Productivity, Trust, and Wages: The Impact of Cognitive and Non-Cognitive Skills on Contracting in a Gift-Exchange Experiment*

Florian Englmaier† Sebastian Strasser‡ Joachim Winter§

6th May 2011

Abstract

A large number of laboratory and field studies indicate that a substantial fraction of the population seems to exhibit social preferences of sorts. Recent work also shows that social preferences affect productivity and the effectiveness of incentives in labor relations. This paper studies whether such effects are taken into account in contracting decisions. We present evidence from a real-effort experiment where we elicit measures of productivity (cognitive skill) and trustworthiness (non-cognitive skill) in the worker's task and make this information available to potential employers. Both, our productivity and trustworthiness measures predict performance in a gift exchange game. Employers anticipate this correctly and are willing to pay a wage premium for these traits. Workers' reaction functions to wage offers differ according to our measure of worker non-cognitive skills.

JEL classifications: C91, J41, D86

Keywords: Cognitive and Non-Cognitive Skills, Laboratory Experiments, Gift Exchange Games, Optimal Contracts

Preliminary version. Please do not circulate.

*We are grateful to Martin Kocher and Julius Pahlke for comments and suggestions on the design. Strasser received financial support from the German Science Foundation (DFG) through GRK 801. The experiments were financed by the DFG via SFB/TR-15.

†Department of Economics, University of Konstanz, florian.englmaier@uni-konstanz.de

‡Corresponding Author: Department of Economics, University of Munich, Geschwister-Scholl-Platz 1, 80539 Munich, Germany. Tel.: +49 89 2180 9776. E-mail: sebastian.strasser@lrz.uni-muenchen.de

§Department of Economics, University of Munich, joachim.winter@lrz.uni-muenchen.de

1 Motivation

Firms spend substantial resources in their hiring procedures to select the "best" candidate for a job. In particular, an increasing fraction of firms uses both ability and personality tests in their hiring processes, see, e.g., Autor and Scarborough (2008). While the rationale for selecting the most "able" candidate is immediately obvious, recent work has highlighted that personality tests can deliver other pieces of information, too, so as to improve the chance to hire a suitable candidate for a given position. One important piece of additional information about a potential candidate would be her social preferences, see for example Englmaier and Leider (2010). In particular in the presence of moral hazard, it is valuable for firms to have access to employees that can be motivated by "social incentives" via gift exchange.

Despite the importance of information acquisition in real world contracting, the understanding of the impact of the availability of different pieces of information on the terms of a contract is surprisingly limited and in particular empirical evidence on the issue is scarce. It is precisely this gap that we want to fill in this project: We provide controlled evidence from the laboratory on how specific pieces of information about a contracting partner are used and interact with the contracting behavior between a principal and an agent in a gift-exchange game.

Research Question When making hiring and contracting decisions, firms naturally desire to minimize the risk of hiring an unsuitable candidate and thus try to learn about the qualification of a candidate, his education, his family background, etc. As a necessary simplification of reality for our experiment, we concentrate on two dimensions of information about a candidate that we consider essential on real world labor markets. First, information about what we call *productivity* henceforth is meant to capture an objective assessment whether the candidate is good at the job he is supposed to accomplish. Second, information about what we call *trustworthiness* henceforth is supposed to encompass all social and reciprocal preferences by the candidate. We consider what we measure as "productivity" in our experiment to capture underlying ability at the task as the most relevant cognitive skill in our setting, we consider trustworthiness the most important non-cognitive skill.

In a situation characterized by moral hazard, both elements play an important part in the effort decision of the agent and hence for the outcome for the principal: Controlling for social preferences, an agent who is more productive at accomplishing a certain task will produce a higher outcome for the principal. Similarly, for given productivity, a reciprocal agent will put in more effort in response to a "generous" wage offer leading to a higher outcome for the principal. We thus focus on the trade-off between these two pieces of information and their impact on contracting behavior by both principals and agents. Hence, the main research question we have in mind is to explore how and if at all these two pieces of information are conditioned upon when writing contracts. Eventually, we want to evaluate how the presence of certain skill sets and available information about them shapes labor market outcomes under moral hazard.

Design Our experiment consists of three parts, which are presented sequentially to subjects such that they do not know what will be the content of the next part. Subjects know in advance, however, that decisions in earlier parts may have an impact on later parts. In the first part, agents work on a real effort task under a piece rate contract. We use their score in this piece rate task as our measure of productivity. In the second part, agents are presented with a binary, neutrally framed, trust game. They are asked to answer in both roles whether as trustor they would trust or not trust and as trustee whether they would return trust or not. The decision as a trustee is our measure of trustworthiness and a proxy for social and reciprocal concerns. Since there are only two choices in the trust game, classification of subjects into "types" is straightforward: either they trust or not, and either they return trust or not. As a third part, half of the players are randomly assigned to be employers and the other half to be employees. Subjects play a gift exchange game where the employer first offers the employee a flat wage and the employee thereafter performs the real effort task from the first part. Before making their wage offers, principals are presented with the information from parts 1 and 2. The amount of information provided to firms is our main treatment variable. In treatment "Productivity and Trustworthiness" (hereafter **PT**) the employers are presented both measures, in treatment "Productivity"(hereafter **P**) they are only presented the productivity measure, and in treatment "Trustworthiness" (hereafter **T**) they are only

presented the trustworthiness measure.

Results We have four main findings. 1) We can predict an agent's behavior in a moral hazard situation using information elicited earlier in the experiment. Performance in the gift exchange game increases in both, the productivity and the trustworthiness measures. 2) Contracts offered by principals systematically vary with the information they have about the agent. If given this information, principals taylor their wage offers to employee types, offering more generous contracts to more productive and more trustworthy subjects. In the absence of information about worker characteristics, wage offers are unaffected. 3) We can estimate the implied willingness to pay for "better" agents in either of the two skill dimensions (cognitive and non-cognitive). As we vary the amount of information accessible to the principals (both measures, only productivity measure, only trustworthiness measure) we obtain coherent estimates of wage premia paid for the traits. 4) Workers' reaction functions to wages in a gift-exchange game differ significantly conditional on our trustworthiness measure. Trustworthy workers provide high levels of effort, but react little to a wage increase. Workers that are not trustworthy have to be motivated by the wage to exert high effort.

By comparing our treatments P and T to the PT treatment, we can rule out that the information revelation in Phase 3 distorts the elicited measures in phases 1 and 2. In P (T) it is communicated that in phase 3 only information on productivity (trustworthiness) is made available. As the elicited measures do not differ across the three treatments, we conclude that no strategic distortion of information in phases 1 and 2 takes place.

Related Literature An extensive experimental literature documents incentives and behavior in gift exchange games, see e.g. Akerlof (1982), Fehr et al. (1993), Fehr et al. (1997). This protocol has proven to be a valuable paradigm that captures incentives on real world labor markets in the laboratory. As a major finding of this literature, preferences for fairness and reciprocity serve as a powerful source of motivation to overcome the informational asymmetry between principals and agents on labor markets[1]. These laboratory studies have also been validated in the field; see e.g. Falk (2007) or Bellemare and Shearer (2009). It is now also widely acknowledged that social preferences like reciprocity or inequity aversion

[1]For references see Fehr and Gächter (2000), Fehr and Schmidt (2003) or Fehr and Falk (2008).

potentially do not only shape market outcomes or the result of bilateral bargaining, but have an important effect on the design of optimal incentive schemes as well, see Englmaier and Wambach (2010) for a theoretical treatment. One additional important empirical finding from both field and laboratory data (e.g. Dohmen et al. (2009)) is that there is substantial heterogeneity with respect to the prevalence of reciprocal inclinations and social preferences among the population (see Fehr and Schmidt (1999) or Fischbacher et al. (2001))[2].

Recent theoretical and experimental work suggests that there are also complementarities from matching incentive structures to preference types; e.g. Englmaier and Leider (2008) or Bartling et al. (2010). However, little work has been done that tests whether social preferences, if measured, predict behavior, see Englmaier and Leider (2009) for an exception, and, more to our point, how social preferences interact with cognitive skills and whether they affect contract design. [3] One recent important exception is a paper by Cabrales et al. (2010) who design an experiment where in the first phase, all subjects choose a payoff vector and play a self-chosen effort game. From these choices their preference parameters in terms of both outcome preferences and reciprocal inclination are estimated, assuming preferences à la Charness and Rabin (2002). In the second phase, it is documented that these estimated preferences predict behavior in a gift exchange game conditional on contract offers. Moreover, contract offers vary systematically with estimated preferences of principals. However, Cabrales et al. (2010) do neither use a real effort task (and hence they do not elicit cognitive skill measures), nor is information on cognitive and non-cognitive skills presented to the principals prior to their contract offers. We consider particularly this last feature essential for our understanding of the functioning of real world labor markets.

The remainder of the paper is structured as follows. Section 2 presents the experimental design. In section 3 we lay out our hypotheses and section 4 presents the results of the experiments. Section 5 discusses our findings and concludes. The appendix contains tables

[2]Preference heterogeneity and its consequences has attracted a lot of attention in the public goods literature, see e.g. Fischbacher and Gächter (2010) and the references therein.

[3]The literature on cognitive and non-cognitive skills has mainly focused on the relation between the two and their interdependence, see e.g. Heckman et al. (2006), Borghans et al. (2008b), and the relationship of the two skill sets on real world outcomes, see e.g. Murnane et al. (1995), Borghans et al. (2008a), Heineck and Anger (2010)

and figures.

2 Experimental Design

The experiment consists of three parts, which are presented sequentially to subjects such that they do not know what will be the content of the next part. Subjects know in advance, however, that decisions in earlier parts may have an impact on later parts. We ran 9 sessions with a total of 216 subjects in December 2010 at the MELESSA laboratory at LMU Munich. The subjects were invited via ORSEE, Greiner (2004), and the experiment was implemented with zTree, Fischbacher (2007). Subjects earned experimental points (EP) during the experiment. The exchange rate from EP to Euros was 1EP = 0.01 €. The experiment lasted about 90 minutes and subjects earned on average 19.5 €.

Elicitation of Productivity and Trustworthiness

In a first part of the experiment, measures of both productivity and trustworthiness are elicited from all subjects. We proxy productivity with a measure from a real effort task that consists in matching words and a four-digit code from a list.[4] Subjects perform this task for 90 seconds and are paid a piece-rate per correct answer of 10 EP that is paid out at the end of the experiment. To prohibit guessing there is a penalty for every wrong answer of 10 EP. Our measure of productivity consists in the number of correctly matched codes after subtracting all wrong answers. Note that we assume all subjects to put in full effort under the piece rate scheme such that our measure of productivity is as closely related to underlying ability as possible. In the remainder of the analysis, we therefore refer to "productivity" as the number of (#right answers − #wrong answers), or equivalently the number of *correct answers*. The corresponding payoff from part one is calculated as follows according to the

[4] A screenshot of the experimental screen of the coding task can be found in Figure 2 in Appendix A2. We conduct a trial period before the elicitation task to familiarize subjects with the computer programme.

number of correct answers:

$$10 \text{ EP} * (\#\text{right answers} - \#\text{wrong answers}) = 10 \text{ EP} * (\#\text{correct answers})$$

There is no particular training required for fulfilling this task, such that we consider our measure a reasonable proxy for a subject's ability in this particular task. Moreover, subjects are presented three screens of 30 seconds one after another with randomly generated words and codes for every new screen, so that we assume intrinsic costs for the task to be linear over time, i.e. there are no effects from fatigue or boredom. The resolution of the number of correct answers is given to subjects only after they gave expectations about average behavior (see below).

Furthermore, subjects play a standard binary trust game in neutral framing to provide a measure of trust and trustworthiness at the individual level. [5] We make use of the strategy method to get data on both trusting behavior and trustworthiness. Subjects take both decisions for both roles and at the end of the experiment they are matched with another subject, roles are randomly determined and payoffs are realized according to the decisions taken in the respective roles. Behavior in the trust game can be seen as indicative whether individual preferences are characterized by high or low levels of trust and trustworthiness when engaging in an interaction with another person. We place particular importance on the trustworthiness of subjects as this serves as a more relevant proxy of social and reciprocal concerns relevant in the gift exchange game than the initial trusting decision.

Expectations

Subsequently, we elicit subjects' expectations about the characteristics in the population in an incentivized manner. First, all participants are asked to estimate the number of correct answers in the productivity task reported before. A correct guess of the session average is rewarded with a prize of 100 EP, from which 10 points are deducted for every correct answer that the guess was away from the true value. If the difference between the guess and the

[5] See Figure 3 in Appendix A2 for the precise amounts used in the trust game.

true value exceeded 10 answers, subjects earned at worst 0 EP from this part.

In a similar vein, we ask subjects how many of the 24 subjects in their session have chosen to reciprocate trust. Subjects are rewarded for the precision of their guesses with a prize of 100 EP if their estimate was correct, and 20 points were deducted for every subject that their guess was away from the true value. Hence, if their guess was more than 5 subjects away from the true value, earnings were 0 EP from this part.

With no feedback about the choices of other participants, individual expectations about the population are a likely candidate to explain the decision on contract choices later on. The resolution of this part took also place at the end of the experiment, such that subjects entered the gift exchange without any information about the behavior of other subjects in the experiment.

Gift-Exchange Game

In the final phase of the experiment, an experimental gift exchange game is implemented in which the task to be fulfilled is identical to the real effort task in the first part of the experiment. Subjects are randomly allocated to a role of either a firm or a worker. Overall, there are 12 workers and 12 firms per session. We employ the strategy method, as firms have to submit a binding wage offer for each of the 12 workers.[6] After all wage offers have been submitted, every worker is matched randomly to a single firm, i.e. every firm has only one worker. The workers learn only the wage offer that their matched firm has determined for them before they start working for their firm. As a consequence, there is no possibility for workers to be influenced by offers that the firm has submitted for other workers or by offers that other firms have submitted to them. Subsequently, workers perform the same real effort task from the beginning for 90 seconds. Their performance then determines the payout to the firm according to the following formula:

$$\text{firm payoff} = 10 \text{ EP} * (\#\text{right answers} - \#\text{wrong answers}) - \text{wage}.$$

[6] A screenshot of the wage setting screen can be found in Figure 4 in Appendix A2.

Workers are paid their predetermined fixed wage and have non-monetary costs of effort from solving the task:

$$\text{worker payoff} = \text{wage}.$$

To avoid that agents can ruin firms by deliberately giving wrong answers we impose a lower limit for the payoff to the firm from the task at 0. This does not preclude firms from making losses if the wage exceeds the revenues generated by their worker in a particular period.

Given the nature of the task and the fact that new words and codes are randomly generated for every screen, there should be virtually no learning possibilities from doing the task a second time. After the real effort task is completed, there is feedback about the number of correct answers and hence the payoff to the firm and the worker. Both firms and workers learn only the details from their interaction, but not from the interaction between another firm-worker pair.

Treatments

Our treatment variations consist in the pieces of information elicited in parts 1 and 2 from the experiment that are made available to firms when submitting their wage profiles. We therefore conduct three treatments:

- **PT**: Information about *productivity* and *trustworthiness* is available.
- **P**: Information about *productivity* only is available.
- **T**: Information about *trustworthiness* only is available.

Information about productivity is given to firms in the form of the number of correct answers from part 1 where all subjects had to do the coding task under a piece rate scheme. Information about trustworthiness is given in the form of the binary decision as trustee in the trust game, i.e. either return trust or not return trust.[7]

[7]Since the trust game was framed in a neutral way, the actual information was labeled "left" or "right" depending on whether subjects opted for the left or the right branch of the game tree.

3 Hypotheses

A Sketch of a Model

In this section, we sketch a simple model of incentive contracts for workers who are heterogeneous with respect to productivity and trustworthiness (assuming stable preferences and productivity types).

A firm hires a worker for a fixed wage w. The interaction is one shot, effort e is not contractible and there are no contingent contracts. Hence the firm relies on gift exchange to elicit performance. Exerting effort has convex costs from effort $c(e)$ for the agent. Output π is assumed to accrue deterministically according to $e*p$ where p is the worker's productivity. Firm profits are then $\pi - w = e*p - w$.

The worker's utility $u(w)$ is given by

$$\begin{aligned} u(w) &= w + \eta(w-o)\,(e*p-w) - c(e) \\ &= w + \eta * w * (e*p-w) - c(e) \\ &= w + \eta * w * e * p - \eta * w^2 - c(e) \end{aligned}$$

where η captures the worker's reciprocal inclination, in our experiment captured by trustworthiness, and o is the outside option which we normalize to 0. In order to elicit a positive effort response the worker has to receive a wage "gift", i.e. a wage exceeding his outside option. When the reciprocal worker receives a positive wage gift his utility increases in the firm's profit. From this, and assuming that the first order condition is necessary and sufficient for an optimal response, we can determine the worker's best response e^*:

$$\begin{aligned} \frac{\partial u(w)}{\partial e} &= \eta * w * p - c^{'}(e) = 0 \\ c^{'}(e^*) &= \eta * w * p. \end{aligned}$$

The first order condition implicitly defines e^* and we immediately see that e^* increases in w, p, and η. I.e., there are complementarities between these variables and a higher wage offer should, ceteris paribus, increase performance more for more productive and more reciprocal subjects. As in treatment PT proxies for both p and η are available, we expect according effects on wages and subsequent performance.

Hypothesis 1 (Treatment PT) *In treatment PT, we expect wage offers and performance in the gift exchange game to be higher for more productive subjects and for more trustworthy subjects.*

In treatment P (T) we expect, naturally as only productivity (trustworthiness) is known, that wage in the gift exchange game increases in productivity (trustworthiness) but not in the other.

Hypothesis 2 (Treatment P) *In treatment P we expect wage offers to be higher for more productive subjects and, due to the lack of information, non-responsive to the trustworthiness measure. Performance should be positively affected by both measures.*

Hypothesis 3 (Treatment T) *In treatment T we expect wage offers to be higher for more trustworthy subjects and, due to the lack of information, non-responsive to the productivity measure. Performance should be positively affected by both measures.*

4 Experimental Results

Part One: Elicitation of Productivity

We start by reporting summary statistics for the coding task performance in part one. Subjects received three screens with 15 matches to be made each such that the maximum attainable was 45 correct answers. Only one out of 216 subjects succeeded in giving all 45 answers within 90 seconds correctly such that time was indeed the limiting factor and the way productivity was measured does not harm high productivity subjects. The average number

of correct answers given was slightly above 29 with a standard deviation of about 6 answers.

treatment	mean	sd	N	median	min	max
PT	29.07	6.73	72	29	9	43
P	29.08	6.05	72	29	12	45
T	29.15	6.08	72	29	12	44
Total	29.10	6.27	216	29	9	45

Table 1: Descriptive statistics for coding task performance in part one

Table 1 shows summary statistics and illustrates that there are no differences across treatments with all treatments being almost identical in terms of the main statistics. The distribution of the correct answers resembles a normal distribution and all standard tests for normality confirm this. As a further robustness check, we regress the coding outcome on a number of socio-demographics to see whether there is explanatory power from gender, age, subject of study or the treatment. We also control for five character traits in the framework of the Big Five Personality Test that we elicited in a control questionnaire at the end of the experiment. Table 9 in the appendix clearly indicates that there is no effect from gender nor from a quantitative orientation in the subject of study (economics, mathematics, natural sciences) on the performance in the task.[8] We do find a significantly negative effect of age on the number of correct answers which can be explained by a lower inclination of a few older subjects for a computerized task. Apart from some effect of the character trait "agreableness", there is no effect from the four other elicited personality traits on coding performance either. More importantly, we confirm that there is no distortion from strategic concerns between the treatments (and hence different levels of information disclosure) on the outcome of the productivity task.

[8]This is confirmed by a Mann-Whitney test on gender differences for coding performance across all treatments which yields a p-value of 0.54.

Part Two: Elicitation of Trustworthiness

For the second dimension of information, we let subjects play a binary trust game presented to them in a neutral frame. Since we employ the strategy method, we have data on choice behavior in both roles of the trust game for every subject. Table 2 displays the percentage of subjects' behavior in the trust game per treatment.

Treatment	no trust no returntrust	trust no returntrust	no trust returntrust	trust returntrust	Total
PT	27.8%	19.4%	11.1%	41.7%	100%
P	33.3%	19.4%	16.7%	30.6%	100%
T	38.9%	19.4%	15.3%	26.4%	100%
Total	33.3%	19.4%	14.4%	32.9%	100%

Table 2: Percentage of subjects trusting and returning trust per treatment

According to our measure, about a third of all subjects can be considered selfish in the sense that they neither trust others nor do they return trust as a trustee. In a similar vein, about one third of subjects appears to have other-regarding concerns such that they both trust and return trust. The remaining third either trusts but does not return trust or vice versa. Although there is some variation across the treatments, these patterns are quite stable in all three treatments. What is important to note is that about half of all subjects trusted and also 50 % of all subjects returned trust while 50 % did not return trust. That is to say that our design succeeds in creating a maximum of variation across subjects which in turns makes information about other subjects very valuable for the contracting phase. If in the population our binary measure of social and reciprocal concerns was distributed less symmetric, the value of the information would clearly decline if not vanish when certain character traits were only to show up in small minorities of the underlying population.

Since we deem the decision of returning trust as more indicative of an individual's concern for reciprocity, we focus in the remainder of the analysis mainly on the behavior of subjects as a second mover in the described trust game. We again control for the impact of the same socio-demographics on returning trust in a probit regression which can be found in Table 10 in the

appendix. A similar picture to above emerges with the absence of a gender and treatment effect, but an increased propensity to return trust with age. A quantitative subject of study decreases the probability of returning trust and apart from the "agreableness" measure all personality traits are insignificant. We also control for the number of correct answers in part one on the propensity to reciprocate trust in part two, but do not find any effect which confirms that there is no relationship between our measure of productivity in part one and reciprocal behavior in part two, which we summarize in our first result.

Result 1 *There are no differences across treatments in personal traits for individual productivity and reciprocal concerns. The two measures quantify two distinct dimensions of a person's characteristics.*

Elicitation of Expectations

Before subjects enter the contracting phase and learn their roles, we elicited expectations in an incentivized manner about the two measures in the population. These may serve as an important factor when deciding about wages and efforts. Expectations were elicited referring to the current session (24 subjects) of the experiment, which we consider sufficiently large that subjects perceive their impact on the session average small enough to enter their expectation about the whole population.

Treatment	# correct choices realized	# correct choices expected	# returntrust realized (per session)	# returntrust expected (per session)
PT	29.1	24.2	12.7	11.9
P	29.1	24.3	11.3	10.8
T	29.2	25.0	10.0	10.1
Total	29.1	24.5	11.3	11.0

Table 3: Subjects' expectations vs. realizations of two elicited measures in parts one and two.

We report in Table 3 both subjects' expectations and actual realizations in the three treatments. As far as the number of correct choices in part one is concerned, subjects underestimate the average number of correct outcomes by a bit less than 5 answers compared to the true average value, which is also confirmed by a highly significant Wilcoxon signed-rank

test between the own performance and the guess of the average performance ($p < 0.01$). Indeed, out of 216 subjects, 178 give a lower expectation of the average than their own coding performance in part one, 29 a higher one and 9 subjects consider themselves to be average. Note however that at this point of the experiment, subjects do not know their performance in part one explicitly but only implicitly from remembering how many correct answers they gave.

	Dep. Var.: expected average of correct choices	**Dep. Var.:** expected # subjects returning trust
# correct choices	0.623***	-0.013
	(0.050)	(0.059)
trust	0.854	1.950**
	(0.641)	(0.756)
return trust	0.738	6.708***
	(0.649)	(0.765)
female	1.147*	1.113
	(0.621)	(0.731)
age	0.083	0.003
	(0.082)	(0.097)
treatment P	0.305	-0.606
	(0.741)	(0.873)
treatment T	1.023	-0.785
	(0.741)	(0.874)
constant	2.434	6.877**
	(2.800)	(3.300)
obs	216	216
R^2	0.44	0.37

Table 4: OLS Regression on elicited expectations. Coefficients show effects relative to answers in the PT treatment. trust and returntrust are dummy variables for behavior in part two. Standard errors in brackets. *** represents significance at p=0.01, ** at p=0.05 and * at p=0.10.

When asked about the number of subjects that returned trust in their session, guesses are far more accurate and subjects correctly predict that about half of the participants chose to return trust. To gain a deeper understanding of what drives the formation of expectations, we regress expectations on behavior in parts one and two. From Table 4 it is immediate to see that subjects are strongly influenced by their own experience and past behavior. For estimating the average number of correct choices in part one, the own result from part one is highly significant. In a similar manner, having trusted and returned trust oneself in part two increases one's expectation of the number of subjects that return trust within a session. We can also see that there are no effects from the respective other measure on expectations, i.e.

subjects with a high productivity in part one are not susceptible to adapt their expectations about reciprocal inclinations in a significant manner and vice versa. This leads us to our second result:

Result 2 *Own past behavior strongly shapes expectations about the population. There are no treatment differences and there is no interdependence in expectations about average behavior between productivity and reciprocal concerns.*

Part Three: Gift Exchange Game

We subsequently present the results from the contracting phase, where initially, i.e. before the measures were elicited, all workers were told which set of information would be disclosed to firms. This set of information consisted in

- the number of correct answers in part one (in treatments PT and P) and
- the binary decision whether to return trust or not in part two (in treatments PT and T).

Firms do not have any experience or knowledge on how workers behave such that we consider firms' wage policies as the cleanest possible measure of their preferences for information about workers.

We begin by looking at wage offers received by workers. Wages were bounded to be not negative and not above 100. Every worker got one offer from each of the 12 firms, but just received and saw one relevant wage level. As a consequence, we can analyze all 12 wage offers a worker got through the strategy method, i.e. we have $108*12 = 1296$ observations. Average wages that were submitted to one single worker are given in the first column, whereas the second column lists the average of the actually chosen one. We find significant differences on the level of average wages submitted to a worker (Mann-Whitney test on average wage offers received per worker; P vs. T: $p < 0.01$, PT vs. P: $p = 0.11$ and PT vs. T: $p = 0.05$), but not on the level of the relevant wage (Mann-Whitney test on received actual wages to workers; P vs. T: $p = 0.15$, PT vs. P: $p = 0.36$ and PT vs. T: $p = 0.72$). As provided levels

Treatment	average wage offer	actual wage offer	average performance	average productivity	average returntrust
PT	52.1	48.6	22.5	29.5	63.9%
P	55.3	54.7	23.2	30.0	44.4%
T	46.4	45.6	23.3	28.6	47.2%
Total	51.3	49.6	23.0	29.4	51.6%
obs	1296	108	108	108	108

Table 5: Summary statistics for workers. The first columns show all submitted wage offers to a specific worker, the second column only the relevant wage. The last column lists the percentage of workers that returned trust in part two.

of performance by agents are significantly lower than the elicited productivity measures in part one, there is evidence that agents do not put in unconditionally full effort levels in the contracting phase.[9] Furthermore, we see that the random attribution of roles to workers and firms has not distorted our two measures in the sense that the sample means of a session (24 subjects) lie close to the means of the workers (12 subjects). To get a deeper understanding on how firms set wages, we run a series of firm fixed effect regressions on the wage offer to a specific worker that we report in Table 6.

Dep. Var.:	**I**	**II**	**III**	**IV**	**V**
Wage offer	All treatments	All treatments	Treatment PT	Treatment P	Treatment T
worker productivity	0.923***	0.648**	1.302***	1.295***	0.031
	(0.213)	(0.277)	(0.109)	(0.100)	(0.015)
worker's return trust	2.979*	2.004	5.132*	-0.747	5.412**
	(1.451)	(1.298)	(1.317)	(0.671)	(0.792)
worker productivity * treatment P		0.699**			
		(0.299)			
worker return trust * treatment T		3.965*			
		(1.764)			
constant	22.638***	23.623***	10.425	16.845**	42.948***
	(5.954)	(5.955)	(3.901)	(2.711)	(0.5344)
obs	1296	1296	432	432	432
R^2	0.04	0.04	0.07	0.09	0.01

Table 6: Regressions include firm fixed effects on wage offers; return trust is a dummy variable for worker behavior in part two. Standard errors (cluster on the session level) in brackets. *** represents significance at p=0.01, ** at p=0.05 and * at p=0.10.

[9]In what follows, we term "performance" the number of correct answers given by workers in the gift exchange relation with firms, to draw a clear semantic distinction to the measure of "productivity" in part one.

We control for the information about the worker given to firms. Already when we pool the data from the three treatments in specification I, we see that the productivity measure increases the wage offer by a bit less than one point per correct answer in part one. This relationship is highly significant. Having returned trust in part two increases the wage offer by roughly three EPs across all three treatments. When we run separate regressions for each treatment, a clear picture about the impact of information on wage setting behavior emerges: When information about productivity is given to firms, there is a wage premium of about 1.3 EPs for every correct answer given in part one (specifications III and IV). In the absence of information about productivity in Treatment T, there is no effect on the wage level. In an identical manner, we find a wage premium of slightly more than 5 points for kind behavior in the trust game, which is persistent in all treatments where information is given about it. Reassuringly, in treatment P we find no effect of information about reciprocal concerns. Accordingly, when we control for interaction effects in a pooled regression (specification II), we find the respective significantly positive effects only when the information is given in the respective treatments. The level of significance is very high for the information about productivity and somewhat lower for the information on trustworthiness. This is also due to the fact that we elicited productivity on a continuous scale and social concerns as a binary measure.

Result 3 *There is a significant wage premium that firms are willing to pay for workers with both productivity and reciprocal concerns when the information is available.*

When we look at effort decisions of workers, the question arises what influences agents most in their decision to provide effort. Along the lines of the gift exchange literature, one can argue that the main driving force will be a high wage offer such that agents reciprocate by exerting high levels of effort. This notwithstanding, the characteristics of a person in terms of productivity and intrinsic willingness to perform well at a given task can similarly affect effort levels.

To find out more about which of these rationales helps to explain worker behavior in the contracting phase, we regress the number of correct answers on the offered wage, both

Dep. Var.: # correct choices	**I**	**II**	**III**
wage	0.234***	0.239***	0.236***
	(0.037)	(0.037)	(0.038)
worker productivity	0.586***	0.607***	0.637***
	(0.162)	(0.162)	(0.172)
worker's return trust	6.425***	6.657***	6.389***
	(2.023)	(2.054)	(2.171)
treatment P		0.277	0.049
		(2.508)	(2.544)
treatment T		3.229	3.228
		(2.495)	(2.582)
female			2.084
			(2.251)
age			-0.036
			(0.239)
quant			-1.657
			(2.825)
constant	-9.136*	-11.275**	-12.080
	(5.298)	(5.602)	(9.218)
obs	108	108	108
R^2	0.39	0.40	0.41

Table 7: OLS Regression on the number of correct choices in the gift exchange game. Standard errors in brackets. *** represents significance at p=0.01, ** at p=0.05 and * at p=0.10

measures of worker characteristics and a set of controls which are reported in Table 7. We find that the offered wage has a highly significant positive impact on the amount of effort, which we take as a clear sign that gift-exchange considerations also play a role in our real-effort experiment. Moreover in a similar manner, a higher productivity measure from part one increases the performance also in the interaction between firms and workers. That is to say that firms correctly anticipate that information about high productivity is indicative of subsequent high performance by workers. A fortiori, reciprocal concerns among workers are also predictive for the effort decision. A worker that has returned trust in part two of the experiment, gives on average 6 correct answers more than a worker who has not returned trust. When we control for treatment differences or other socio-demographic characteristics of workers, we do not find any effect on the effort decision which underlines the importance of the information contained in the two measures we elicited. We summarize these findings in our next result.

Result 4 *The wage offer is an important driving force of performance in the gift exchange game. Similarly, underlying individual characteristics are significant determinants for performance.*

We next pool all contracts concluded and allocate all workers into four broad categories. To do so, we classify a worker as being of "high productivity" when her productivity in part one of the experiment exceeded the mean of all subjects and as a "low productivity worker" if it is below. Although this classification is not element of the experiment, it nevertheless allows to compare the effects of an interaction with the respective workers also in a binary way for the productivity measure. Table 8 shows the key summary statistics for all 108 concluded

Worker Type	**All Treatments**				**PT**	**P**	**T**
	Effort	**Wage**	**Firm's Profits**	**Obs**	**Effort**	**Effort**	**Effort**
high productivity & returntrust	29.1	48.1	242.8	24	28.2	30.1	29.3
high productivity & no returntrust	21.8	47.9	170.5	25	11.0	23.1	27.3
low productivity & returntrust	24.0	53.3	186.7	32	25.2	22.0	24.3
low productivity & no returntrust	17.6	48.2	127.3	27	18.3	18.1	16.8
Total	**23.0**	**49.6**	**180.6**	**108**	**22.5**	**23.2**	**23.3**

Table 8: Statistics of all contracts for the four different worker types across all treatments.

contracts and confirms the positive impact of the two elicited measures on workers' effort levels (i.e. the number of correct choices given) and hence firm profits. When we pool all treatments, we find that both characteristics increase the number of correct choices given in the gift-exchange game and hence the profit of the firm. We rule out that workers act differently when responding to a wage offer that was made by the employer under different sets of information (i.e. the different treatments) since the three columns at the right replicate the same pattern for average effort choices in PT, P and T separately. As argued above, we hence find our two measures of cognitive and non-cognitive skills to have explanatory power for the performance in the contracting phase *regardless of the pieces of information disclosed*. It is important to note that the interaction of both skills does not reduce the effort premium workers provide for each skill separately. Controlling for one skill measure, a higher scale of the other skill measure induces higher efforts from workers. While this is no news regarding cognitive skills, our results confirm an equally important impact of non-cognitive skills on the effort decision.

Result 5 *Regardless of the information disclosed, high productivity workers and workers with reciprocal inclinations provide higher levels of effort significantly.*

We end this section by presenting evidence about the interaction of the worker types with the profitability of employment strategies of firms. From the influence of the two measures on worker performance, it follows that firm's profits are increased when they interact with a high productivity worker or a worker with reciprocal concerns. This holds true regardless of the amount of information disclosed to firms for the contracting phase in the three different treatments.[10] When we compare firm profits between the different types of workers, there is a significant difference both for the binary productivity measure (Mann-Whitney, $p = 0.03$) and worker behavior in the trust game (Mann-Whitney, $p = 0.01$). An interaction with a high type in the two dimensions yields almost twice the profits to a firm than an interaction with the low type in both dimensions. Since firms cannot actively choose their worker in our design, but are allocated a worker at random and can only offer different wages, we next proceed to analyzing how the different types of workers react to an offered wage. As we are mainly interested in the effects of non-cognitive skills on contracting, we focus henceforth on the effects of different levels of trustworthiness amongst workers on contracting behavior between firms and workers.

Not surprisingly and in line with the literature, we find a significantly positive relationship between the wage and the provided effort level. Figure 6 in the appendix displays the scatterplot of the wage-effort relationship for trustworthy workers. The slope can be interpreted as the wage increment required to induce the worker to provide one additional unit of effort. Both slopes are significantly positive at 0.11 (trustworthy workers) and 0.35 (not trustworthy workers), thus making it optimal for firms to set the maximum wage of 100 for both types of workers.[11] Workers with a high non-cognitive skill display a high intrinsic propensity to

[10] E.g. the increase in profits from interaction with a reciprocal worker is about 111 EPs in the PT treatment, 38 EPs in the P treatment and 52 EPs in the T treatment.

[11] The efficiency factor for the firm of one unit of effort is by design 10 such that one unit of effort yields an additional revenue of 1.1 and 3.5 respectively for the firm at costs of 1. Table 11 in the appendix reports an OLS regression on firm's profit in the contracting phase. The regressions confirm that in our parametrization a higher wage is always profitable for firms. The revenues generated through one additional EP for the wage exceed the cost incurred by the firm in all treatments and across all types of workers.

Figure 1: Scatterplot of the wage-profit relation of firms for trustworthy workers (left panel) and not trustworthy workers (right panel). The shaded areas indicate the 95 % confidence interval around the linear regression line.

provide effort regardless of the wage level (i.e. the intercept), but react only little to extrinsic incentives through the wage offer. In contrast, workers with low cognitive skills strongly react to an increase in the offered wage, but work only little out of intrinsic motivation in the absence of a high wage offer.

As a consequence, firms' profits from an interaction with a worker that did return trust is essentially flat for all wage levels at about 200 EP. Figure 1 displays the wage-profit relationship for firms for trustworthy types (left panel) and not trustworthy types (right panel). In terms of profits, a wage increment of 1 EP leads only to an increase of 0.1 EP in profits, which is not significantly different from zero. When firms interact with a worker that did not return trust, however, there is a steep relation between the offered wage and firms' profits. An additional increment in the wage offer leads to an increase of 2.5 EP in profits, which is highly different from zero. While setting the highest possible wage is indeed the optimal

strategy for firms regardless of the worker type ex-post, we find a markable difference in reaction on behalf of workers when offered different wage levels. Trustworthy workers react very little to the wage level, but provide high levels of unconditional effort, while not trustworthy workers need a high wage offer to exert high levels of effort.[12]

The wage setting of firms ex-ante naturally also interacts with firms' individual characteristics e.g. risk aversion, since there is always the danger that a worker shirks completely. This provides a rationale why firms do not set the maximal wage of 100 - which would be optimal ex-post - but offer lower wages in the first place. Our results nicely explain why a segregation of existing contracts on the labor market prevails which is based on workers' types, see also Bartling et al. (2010). To see this, we refer to differences between the non-profit and the profit sector as an example, although the mechanism is applicable also for the labor market in general.[13] In non-profit organizations, many people provide high levels of effort while being paid a comparably low wage. If there is selection into these jobs by types with certain skills and reciprocal inclinations, it is indeed optimal for firms to pay low wages, since a wage increase would not foster agents to provide even higher effort levels. When interacting with a worker without reciprocal concerns, however, the level of the wage is indeed the main source of motivation, which makes it optimal for firms to set a high wage in the first place if the productivity of effort is high enough. A wage reduction for these workers would lead to a substantial decrease in efforts. Information about reciprocal inclination or non-cognitive skills is hence important for firms to determine their wage policy. The importance of information about workers in increased, when selection from different types of workers into jobs is absent. These arguments should hold a fortiori if we allow for more general than flat-wage gift-exchange contracts. We summarize this finding in our last result.[14]

[12]This finding of a flat wage-profit curve for reciprocal types and a steep curve for non-reciprocal types is identical across all three treatments. The p-values of the slope in separate OLS regression of profits on wages are 0.002, 0.015 and 0.016 for the latter and 0.613, 0.377 and 0.850 for reciprocal types (PT, P and T). The same holds true for the intercept where we find the identical pattern in each treatment separately.

[13]See Benabou and Tirole (2003) and Besley and Ghatak (2005) for studies of differing incentive mechanisms in the for-profit and the not-for-profit sector.

[14]See Figure 7 in the appendix for the scatterplot of the wage-effort relationship for workers with a high and a low productivity from part one. Both have a significantly positive slope, but we do not find any difference in the reaction curve of workers in this classification.

Result 6 *Different levels of trustworthiness among workers lead to a segregation of the labor market for optimal contracts. Trustworthy types react very little to the wage, but unconditionally provide high effort levels. Not trustworthy types have to be extrinsically motivated by generous wage offers and shirk in response to low wage offers.*

5 Discussion and Conclusion

We present evidence from a laboratory experiment indicating that worker performance under moral hazard can be predicted by their previously measured "type" in terms of productivity and trustworthiness. Firms seem to be aware of this and offer more generous wages to workers who are, according to the elicited measures, more productive and more trustworthy. Our results suggest that workers with better skill sets earn wage premia on the labor market. While this is no news regarding cognitive skills (productivity), we find also significant wage premia for non-cognitive skills (trustworthiness). These wage premia are, however, consistent with rational firm behavior, as under moral hazard performance is significantly positively affected by both these cognitive and non-cognitive skills. We furthermore find a heterogeneous reaction curve of workers to wage offers according to their classification into high or low non-cognitive skill types, which rationalizes different wage setting behavior in different sectors of the labor market. While moral hazard is for sure an important friction that governs contracts and incentives on labor markets, we argue in this paper that elements of asymmetric information in the form of different "types" of workers are similarly important for our understanding of real world phenomena. Experimental economics has so far mainly focused on the moral hazard component of labor market contracts, but our results indicate that elements of adverse selection in the form of heterogeneity of workers and available information about interaction partners shapes final market outcomes as well. Given the resources spent on information acquisition by firms for the hiring decision, we consider the role and the acquisition of information on labor markets both a field of high economic relevance and a promising topic for future research to foster our understanding of the functioning of labor markets in the presence of incomplete contracts.

References

Akerlof, G. (1982). Labor contracts as partial gift exchange. *Quarterly Journal of Economics*, 97:543–569.

Autor, D. H. and Scarborough, D. (2008). Does job testing harm minority workers? evidence from retail establishments. *The Quarterly Journal of Economics*, 123(1):219–277.

Bartling, B., Fehr, E., and Schmidt, K. M. (2010). Screening, competition, and job design: economic origins of good jobs. IEW - Working Papers iewwp470, Institute for Empirical Research in Economics - IEW.

Bellemare, C. and Shearer, B. (2009). Gift giving and worker productivity: Evidence from a firm-level experiment. *Games and Economic Behavior*, 67(1):233–244.

Benabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *Review of Economic Studies*, 70(3):489–520.

Besley, T. and Ghatak, M. (2005). Competition and incentives with motivated agents. *American Economic Review*, 95(3):616–636.

Borghans, L., Duckworth, A. L., Heckman, J. J., and ter Weel, B. (2008a). The economics and psychology of personality traits. *Journal of Human Resources*, 43(4):972–1059.

Borghans, L., Meijers, H., and Weel, B. T. (2008b). The role of noncognitive skills in explaining cognitive test scores. *Economic Inquiry*, 46(1):2–12.

Cabrales, A., Miniaci, R., Piovesan, M., and Ponti, G. (2010). Social preferences and strategic uncertainty: An experiment on markets and contracts. *American Economic Review*, 100(5):2261–78.

Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869.

Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2009). Homo reciprocans: Survey evidence on behavioural outcomes. *Economic Journal*, 119(536):592–612.

Englmaier, F. and Leider, S. (2008). Contractual and organizational structure with reciprocal agents. CESifo Working Paper Series 2415, CESifo Group Munich.

Englmaier, F. and Leider, S. (2009). Managerial payoff and gift exchange in the field. mimeo.

Englmaier, F. and Leider, S. (2010). Gift exchange in the lab - it is not (only) how much you give .. CESifo Working Paper Series 2944, CESifo Group Munich.

Englmaier, F. and Wambach, A. (2010). Optimal incentive contracts under inequity aversion. *Games and Economic Behavior*, 69(2):312–328.

Falk, A. (2007). Gift exchange in the field. *Econometrica*, 75(5):1501–1511.

Fehr, E. and Falk, A. (2008). Reciprocity in experimental markets. In Plott, C. R. and Smith, V. L., editors, *Handbook of Experimental Economics Results*, volume 1, chapter 38, pages 325–334. Elsevier.

Fehr, E., Gachter, S., and Kirchsteiger, G. (1997). Reciprocity as a contract enforcement device: Experimental evidence. *Econometrica*, 65(4):833–860.

Fehr, E. and Gächter, S. (2000). Fairness and retaliation: The economics of reciprocity. *Journal of Economic Perspectives*, 14(3):159–181.

Fehr, E., Kirchsteiger, G., and Riedl, A. (1993). Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics*, 108:437–460.

Fehr, E. and Schmidt, K. (1999). A theory of fairness, competitition and cooperation. *Quarterly Journal of Economics*, 114:817–868.

Fehr, E. and Schmidt, K. (2003). Theories of Fairness and Reciprocity - Evidence and Economic Applications. In Dewatripont, M., Hansen, L. P., and Turnovsky, S. J., editors, *Advances in Econonomics and Econometrics*, volume 1, chapter 6, pages 208–257. Cambridge University Press, 8th world congress of the econometric society edition.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10:171–178.

Fischbacher, U. and Gächter, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review*, 100(1):541–56.

Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, 71(3):397–404.

Greiner, B. (2004). The online recruitment system ORSEE 2.0 - A guide for the organization of experiments in economics. University of Cologne, Working Paper Series in Economics.

Heckman, J. J., Stixrud, J., and Urzua, S. (2006). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics*, 24(3):411–482.

Heineck, G. and Anger, S. (2010). The returns to cognitive abilities and personality traits in germany. *Labour Economics*, 17(3):535–546.

Murnane, R. J., Willett, J. B., and Levy, F. (1995). The growing importance of cognitive skills in wage determination. *The Review of Economics and Statistics*, 77(2):251–66.

Appendix A1: Tables

Dep. Var.: # correct choices	**I**	**II**
female	-0.378	-0.574
	(0.870)	(0.912)
age	-0.378***	-0.374***
	(0.111)	(0.112)
quant	0.504	0.492
	(1.069)	(1.097)
treatment P	-0.281	-0.214
	(1.031)	(1.042)
treatment T	-0.001	0.066
	(1.027)	(1.030)
Big Five (Extraversion)		-0.206
		(0.415)
Big Five (Agreableness)		0.877**
		(0.409)
Big Five (Conscientiousness)		0.119
		(0.385)
Big Five (Emotional Stability)		-0.288
		(0.341)
Big Five (Openness)		-0.404
		(0.425)
constant	38.232***	37.773***
	(2.855)	(3.784)
obs	216	216
R^2	0.06	0.09

Table 9: OLS Regression on the number of correct answers in part one. Coefficients show effects relative to answers in the PT treatment. quant is a dummy for quantitative orientation of studies. All Big Five measures are on a scale from 1 to 7 indicating the strength of the individual personality trait. Standard Errors in brackets. *** represents significance at p=0.01, ** at p=0.05 and * at p=0.10.

Dep. Var.: Returntrust	**I**	**II**	**III**
female	0.016	0.010	-0.036
	(0.182)	(0.182)	(0.193)
age	0.062**	0.057*	0.053*
	(0.027)	(0.028)	(0.028)
quant	-0.548**	-0.543*	-0.612**
	(0.230)	(0.230)	(0.241)
treatment P	-0.111	-0.118	-0.109
	(0.214)	(0.214)	(0.221)
treatment T	-0.253	-0.254	-0.278
	(0.215)	(0.215)	(0.219)
Big Five (Extraversion)			0.129
			(0.088)
Big Five (Agreableness)			0.203*
			(0.089)
Big Five (Conscientiousness)			-0.005
			(0.082)
Big Five (Emotional Stability)			-0.073
			(0.072)
Big Five (Openness)			-0.129
			(0.091)
# correct choices		-0.015	-0.021
		(0.014)	(0.015)
constant	-1.324*	-0.736	-0.929
	(0.691)	(0.882)	(1.027)
obs	216	216	216
Pseudo R^2	0.05	0.06	0.08

Table 10: Probit Regression on returning trust. Coefficients show effects relative to answers in the PT treatment. quant is a dummy for quantitative orientation of studies. All Big Five measures are on a scale from 1 to 7 indicating the strength of the individual personality trait with 1 being very weak and 7 being very strong. Standard Errors in brackets. *** represents significance at p=0.01, ** at p=0.05 and * at p=0.10.

Dep. Var.: Firm's Profit	**I**	**II**	**III**
wage	1.341***	1.389***	1.428***
	(0.367)	(0.371)	(0.377)
worker productivity	5.862***	6.070***	6.694***
	(1.616)	(1.622)	(2.086)
worker's return trust	64.253***	66.575***	76.469***
	(20.228)	(20.541)	(25.314)
treatment P		2.765	79.203
		(25.086)	(109.882)
treatment T		32.294	34.635
		(24.951)	(25.237)
worker productivity * treatment P			-2.004
			(3.426)
worker's return trust * treatment T			-33.689
			(44.184)
constant	-91.360*	-112.751**	-139.394**
	(52.983)	(56.017)	(69.518)
obs	108	108	108
R^2	0.26	0.28	0.28

Table 11: OLS Regression on firm's profits in the gift-exchange game. Worker productivity indicates productivity from part one. Worker's return trust indicates behavior in part two. Standard errors in brackets. *** represents significance at p=0.01, ** at p=0.05 and * at p=0.10

Appendix A2: Figures

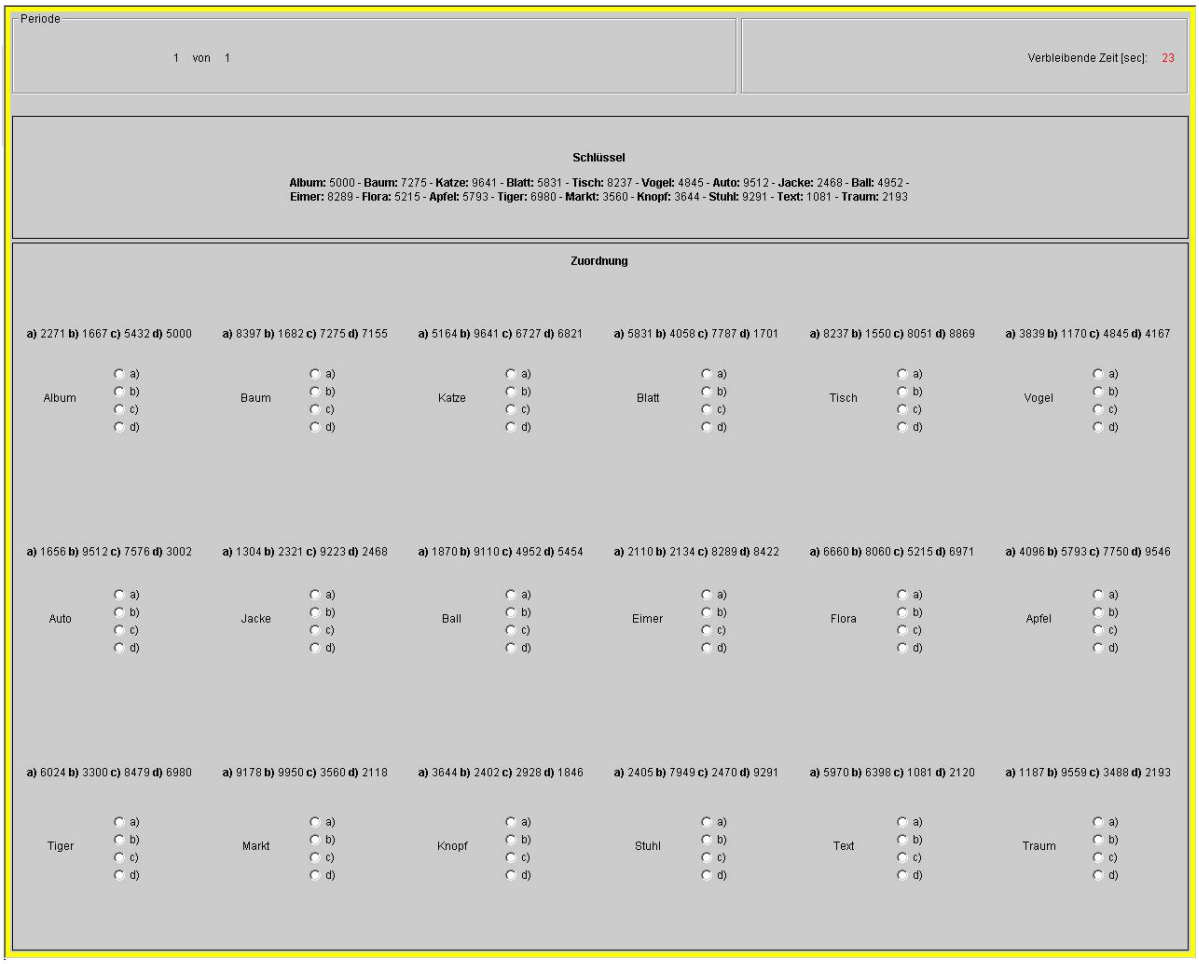

Figure 2: Screenshot of the real-effort task: The key is shown in the upper half of the screen, the matching is done in the lower half of the screen. Subjects had 30 seconds for each of the three screens.

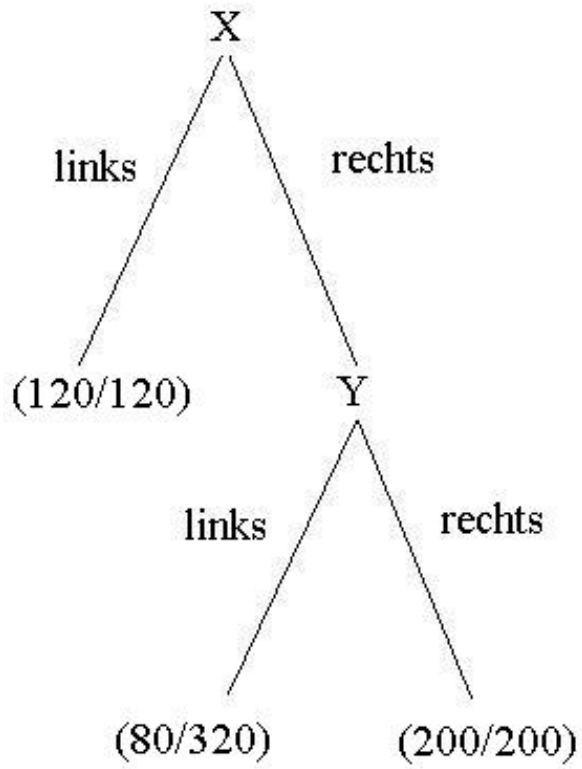

Figure 3: Presentation of the trust game to subjects. Subjects had to choose as person X (first mover) and as person Y (second mover), where at each point they could choose between "left" and "right". The corresponding payoffs are given in experimental points (EP) with the first mover's payoff listed first.

Figure 4: Wage entry screen for firms: Worker Characteristics are shown in brackets; firms enter one wage for every worker.

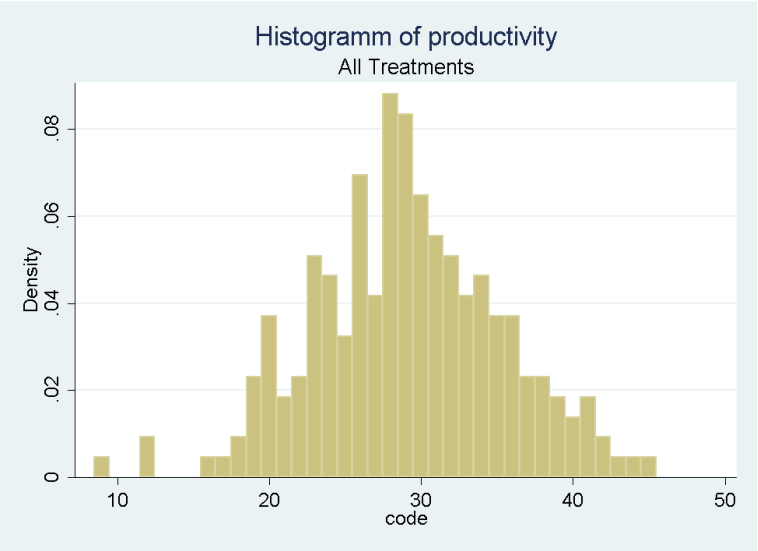

Figure 5: Histogram of the coding task performance in part one for all treatments with density of the number of correct answers.

Figure 6: Scatterplot of the wage-effort relation for trustworthy types (left panel) and not trustworthy types (right panel). The shaded areas indicate the 95 % confidence interval around the linear regression line.

Figure 7: Scatterplot of the wage-effort relation for high productivity types (left panel) and low productivity types (right panel).