adjustment costs across the two sectors. In these models technological factors are a source of dual labor markets. In our experiment all jobs are identical. Our findings suggest that agents' individual characteristics (their reciprocal or narrowly self-interested behaviour) may also contribute to the segmentation of the labor market.

5 Competition

Introducing scope for reputation building increases the average effort of the agents by 38 percent (from 3.7 in the base to 5.1 in the reputation treatment). However, this is still far off the efficient effort level of 10. The remaining inefficiency is due to the fact that about 20 percent of the subjects do not react to the possibility of reputation building. On the one hand there are the principals who never trust and always offer control contracts with low wages. They induce agents who would have worked hard for a higher wage to withdraw their efforts. On the other hand there are the agents who always shirk. They induce all principals to turn to control contracts with low wages.

In our experiments one principal is matched with one agent in every period, and the two parties form a bilateral monopoly. In the real world there is often competition of employers for employees and of employees for jobs. Competition may have two effects. First, it may increase the returns to reputation. Principals will overbid each other to attract the agent with the best reputation, and the agents with the best reputation will get the best paid jobs. Second, competition may foster learning. Principals who offer control contracts with low wages will be left with the agents with the poorest reputations, and agents who always shirk will get the worst paid jobs. In a final treatment with competition we analyze how strong these

effects are. What is the effect of competition on contract choices, effort provision, and the segmentation of the labor market?

5.1 Experimental Design

We implement competition as follows: At the beginning of each period four principals are matched with four randomly selected agents. The principals observe the agents' reputation indices (i.e. the average of each agent's effort choices over the last three periods). Each principal has to make four (possibly different) contract offers to each of the four agents. He also has to specify the order in which the offers are made. The first offer of each principal is presented to the respective agents. Thus, in the first round of offers an agent may get several offers, just one offer, or none. The agent can now decide whether to accept any of the offers he received, but not more than one. If he accepts an offer, he and his principal disappear from the market and the agent has to make his effort choice. Those principals whose offers got rejected continue with their second offers, and so on. This algorithm guarantees that each agent receives at least one offer. He does not observe the offers received by the other agents. However, he may have to wait a while before he gets his first offer in which case he may conclude that he is not the first choice of any principal. Similarly, principals do not observe the contract offers made by other principals. However, if their offers get frequently rejected they may conclude that other principals offered more attractive contracts.

Note that in the base treatment competition cannot make any difference. Without reputation building all agents look identical and principals cannot discriminate between them. Furthermore, because there are equally many principals and agents in each matching group, the situation is strategically equivalent to a bilateral monopoly. This is why we introduced competition only in the reputation treatment.

5.2 Experimental Results

Result 11 (Competition and Contract Choice): *In the competition treatment almost all principals condition their contract choices on agent’s reputation. On average 77 percent of all accepted contracts are trust contracts.*

Figure 8 displays the fraction of trust contracts in the competition treatment over time and compares it to the reputation and the base treatment. In the competition treatment the fraction of trust contracts starts out at 65 percent in period 1 and increases to 90 percent in period 10. It remains high until the very last period, where there is a sharp drop to about 60 percent. In all periods the fraction of trust contracts in the competition treatment is at least 20 percentage points higher than in the reputation treatment which is in turn much higher than in the base treatment.




FIGURE 8—FRACTION OF TRUST CONTRACTS IN EACH OF THE 15 PERIODS

Why are trust contracts so popular with competition? Figure 9 looks at agents with low, medium and high reputations and reports the fraction of trust contracts they receive with and without competition. The effect is dramatic. With competition almost 90 percent of the agents with a high reputation are offered trust contracts (with an average wage of 25.6), while 90 percent of the agents with a low reputation are offered control contracts (with an average wage of 11.5). Thus, almost all principals make their contract choices contingent on the observed reputation of the agent. In fact, only one principal out of 32 (3 percent) never offers a trust contract (as compared to 20 percent in the competition treatment).




FIGURE 9—FRACTION OF (ACCEPTED) TRUST CONTRACTS IN REPUTATION CLASSES

This in turn induces the agents to care much more about their reputations. Figure 10 shows the distribution of (average) reputation indices of agents in the reputation and in the competition treatment. In the reputation treatment the bulk of agents had a medium reputation between 4 and 6. With competition almost 80 percent of the agents acquire a high reputation of at least 8.




FIGURE 10—DISTRIBUTION OF AGENTS' AVERAGE REPUTATION INDEX

The question arises whether the popularity of trust contracts is partly due to hidden costs of control. Figure 11 displays the average effort chosen by the agents for different wage levels. As we have seen in the base and the reputation treatment before, average effort increases with wages for both, CCs and TCs, but there is no significant difference in effort provision between the two types of contracts.




FIGURE 11—AVERAGE EFFORT IN WAGE INTERVALS

Result 10 (Effort Provision with Competition): *With competition agents spend an average effort level of 7.3. There is no significant difference of effort between CCs and TCs. Thus, the higher profitability of TCs is not due to hidden cost of control but entirely to higher productivity.*

The average effort in the competition treatment (7.3) is significantly higher than the average effort in the reputation treatment (5.1). Total surplus also significantly increases from 16.2 in the reputation treatment to 27.9 with competition. Furthermore, both the introduction of reputation and of competition yields a Pareto-improvement because principals and agents are both better off (see Figure 12).




FIGURE 12—PAYOFFS AND SURPLUS IN THE DIFFERENT TREATMENTS

The main findings of the competition treatment are summarized as follows.

Result 11 (Contract Choices and Efficiency): *In the competition treatment almost all principals offer trust contracts with high wages to agents with a high reputation and control contracts with low wages to agents with a low reputation. Almost all agents understand this, work hard and acquire a high reputation. The*

segmentation of the labor market disappears. Both, principals and agents benefit strongly from reputation and from competition.

Competition fosters learning and raises the incentives to acquire a good reputation. Thus, the results of the competition treatment are consistent with Hypothesis 2 derived from the predictions of models of outcome-based fairness.

6 Conclusions

In this paper we have shown that information about past behavior can have a crucial effect on optimal contract design. If the principal cannot observe past behavior, it is optimal to control the agent and to pay low wages. If the principal can observe past behavior and if the agent's reputation is moderate or high, it is optimal to offer trust contracts and to pay generous wages.

We find that most (but not all) principals understand this. In the base treatment, the fraction of trust contracts falls to about 10 percent in the last periods, and wages are on average so low that they just ensure the acceptance of contract offers by the agents. In contrast, in the reputation treatment, there is a constant fraction of more than 40 percent of trust contracts with generous wages. As a result, the average effort is 50 percent higher in the reputation treatment than in the base treatment. However, there still is a significant fraction of subjects who do not behave optimally. About 20 percent of the principals always offer control contracts (even to agents with a high reputation) and a similar fraction of agents always shirk. This gives rise to a segmentation of the labor market with some agents working hard and being offered generous trust contracts most of the time and some agents choosing low effort levels and being offered control contracts with low wages.

The introduction of competition fosters learning and increases the incentives to build a good reputation. With competition the fraction of trust contracts rises to more than 75 percent, and almost all agents acquire a good reputation.

Models of outcome-based fairness are largely consistent with our experimental findings, while the predictions of the classic self-interest model and the notion of strong reciprocity go astray.

One might have suspected that principals would prefer to observe the agents' past behavior without the agents knowing that they are observed. After all, in this case agents' do not try to mimic trustworthy behavior but rather reveal their "true" types. Our results show that this is not the case. If agents don't know that they are observed all of them expect to be offered control contracts with low wages, so none of them has an incentive to show that he would work hard if he could expect higher wages in the future. Therefore, principals are strictly better off if agents know that they are observed.

References

- Bolton, Gary; Katok, Elena and Ockenfels, Axel (2004), "How Effective are Electronic Reputation Mechanisms? An Experimental Investigation," *Management Science* 50(11), 1587-1602.
- Brown, Martin, Falk, Armin and Fehr, Ernst (2004), "Relational Contracts and the Nature of Market Interactions", *Econometrica*, Vol 72 (3), 747-780.
- Bulow, Jeremy I. and Summers, Lawrence H. (1986), "A Theory of Dual Labor Markets with Applications to Industrial Policy, Discrimination and Keynesian Unemployment", *Journal of Labor Economics*, Vol. 4 (3), 376-414.
- Doeringer, P.B. and Piore, Michael J., *Internal Labor Markets and Manpower Analysis*. Lexington, Mass.: Heath, 1971.
- Dufwenberg, Martin and Kirchsteiger, Georg (2004), "A Theory of Sequential Reciprocity," *Games and Economic Behavior*, Vol. 47, 268-98
- Falk, Armin and Kosfeld, Michael (2006), "The Hidden Costs of Control," *American Economic Review*, Vol. 96(5), 1611-30
- Fehr, Ernst, Klein, Alexander and Schmidt, Klaus M. (2007), "Fairness and Contract Design," *Econometrica*, Vol. 75(1), 121-154

- Fehr, Ernst and Rockenbach, Bettina (2003), "Detrimental Effects of Sanctions on Human Altruism," *Nature*, Vol. 422, 137-40
- Fehr, Ernst, Kirchsteiger, Georg, and Riedl, Arno (1993), "Does Fairness Prevent Market Clearing? An Experimental Investigation," *Quarterly Journal of Economics*, 58, 437-460.
- Fehr, Ernst and Schmidt, Klaus M. (1999), "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, Vol. 114, 817-68
- Fehr, Ernst and Schmidt, Klaus M. (2006), "The Economics of Fairness, Reciprocity and Altruism - Experimental Evidence and New Theories," *Handbook on the Economics of Giving, Reciprocity and Altruism*, ed. by Serge-Christophe Kolm and Jean Mercier Ythier, Elsevier/North-Holland, Amsterdam
- Fischbacher, Urs (2007), "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, Vol. 10, 171-78
- Gintis, Herb (2000), "Strong Reciprocity and Human Sociality", *Journal of Theoretical Biology*, Vol. 206, 169-179.
- Greiner, Ben (2004), "An Online Recruitment System for Economic Experiments," in: *Forschung und wissenschaftliches Rechnen 2003*, ed. by Kurt Kremer and Volker Macho, GWDG Bericht 63, Göttingen, Ges. für Wiss. Datenverarbeitung
- Huck, Steffen, Ruchala, Gabriele K. and Tyran, Jean-Robert (2006), "Competition Fosters Trust," CEPR Discussion Paper No. 6009.
- Huck, Steffen, Ruchala, Gabriele K. and Tyran, Jean-Robert (2007), „Pricing and Trust“, mimeo, University College London.
- Keser, Claudia (2002), "Trust and Reputation Building in E-Commerce," *CIRANO WorkingPaper 2002s-75*.
- Kreps, David M., Milgrom, Paul R., Roberts, John, and Wilson, Robert B. (1982): "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory* 27(2), 245-252
- Milinski, Manfred and Rockenbach, Bettina (2007), "Spying on Others Evolves," *Science*, Vol. 317, 464-65
- Rabin, Matthew (1993), „Incorporating Fairness into Game Theory and Economics,“ *American Economic Review* 83, 1281-1302.
- Saint-Paul, Gilles (1997), *Dual Labor Markets*, Boston: MIT-Press.

Appendix

I Tobit Regressions

TABLE A1—TOBIT REGRESSIONS IN BASE TREATMENT

	Tobit					
	Trust Contracts			Control Contracts		
	effort	effort	effort	effort	effort	effort
wage	0.38***	0.35***	0.35***	0.47***	0.41***	0.41***
desired effort	—	0.11	0.21	—	0.19	0.20
period dummies	—	—	insig. ²³	—	—	insig. ²⁴
constant	-5.29***	-5.64***	-7.36***	-5.23***	-5.78***	-5.20***
pseudo R ²	0.05	0.05	0.06	0.25	0.26	0.28
observations		237			421	
left censored		91 (at effort ≤ 1)			358 (at effort ≤ 3)	
uncensored		108			59	
right censored		38 (at effort ≥ 10)			4 (at effort ≥ 10)	

Notes: We consider only accepted contracts as there is no effort choice if a contract is rejected. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

TABLE A2—TOBIT REGRESSIONS IN REPUTATION TREATMENT

	Tobit							
	effort	Trust Contracts			Control Contracts			
		effort	effort	effort	effort	effort	effort	effort
wage	0.31***	0.40***	0.40***	0.40***	0.24***	0.30***	0.26***	0.25***
medium-rep*wage	0.10***	-0.04	-0.04	-0.05	0.18***	0.09	0.10	0.10
high-rep*wage	0.15***	0.07	0.07	0.06	0.22***	0.14	0.14	0.15
last period	-5.81***	-5.89***	-5.90***	-5.25***	-3.31***	-3.31***	-3.26***	-2.13
medium-reputation	—	2.98**	3.00**	3.23**	—	1.23	1.07	1.03
high-reputation	—	1.65	1.69	1.79	—	1.14	0.92	0.86
desired effort	—	—	-0.02	-0.02	—	—	0.13	0.13
other period dummies	—	—	—	insig. ²⁵	—	—	—	insig.
constant	-2.12***	-4.07***	-4.02***	-4.54***	-1.84***	-2.57**	-2.79**	-3.74**
pseudo R ²	0.21	0.22	0.22	0.21	0.16	0.17	0.17	0.20
observations		331				324		
left censored		42 (at effort ≤ 1)				202 (at effort ≤ 3)		
uncensored		215				112		
right censored		74 (at effort ≥ 10)				9 (at effort ≥ 10)		

Notes: We consider only accepted contracts as there is no effort choice if a contract is rejected. ‘Medium-reputation’ and ‘high-reputation’ are dummy variables that take on value 1 if the reputation index is in [3.5,6.5) or [6.5,10], respectively. ‘Last period’ is a dummy variable that takes on value 1 for observations in period 15. *** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

²³ Only period 7 is significant at the 5 percent level.

²⁴ Only periods 10, 12, and 15 are significant at the 10 percent level.

²⁵ Only period 11 is significant at the 10 percent level and period 12 is significant at the 5 percent level.

II Order Effects

TABLE A3—CHECK FOR ORDER EFFECTS IN BASE TREATMENT

	OLS (robust) base pooled effort	OLS (robust) base first effort	OLS (robust) base second effort	OLS (robust) base pooled effort
wage	0.20***	0.20***	0.20***	0.20***
control*wage	-0.07***	-0.06**	-0.10**	-0.06**
control	2.62***	2.39***	3.45***	2.39***
wage*base-second	—	—	—	-0.01
control*wage*base-s~nd	—	—	—	-0.03
control*base-second	—	—	—	1.06
base-second	—	—	—	-0.75
constant	-0.42	-0.28	-1.03	-0.28
observations	1353	658	695	1353
R ²	0.26	0.29	0.17	0.27

Notes: We consider only accepted contracts as there is no effort choice if a contract is rejected. ‘Base-second’ is a dummy variable that takes on value 1 if the base treatment is played after the reputation treatment.

*** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

TABLE A4—CHECK FOR ORDER EFFECTS IN REPUTATION TREATMENT

	OLS (robust) reputation pooled effort	OLS (robust) reputation first effort	OLS (robust) reputation second effort	OLS (robust) reputation pooled effort
wage	0.20***	0.22***	0.20***	0.22***
control*wage	-0.06***	-0.13***	0.00	-0.13***
control	1.83***	2.84***	0.74	2.84***
medium-reputation*wage	0.07***	0.08***	0.04**	0.08***
high-reputation*wage	0.10***	0.11***	0.07***	0.11***
last period	-1.87***	-1.81***	-1.97***	-1.81***
wage*rep-second	—	—	—	-0.03
control*wage*rep-second	—	—	—	0.12***
control*rep-second	—	—	—	-2.10***
med-rep*wage*rep-s~nd	—	—	—	-0.04
high-rep*wage*rep-s~nd	—	—	—	-0.04
last period*rep-second	—	—	—	-0.15
rep-second	—	—	—	1.65***
constant	0.50*	-0.23	1.42***	-0.23
observations	1290	655	635	1290
R ²	0.62	0.62	0.62	0.63

Notes: We consider only accepted contracts as there is no effort choice if a contract is rejected. ‘Control’ is a dummy variable that takes on value 1 if the CC is offered. ‘Medium-reputation’ and ‘high-reputation’ are dummy variables that take on value 1 if the reputation index is in [3.5,6.5) or [6.5,10], respectively. ‘Last period’ is a dummy variable that takes on value 1 for observations in period 15. ‘Rep-second’ is a dummy variable that takes on value 1 if the reputation treatment is played after the base treatment.

*** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

III Incentive and Sorting Effects of Reputation

To calculate the pure sorting effect we take the contract offers in the reputation treatment, which were directed towards agents with a given reputation (sorting effect is present), and “offer” them to the agents in the base treatment (no incentive effect), for whom we can construct a reputation since we know their past effort choices. The “effort choices” of the agents in the base treatment are estimated by a simple OLS regression of effort on wage, contract type, the interaction of wage and contract type, and reputation index (the average of last three effort choices), considering accepted contracts only. The regression results are reported in Table A5. With these hypothetical effort choices we can calculate the surplus that would result as the consequence of the pure sorting effect.

The calculation consists of several steps. First, we know the frequency with which the agents in the base treatment rejected contracts of a certain type and wage level. Second, we apply these rejection rates to calculate for each reputation index the average wage that was offered both in accepted trust and control contracts. Third, we calculate the frequency with which agents of a given reputation index were offered either CCs or TCs. Forth, we determine how many agents in the base treatment had a given reputation index. Given these numbers, we finally weight the hypothetical effort choices (estimated by our regression) with their relative frequency, multiply them with the corresponding efficiency parameters, and account for effort costs to arrive at the average surplus.

To calculate the pure incentive effect we take the contract offers in the base treatment (no sorting possible) and “offer” them to the agents in the reputation treatment (incentive effect is present) at random, i.e. ignoring their reputations. The “effort choices” of the agents in the reputation treatment are estimated by a simple OLS regression of effort on wage, contract, the interaction of wage and contract, considering accepted contracts only. The regression results are reported in Table A5. With these hypothetical effort choices we can calculate the surplus that would result as the consequence of the pure incentive effect. Here

the calculation is simpler. We know the frequency with which the agents in the reputation treatment rejected contracts of a certain type and wage level. Given the contracts from the base treatment and the rejection rates from the reputation treatment, we can weight the hypothetical effort choices (estimated by our regression) with their relative frequency, multiply them with the corresponding efficiency parameters, and account for effort costs to arrive at the average surplus.²⁶

TABLE A5—EFFORT CHOICES IN BASE AND REPUTATION TREATMENT

	OLS (robust) effort in base treatment		OLS (robust) effort in reputation treatment
	index available	no index available	
wage	0.22***	0.10	0.30***
control*wage	-0.07**	-0.08	-0.10***
control	2.63***	1.95**	2.19***
index	0.33***	–	–
constant	-1.80***	1.29	-0.36
observations	596	62	711
R ²	0.39	0.08	0.51

Notes: We consider only accepted contracts as there is no effort choice if a contract is rejected. ‘Control’ is a dummy variable that takes on value 1 if the CC is offered. ‘Index’ is the average of last three effort choices rounded to natural numbers. For the base treatment we estimate effort separately for the cases in which an index is available (at least one past effort choice) and the cases where no such index is available (mainly in the first period).

*** denotes significance at 1 percent, ** at 5 percent, and * at 10 percent.

²⁶ The exact calculations are available from the authors upon request.