Gift exchange in a multi-worker firm*

Sandra Maximiano[a], Randolph Sloof[a] and Joep Sonnemans[b]

[a] SCHOLAR, [b] CREED
Faculty of Economics and Business
University of Amsterdam
Roetersstraat 11
1018 WB Amsterdam
Email: S.M.S.N.Maximiano@uva.nl

Revised version: May 2006

Abstract

One of the main findings of a large body of gift exchange experiments is that in an incomplete contracts environment workers on average do not shirk and usually provide more than the minimum enforceable effort level. In general, 40 to 60 percent of the workers reward higher wages with higher effort. These results are observed for simple one-employer – one-worker relationships. In this paper we investigate whether they generalize to the more realistic situation in which the employer employs several workers. We compare a bilateral gift exchange game with a treatment in which each employer has four workers. We find that effort levels in the latter treatment are only marginally lower. Gift exchange thus appears to be robust to increases in the size of the workforce.

* We thank Hessel Oosterbeek for his help in designing this study and Robert Dur, Arno Riedl and two anonymous referees for valuable comments that improved the paper considerably. An earlier version of this paper was presented at conferences in Amsterdam, Steyr, Cordoba and San Francisco

1. Introduction

The main objection mainstream economists typically raise against laboratory experiments is their (supposed) lack of external validity: To what extent can the results of such experiments be generalized to more complex real world phenomena? One of the major criticisms concerns the artificiality and simplicity of the conditions and relations created in the laboratory, which do not prevail, and are also not representative of those in reality (cf. Loewenstein, 1999, Schram, 2005). Experimentalists typically defend and motivate their simplified design by referring to the economic theory to which the experiment relates. Any theoretical model necessarily has to simplify matters by leaving out many elements of reality and has to make simplifying assumptions about rationality, (homogeneity of) preferences, type of institutions etc. For practical reasons, the experimentalist who subsequently wants to test the model has to simplify even more and will typically focus only on what s/he considers to be the core elements of the model. In describing the design of the experiment the experimenter will argue that according to standard theory these simplifications are reasonable. The typical justification given is that "According to standard theory, it should not make a difference that…". But what if the experimental results point out that the behavior of subjects is not in line with standard theory? This may fire back and invalidate the experiment as (simplified) representation of the economic reality of the outside world.

A case in point is Akerlof's (1982) theory of the labor market as a gift exchange and the experimental tests of this model (see e.g. Fehr et al., 1993; Fehr and Gächter, 2000, summarize a series of gift exchange experiments). In the gift exchange model employers pay non-minimal wages to workers, who in response choose higher than minimum effort levels. As in all models, reality is necessarily simplified. For example, in reality most firms have multiple hierarchical levels and the owners of the firm (shareholders) are often not directly responsible for setting wages. In the experimental tests the situation is further simplified by limiting each employer to hire at most one worker, assigning participants randomly to the role of either employer or worker, by the choice of a specific (simple) payoff scheme etc. These design features are justified by common assumptions about individual rationality and preferences: agents are exclusively money maximizers. The typical experimental findings are that

wages and effort are above the minimum levels, supporting Akerlof's theory, but refuting the very assumptions that were used to justify the experimental design.

However, if players have social preferences,[1] generalizing the experimental results to the real world can be problematic. For example, it is not self-evident that the findings can be generalized to situations with many workers per employer. An individual with outcome-oriented social preferences who prefers equal earnings for everybody, will as a worker in the standard experiments respond to a high wage with a high effort, but will be (much) less inclined to do so when the employer has many workers. This holds because in the latter case the employer is likely to earn much more than any of her workers. Clearly, most labour relationships are more like the second situation than the first one. A similar remark applies to generalizing the results to multi-layer hierarchies in which the manager who decides on the wage is not residual claimant (cf. Baker et al., 1988).

Acknowledging the general problem of external validity, in this paper we focus on one particular aspect of it already mentioned above. We do so in the context of the gift exchange game. In particular, we compare a one-employer – one-worker (1-1) treatment with another one-employer – four-workers (1-4) treatment. In the 1-4 treatment the employer is playing the same game as in the 1-1 treatment, but with four workers at the same time. If wages and efforts would be the same in both treatments, the employer would make four times more money in the 1-4 treatment and, like in the real world, employers would make much more money than workers.

Frankly, we did not believe that the gift exchange model would survive the more realistic "multiple workers per employer" design. We were wrong. In the 1-4 treatment effort levels are only marginally lower than in the 1-1 treatment. It thus appears that the gift exchange relationship is quite robust to increases in the size of the workforce. Our results also suggest that an important driving force behind gift-exchange is intention-based reciprocity, rather than exclusively outcome-oriented social preferences. The fact that the employer earns much more than the worker does in the 1-4 treatment seems only a minor consideration in a worker's decision whether to reciprocate (either negatively or positively) a given wage offer.

[1] Models of social preferences assume that players do not care solely about their own material payoffs. We distinguish between two types of social preferences (see Camerer, 2003, for an overview). The first type refers to *outcome-oriented social preferences*. Here players care about the distribution of payoffs. The second type concerns *intention-based reciprocity* models in which players care about the intentions that drive other players' actions. Here a player is willing to sacrifice own material payoffs in order to reciprocate, either rewarding kind (fair) or punishing unkind (unfair) behavior.

Some experimental studies already exist that explore fairness and reciprocity issues in a one-principal – two-agents incomplete contracting framework. Both Meidinger et al. (2000, 2003) and Warglien and Rossi (2001) investigate whether cooperation between two team members is affected by the principal's (common) contract offer. They observe that the more generous the offer, the higher the level of cooperation among them. Güth et al. (2001b) consider a situation in which the principal can offer a menu of contracts to two independent agents with different productivities. It appears that public observability of the co-worker's contract induces the employer to compress compensation schemes (i.e. fixed wage plus a return share). Because the worker benefits from higher effort levels himself through the return share, actual effort provided is not driven by pure gift exchange. The study that comes closest to ours is Charness and Kuhn (2004). They focus on pure gift exchange in a setup where the two workers have different productivities. Their main interest lies in horizontal fairness concerns between workers and whether this leads to wage compression by the employer. In contrast, we are concerned with vertical fairness between an employer and her workers. Moreover, in our 1-4 treatment there are four (rather than two) workers per firm and each worker receives the same wage by design.

Other aspects of gift exchange robustness are addressed in a number of recent studies. Charness et al. (2004) observe that when subjects are provided with a comprehensive payoff table, employers offer lower wages and workers are less reciprocal compared to a treatment in which subjects just had the necessary information to compute their payoffs. Engelmann and Ortman (2002) make three comparisons: corner versus interior equilibrium solution, high versus low efficiency gains and abstract versus labor market frame. They find that efficiency gains interact with framing in affecting wage choices. For given wage offers, however, effort levels are unaffected by both efficiency gains and framing. Concerning subject pool effects, Hannan et al. (2002) find significant differences between Pittsburgh undergraduates and MBAs. The latter choose effort levels that are larger or equal to those reported by Fehr et al. (1993, 1998a), but the undergraduates provide less effort.[2] Hannan et al. also consider productivity differences between firms and the impact of adding an unenforceable 'requested effort level' to the wage offer. By and large they find that productivity differences do not affect worker's reciprocal effort choices, while adding

[2] On the other hand, Charness (2004) uses Berkeley students and gets results that are similar to the European subject pool experiments.

non-binding requests resulted in increased overall effort from undergraduates (but not for MBAs). Finally, Fehr et al. (2002) report no evidence that less effort is provided when stakes are increased.

The remainder of this paper is organized as follows. In the next section we present the version of the gift exchange game that we consider, together with an intuitive discussion of the theoretical predictions under various assumptions about social preferences. Section 3 describes the experimental design and the main differences with earlier experiments. Results are presented in Section 4. Section 5 discusses the robustness of these results and Section 6 summarizes and concludes.

2. The gift exchange game and the theoretical predictions

2.1 The gift exchange game

We consider the following setting, based on the bilateral gift exchange game of Fehr et al. (1998a).[3] There is a firm consisting of one (female) employer and k (male) workers, with k≥1. The employer has an initial amount of capital (endowment) $k \cdot \omega$ available for employing the k workers and paying their salary, while workers have an initial wealth of $\omega \geq 0$. The employer moves first and chooses a wage $w \geq 0$ for each of her k workers; she has to pay the *same* wage to all her workers. (Workers cannot reject the wage offer made.) Subsequently, workers simultaneously decide how much effort e_i to provide at cost $c(e_i)$, for i=1,..,k. Here $c(e_i)$ denotes the costs of effort in monetary terms. Note that workers choose their individual effort without knowing the effort choices of their co-workers. The marginal value product of effort is fixed per unit of effort provided and equals $v>0$.

Under the above assumptions, the monetary payoffs m_i of worker i equal:

$$m_i = w - c(e_i) + \omega \qquad \text{for i=1,..,k} \tag{1}$$

[3] Our version differs in some notable ways from the one used in the 'typical' gift-exchange experiment, see the next section for a discussion.

For notational convenience we will refer to the employer as being player 0. Her payoffs in monetary terms are then given by:

$$m_0 = v \cdot \sum_{i=1}^{k} e_i - k \cdot w + k \cdot \omega \qquad (2)$$

In the experiment agents' choices are restricted in the following way. For the employer, the wage w has to be a multiple of 5 and between 0 and 100. The effort choice of each worker has to be an integer between 1 and 10. The value of v is set equal to 10. The cost of effort function is given by Table 1.[4] The initial wealth parameter ω is set equal to ω=90.

Table 1. The cost of effort function

Effort e	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*
Costs c(e)	0	1	2	4	6	8	10	12	15	18

Our treatment variable is the number of workers k per employer. We consider two treatments, viz. k=1 and k=4. The first one is referred to as the 1-1 treatment and provides the baseline. The second one is labelled the 1-4 treatment. Note that treatment 1-4 is just a four times replica of treatment 1-1; the employer has four times as much capital available and employs four rather than just one worker.

2.2 Theoretical predictions

In this subsection we discuss the theoretical predictions under various assumptions about social preferences in an informal and intuitive way. A more formal discussion is relegated to Appendix A. Our main interest lies in worker behavior. When workers

[4] The same range of effort levels and the same cost function is used in e.g. Fehr et al. (1993, 1996a, 1996b, 1998a, 1998b) and Hannan et al. (2002). However, the employer's payoff function that we use differs from the one used in a large part of previous gift-exchange experiments. We will return to this point in Subsection 3.1.

are selfish and care about their own monetary payoffs only, they will provide the minimum effort level regardless of the wage offered. This prediction is independent of the number of workers k employed, i.e. the size of the workforce is irrelevant for the predicted outcome. As discussed in the Introduction, this fact is typically used to justify the focus on the simpler *bilateral* gift-exchange games where k=1.

This independence no longer holds when workers also care about the payoffs of others. In general two types of social preferences can be distinguished. First, under *outcome-oriented social preferences* agents care about the distribution of payoffs, and possibly about overall payoffs as well. Second, in case of *intention-based reciprocity* agents may want to punish unkind behavior and reward kind behavior, even when this lowers their own material payoffs. Both types of preferences may induce the worker to provide more than minimum effort, although they do lead to somewhat different predictions about the wage-effort relationship in the two treatments.

First assume that workers have outcome oriented-social preferences and dislike inequality like in Fehr and Schmidt (1999). They may then be willing to reduce inequality by choosing higher than minimum effort. In the 1-1 treatment the worker is ahead of the employer whenever $10e + c(e) \leq 2w$. For $w=0$ this inequality cannot be satisfied and the worker is necessarily behind. Because higher effort would just increase inequality, he will choose $e=1$ in response. However, for non-minimal wages $w\geq 5$ the worker earns more than the employer when the minimum effort is chosen. In that case only a weak aversion to inequality is needed to induce the worker to provide one unit of effort more. This will cost him 1 and yields the employer an additional payoff of 10. More generally, as long as the worker stays ahead he may be willing to put in extra units of effort, given the attractive ratio of low marginal effort costs (at most 3, cf. Table 1) to high marginal rewards (10). For wages of 60 and higher the worker is ahead for all feasible effort levels. A sufficiently inequality-averse worker will thus provide the maximum effort level of $e=10$ for all $w\geq 60$. Overall, a significantly steep wage effort relationship with an intercept at $e=1$ results.

Next consider what the same inequality-averse worker will do in treatment 1-4. In this case the employer's endowment equals 360, four times as much as the amount of 90 the worker gets. As a result, the worker is necessarily behind whenever $w\leq 60$ and for these wages he just provides $e=1$. For higher wages he may be willing to supply more than minimum effort, depending among other things on his beliefs

about what the three other workers within the firm will do (cf. Appendix A). Assuming he believes that they will behave just like him, the worker is ahead of the employer whenever $40e + c(e) + 270 \leq 5w$. For the maximum wage w=100 the worker is then willing to choose e=6 at most, because for $e \geq 6$ he is necessarily behind. A sufficiently inequality-averse worker will monotonically increase his effort from e=1 to e=6 when wages increase from 60 to 100. Over the full range of wages from 0 to 100 the wage-effort relationship is thus significantly *less* steep than in treatment 1-1 (but also has an intercept at e=1). Given that this prediction is driven by the difference in initial endowments between employers and workers, it is important to understand how subjects perceived these endowments. In Section 5 we return to this.

Besides equality workers may care about total surplus as well. In that case they may be inclined to sacrifice to help the employer even when the latter is ahead. The model of *quasi-maximin* preferences introduced by Charness and Rabin (2002) allows for such efficiency-based concerns.[5] It assumes (among other things) that agents prefer a higher sum of total payoffs and also have a special concern for helping the worst-off player. In the context of the gift exchange game this implies that (all else equal) the worker always prefers higher earnings for the employer, but puts relative more weight on her well-being when she is behind. This causes that the worker is more inclined to choose higher effort levels in treatment 1-1 than in treatment 1-4, just like under inequality-aversion. The difference in slopes is likely to be smaller though, because efficiency-based preferences may induce the worker to supply more than minimum effort even for a wage of zero. This depends on how strong he cares about efficiency. In case this concern is weak the wage-effort relationship has a significantly smaller slope in treatment 1-4 than in 1-1 and in both treatments the intercept is at e=1. This coincides with the predictions under inequality-aversion. When the concern for efficiency increases, the lowest effort the worker supplies shifts up. This decreases the difference in slopes between the two treatments, but necessarily increases the intercept of the wage-effort relationships at the same time.

From the above it follows that neither inequality-aversion nor quasi-maximin preferences *per se* predict the simultaneous occurrence of both an equally steep wage-

[5] Several gift-exchange experiments have confirmed that employer-worker interactions generate substantial efficiency gains, given the attractive ratio of low marginal effort costs to high marginal rewards imposed. Therefore, the efficiency-enhancing role of gift-exchange has been considered as a possible cause for the behavioral deviations from standard equilibrium predictions (see e.g. Engelmann and Strobel, 2004). We are thankful to a referee for bringing this point to our attention.

effort relationship in our two treatments and minimum effort returned for low wages (i.e. intercepts at e=1). The second type of social preferences models may do so. These intention-based reciprocity models assume that agents are reciprocal and willing to punish unkind actions and/or reward kind ones, even when these punishments and rewards are costly to carry out. In the present context, workers may want to punish low wages and/or reward high ones. The model of Dufwenberg and Kirchsteiger (2004) emphasizes the latter, i.e. *positive* reciprocity. They measure the kindness of a particular wage offer relative to what the employer could give the worker in principle. In both treatments higher wages increase the payoff of the worker and thus will be considered kind. The worker will therefore choose higher effort levels in return. Because there is no reason to interpret a given wage differently in the two treatments, reciprocal workers will act the same; the slope of the wage-effort relationship will be identical in both treatments.[6]

The full model of Charness and Rabin (2002) includes *negative* reciprocity. Starting from quasi-maximin preferences a worker is inclined to sacrifice in order to help the employer. Negative reciprocity enters the model if he is only willing to do so when the employer did not "misbehave". Put differently, workers care about efficiency (social welfare) but withdraw this concern when the employer intentionally offers a low wage. This draws the effort levels provided for low wages downwards in both treatments, as compared to the predictions under quasi-maximin preferences per se. We may then observe an equally steep wage-effort relationship in both treatments, with minimum effort offered for low wages at the same time.

The above discussion makes clear that our experiment is not particularly suitable to discriminate between various models of social preferences. We also did not design our experiment with that goal in mind. Numerous studies have already established the importance (and limits) of the motivational factors included in these models and a number of experiments have been purposely designed to discriminate between them, see e.g. Charness (2004), Cox (2004), Falk et al. (2000) and McCabe et al. (2003). Our interest rather lies in testing the robustness of gift exchange to

[6] Within the context of the gift exchange game the reciprocity theory of Falk and Fischbacher (2006) also stresses positive reciprocity (see their Section 4.2). Like the full model of Charness and Rabin (2002) discussed below, they combine outcome-oriented social preferences and intention-based reciprocity. In particular, they assume that the (un)kindness of a given wage offer depends on both the employer's intentions and on the payoff (in)equality that results. The latter aspect causes that a given wage offer is interpreted differently in the two treatments. As a result, the Falk and Fischbacher model predicts a less steep wage-effort relationship in the 1-4 treatment.

expanding the number of workers per employer. We have discussed some leading models of social preferences in this subsection just to point out that they make somewhat different predictions in this regard.

3. Experimental design and procedures

In Subsection 2.1 we discussed the particular game (and its parameterization) that we study in our experiment. In this section we first discuss some other important features of our experimental design and compare these with earlier gift-exchange experiments. Subsequently, we elaborate on the experimental procedures.

3.1 Experimental design

In general gift exchange experiments come in two forms. The most commonly used one is the gift exchange *market* (GEM), in which employers compete for workers and workers for employers. The wage is determined either in a one-sided oral auction (e.g. Fehr et al., 1993), a one-sided posted offer market (e.g. Brandts and Charness, 2004) or in a double auction (e.g. Fehr and Falk, 1999). GEMs are useful to isolate the role of social norms, because in other (complete contract) market experiments there is typically quick convergence to the competitive equilibrium. The second type is the *bilateral* gift exchange game (BGE). In BGEs the number of employers and workers is the same and matching occurs exogenously. By comparing GEMs with BGEs, the effect of competition on social norms can be isolated. Fehr et al. (1998a) do not find significant differences between the two.[7]

Our experiment is based on the bilateral gift exchange game of Fehr et al. (1998a). Just like them we use a labor market frame, creating in the laboratory a small-scale replica of employer-worker(s) interactions. However, our design differs in a number of respects from theirs. Most importantly, in our experiment the interaction is really one-shot and the strategy method is used.[8] The strategy method implies that

[7] The Spearman rank correlation between wages and effort for each worker is significantly positive at the 10 percent level for 60 percent of all workers in GEM and 57 percent in BGE. Wages in BGE are initially higher than wages in GEM but after a few periods wages in these two treatments coincide.

[8] Our design also differs in other, less important ways. Mainly for practical reasons we do not allow workers to reject the employer's wage offer. This guarantees that in the 1-1 (1-4) treatment each

each worker has to make a contingent effort choice, i.e. each worker has to indicate his effort choice for every possible wage offer. The main advantage of using this method is that it allows us to study the wage-effort relationship of every individual worker. In particular, for every possible wage level, we observe a worker's actual effort choice. At the same time, we also obtain sufficient observations with respect to effort choices for very low and very high wages. In standard gift exchange experiments that do not make use of the strategy method, the bulk of the wage offers made by employers are in the intermediate range. Especially very high wages are hardly observed at all and hence little data is obtained concerning effort returned for the upper wage levels.

The theoretical predictions of Subsection 2.2 are unaffected by our use of the strategy method.[9] It is possible though that subjects in practice do behave differently, because the strategy method forces workers to think about all possible wage offers and not just about the employer's actual wage offer. This may for instance induce them to display (more) monotonic behavior. A number of experimental studies find that the impact of the elicitation method on subjects' choices is limited, especially in situations of low complexity (cf. Bosch-Domènech and Silvestre, 2005, Brandts and Charness, 2000, Cason and Mui, 1998, Falk and Kosfeld, 2005, Oxoby and McLeish, 2004). From this perspective our use of the strategy method seems justified. Some other studies do find that the strategy method leads to different results (cf. Brosig et al., 2003; Güth et al., 2001a). But, even when the strategy method by itself affects behavior, there is a priori no reason to expect that this impact will differ between our two treatments. Treatment comparisons are thus likely to be unaffected. As it appears, our results for the 1-1 treatment are very similar to previous gift exchange experiments that did not use the strategy method (cf. Section 4). Given this, it seems rather implausible that the strategy method would only influence behavior in the 1-4 treatment, and in such a way that it explains our results. In Section 5 we return to this issue when we discuss the robustness of our findings.

Rather than the usual practice of having subjects play the gift exchange game repeatedly in a row, we let them play it only once. We do so to make the two

---

employer always has one worker (four workers). Because in both treatments workers cannot punish low wage offers by rejecting employment, this difference is immaterial for our comparison across treatments.

[9] Note that the subgame perfection criterion implicitly employed in Subsection 2.2 deletes Nash equilibria of the sequential move game in which workers use a weakly dominated strategy.

treatments better comparable. Repeated interaction would lead to different learning and information updating possibilities in the two treatments.[10] Moreover, matching of subjects would also be problematic, because in treatment 1-4 it would be impossible to match subjects only once (without contamination) when more than just a couple of rounds are played. Reputational effects then cannot be avoided. A single round has the additional advantage that marginal incentives are particularly salient.

A final aspect of our design that deserves brief discussion is the employer's profit function. The typical gift exchange experiment calculates the employer's profit as $(v-w)\cdot e$, whereas in our design the profit per worker equals $v\cdot e-w$. The former specification is typically used to prevent that loss aversion pollutes fairness effects. For our profit function the employer makes a loss when she pays a high wage and a worker chooses a low effort in response. According to standard theory the choice between the two profit functions is immaterial, but under alternative motives the type of profit function used does matter (see below). Fehr et al. (1998a) compare a treatment where they completely avoid losses (i.e. profit equals $(v-w)\cdot e$) with another one in which the employer can make losses (profit equals $v\cdot e-w$). They observe that the pattern of reciprocal interactions is basically the same in both treatments.

In our view a profit function equal to $v\cdot e-w$ is more appropriate. With this specification the surplus per worker equals $v\cdot e-c(e)$ and only depends on the effort level chosen. The wage itself simply acts as a transfer payment. In those papers that use the multiplicative specification $(v-w)\cdot e$, surplus equals $v\cdot e+w\cdot(1-e)-c(e)$ and is increasing in w for a given level of effort.[11] Employers may then choose a higher wage just to increase the surplus, rather than to elicit higher effort levels from workers. A high wage policy is then more difficult to interpret. Another disadvantage is that the multiplicative profit function hinders the effectiveness of high effort as a reciprocation device: the value of an additional unit of effort is substantially lower in the upper wage ranges (i.e., the effective multiplier of effort decreases with the wage). Choosing a higher effort level at a higher wage then does not necessarily indicate that

[10] This holds because informing a worker about his payoffs in treatment 1-1, enables him to infer the payoffs of the employer precisely. When we then also inform workers in treatment 1-4 about employer's payoffs, these workers can infer the average effort level chosen by others within the group.

[11] This holds because in these papers $e \leq 1$. For ease of presentation to the subjects we have scaled effort levels upwards (and marginal revenue v proportionally downwards), such that effort can take integer values only.

the worker provides a larger gift in return. For the profit function that we use the interpretation of higher effort levels is more clear-cut.

To avoid negative payoffs, we provided all subjects with a sufficiently high initial endowment (besides a show up fee). This explains our choice of $\omega=90$ (cf. Subsection 2.1).

3.2 Experimental procedures

Overall we conducted 7 sessions. Four of them considered treatment 1-1 (84 participants), the other three sessions concerned treatment 1-4 (60 participants). Participation was restricted to one session. The subject pool consisted mainly of the undergraduate student population of the University of Amsterdam. Most of them were students at the Department of Economics and Econometrics. 58% of the participants were male. The average earnings were 17,09 euros (including a show-up fee of 5 euros) for about one hour on average.

The experiment was computerized using z-Tree (Fischbacher, 1999). Subjects started with on-screen instructions. They also received a summary of the instructions on paper (see Appendix B). To ensure that subjects understood the experiment, in particular how their payoffs were calculated, all subjects had to answer a number of control questions correctly before the experiment started.[12] Then, at the start of the experiment, subjects learned their roles. In treatment 1-1, half of the participants were assigned the role of employer, the other half the role of worker. Both of them received an initial endowment of 90 points (besides the show up fee). In treatment 1-4, 20% of the participants had the role of employer, the other 80% the role of worker. Here employers received an endowment of 360 points, while workers got 90 points. The conversion rate was 1 euro for 9 points.

In both treatments subjects were randomly and anonymously matched into firms. In treatment 1-1 a firm was composed of one employer and one worker, in treatment 1-4 a firm consisted of one employer and four workers. The employers had to set the wage without knowing the effort choices made by their worker(s). Recall

[12] Subjects generated the numerical examples for the control questions themselves, by first making hypothetical choices for both the employer's and the worker's role. The control questions then asked them to calculate the payoffs. By using this procedure, subjects also got familiar with the decision task they would face in the actual experiment (although there they would take decisions for one of the two roles only).

that wages were restricted to multiples of 5 between 0 and 100. Workers had to fill in a wage-effort table without knowing the actual wage set by their employer (and the effort levels chosen by their co-workers). In particular, for every possible wage they had to indicate their effort choice, an integer between 1 and 10. When all subjects had made their decisions, workers were informed about the actual wage offered by the employer, and their own and their employer's payoffs. Employers learned the effort choices of their workers for the wage chosen. No subject was ever informed about the choices of subjects in other firms. Finally, each participant learned his/her own earnings and was paid accordingly (after filling in a short ex post questionnaire).

4. Results

In this section we present the findings of our experiment. Because our main interest lies in the robustness of workers' willingness to reciprocate, we will primarily focus on their effort choices. After that we briefly turn to employer behavior and overall earnings. Recall that any influence of employers' *actual* wage decisions on the conditional effort choices of workers is ruled out by our one-shot strategy design.

4.1 Worker behavior

Standard theory based on selfish preferences predicts that workers will choose the minimum effort level of one for every wage amount offered. In clear contrast to this, the actual mean average effort of the workers over all possible wages is 3.81 in treatment 1-1 and 3.31 in treatment 1-4. Moreover, the mean efforts increase with the wage level: for wages below 50 the average effort equals 2.27 in treatment 1-1 and 2.01 in treatment 1-4, while for wages of 50 and higher these numbers increase to 5.21 and 4.46 respectively.

More detailed information about the relationship between wages and effort is provided in Figure 1. This figure displays the average effort by wage for each of the two treatments separately. (It also depicts the wage-effort relationship for an additional treatment that will be discussed in Section 5.) In both treatments a clear pattern emerges: workers on average choose a higher effort level when the wage is higher. The Spearman rank correlation between effort and wages is highly significant

and around 0.45 in both treatments (0.50 and 0.42 in treatments 1-1 and 1-4 respectively; more detailed regression results will be discussed below). We summarize this finding in the following result:

Result 1. In both treatments, mean effort levels are increasing in the wage.

Although our one-shot strategy design differs from previous gift-exchange experiments (cf. Subsection 3.1), for the 1-1 treatment Result 1 replicates earlier findings reported in the literature (Fehr and Gächter, 2000). This also holds for our findings concerning employer behavior and earnings, which will be discussed in the next subsection. For the 1-1 treatment our results are thus comparable to those of previous studies.

Figure 1: Average effort by wage

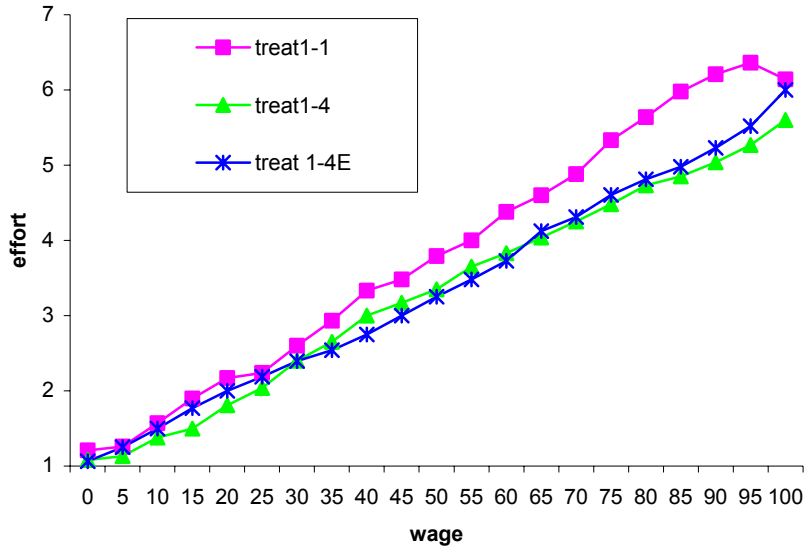

Our primary interest lies in comparing the two treatments. Figure 1 indicates that average conditional effort in the 1-1 treatment is higher than in the 1-4 treatment for all 21 wage levels. The difference between treatments also increases as the wage increases (except for a wage of 100). However, the differences are small and insignificant. Nonparametric (Mann-Whitney) tests for each wage separately show no statistical significance at the 5% level. The same conclusion can be drawn when we compare *individual mean effort levels* across treatments. Comparing these individual

means, no significant differences are found between the two treatments (two-sided Mann-Whitney test; p=0.3).

Non-parametric tests are based on order statistics only and thus neglect some of the information in the data. Regression analysis does not have this disadvantage. We therefore also consider the results from regressing effort on the wage level, a treatment dummy and an interaction term.[13] To account for the panel structure of our data (we have 21 data points per individual worker), a random effects model is estimated. Table 2 shows the results. The intercepts in both treatments are equal and figure around one, but the slope is less steep in the 1-4 treatment. This is of course much like Figure 1. The difference in slope is small but statistically significant. Increasing the wage by 10 points leads to an estimated increase in effort of around 0.55 in the 1-1 treatment and of around 0.46 in the 1-4 treatment.

Table 2. Random effects linear regression with effort as dependent variable

Independent variable	*coefficient*
Constant	1.032***
	(0.3340)
Wage	0.0556***
	(0.0018)
Treatment 1-4	-0.0559
	(0.4654)
Treatment 1-4 * Wage	-0.0091***
	(0.0025)

Remark: Robust standard errors appear in parentheses. *** Indicates significance at the 1%-level. There are N=1890 observations in total, with 90 clusters (workers) and 21 observations per cluster. The Wald statistic equals χ^2=1653.03; p=0.000.

Result 2. The wage-effort relationship is significantly less steep in the 1-4 treatment than in the 1-1 treatment. Differences in effort levels are small though.

[13] Incorporating wage squared as an additional explanatory variable does not yield a significant coefficient at the 5% level and hardly affects the estimates of the other coefficients. Table 2 therefore only presents the simpler specification.

The analysis up till now is based on mean conditional effort levels per treatment. Clearly, these average efforts may conceal large individual differences between subjects within a treatment. We therefore explore individual worker behavior in more detail. Like in other gift exchange studies, substantial heterogeneity among subjects is observed in the data. The following general picture emerges. A number of workers can be considered selfish; they choose the minimum effort level regardless of the wage. A second group acts reciprocally; these workers increase their effort with wages in a monotonic way. A small third group can be labelled erratic; they vary their effort, but with no clear relation to the wage.

To arrive at our exact classification, we fit a linear effort function $e_i = a + b \cdot w_i$ for each individual subject. The estimated wage coefficient is then used together with the observed monotonicity of the relationship to classify the individual workers. The precise definition of our three different categories is as follows:

Selfish behavior: This group consists of workers that provide the minimum effort level for all possible wages.

Reciprocal behavior: These subjects choose a monotonic and increasing wage-effort schedule. We also included two subjects with a monotonic wage-effort schedule for all wages, except for the highest and the two highest wages respectively. Workers within this category have a significantly positive wage coefficient; $b = \partial e/\partial w > 0$.

Erratic behavior: These subjects do not make monotonic effort choices. Possibly these subjects made some typing mistakes, did not understand the experiment, or were not taking it seriously (or only filled in a serious effort level for wages they considered plausible).

Table 3 shows the classification of subjects in these three categories by treatment. In treatment 1-4 there appear to be more selfish workers than in treatment 1-1. There are two possible causes for this finding. First, it can represent a treatment effect; some workers who would have acted reciprocal in the 1-1 treatment are acting selfish in the 1-4 treatment. Second, the difference may be caused by chance. The latter explanation cannot be ruled out, because the difference in the proportion of selfish players

between treatments appears insignificant using a binominal test (although it must be noted that this test is not very powerful with these low numbers).[14]

Table 3. Classification of workers by treatment and average contribution per worker type

Type of behavior	*Treatment 1-1*		*Treatment 1-4*	
	frequencies	*average contribution*	*frequencies*	*average contribution*
Selfish	23.8% (10)	1.00	33.3% (16)	1.00
Reciprocal	64.3% (27)	4.75	60.4% (29)	4.61
Erratic	11.9% (5)	4.35	6.3% (3)	2.83
Total	100% (42)	3.81 (42)	100% (48)	3.31 (48)
Non-erratic workers	88.1% (37)	3.74 (37)	93.8% (45)	3.33 (45)

Remark: Number of individual workers in parentheses

Result 3. The percentage of selfish workers is higher in treatment 1-4 than in treatment 1-1, but not significantly so.

The fact that there are less selfish workers in treatment 1-1 provides one explanation for our finding that the wage-effort relationship is slightly steeper in this treatment (cf. Result 2). Another explanation would be that the reciprocal workers behave less reciprocally in treatment 1-4. To explore the latter possibility, Table 3 also compares average efforts of the different worker types across treatments. It appears that the reciprocal workers make the same average efforts in both treatments. The same conclusion follows from estimating a random effects model like in Table 2, but now restricted to the reciprocal types only. For this subsample the wage-effort relationship is not significantly different in the 1-4 treatment. The (small) differences observed in Figure 1 can thus be accounted for by the fact that fewer subjects act reciprocal in

[14] Also when we compare the complete type distribution (i.e. consisting of the three categories 'Selfish', 'Reciprocal' and 'Erratic') between treatments by means of a chi-square test, no significant differences are found (p=0.46).

treatment 1-4, rather than by reciprocal subjects acting less reciprocal in the 1-4 treatment. A possible explanation for this finding is that some workers may be willing to free ride on their fellow workers' reciprocal behavior, because they do not feel entirely responsible for the employer's earnings.[15] Charness (2000) finds experimental evidence for this type of *responsibility-alleviation* effect within the context of a bilateral gift exchange game.

Summing up, Results 1 through 3 clearly show that a large fraction of workers have motivations going beyond material self interest. Effort levels are typically above minimum and we find a positive correlation between effort and wages. Although the relationship is somewhat steeper in treatment 1-1 as compared to treatment 1-4, the difference is much smaller than one would expect on the basis of inequality-aversion preferences (cf. Subsection 2.2 and Appendix A). Overall the wage-effort relationship appears to be quite robust to increasing the number of workers within a firm.

Our results can neither be explained by efficiency-based preferences (like quasi-maximin) *per se*. Although concerns for efficiency make the predictions more similar across treatments, they also increase the lowest effort the worker supplies at the same time. In particular, subjects with a strong enough preference for efficiency are predicted to supply higher than minimum effort even for very low wages. This is not what we observe in Figure 1. Also the estimated intercepts in Table 2 do not differ significantly from one (i.e. the minimum effort) in both treatments.

The joint observation of an almost equally steep wage-effort relationship in the two treatments and minimum effort returned for low wages suggests that intention based reciprocity plays a major role in gift exchange. This interpretation is in line with the findings reported in Charness (2004). He compares the standard bilateral gift exchange game with a treatment in which the wage is determined exogenously (either by a random draw or by the experimenter). In the latter case the wage-effort relationship has a significantly lower slope and a substantially higher intercept. This unambiguously reveals that intentions matter for gift exchange. Charness estimates that about 35% of the slope in the standard gift exchange game is a direct consequence of intention-based reciprocity. His results also tentatively suggest that negative reciprocity in response to intentionally low wages is more important than positive reciprocity in response to intentionally high wages is.

[15] We are thankful to an anonymous referee for suggesting this possible explanation.

As already noted in Subsection 2.2, our experiment is neither very suitable for identifying the different models of social preferences nor does it intend to do so. Our results are in line with both the importance of positive reciprocity as stressed by the model of Dufwenberg and Kirchsteiger (2004) and Falk and Fischbacher (2006), and with efficiency-based preferences together with negative reciprocity (concern withdrawal) as emphasized by Charness and Rabin (2002). Although the results of Charness (2004) discussed above suggest that the latter is the more likely explanation, the main point we want to make here is that reciprocity – being it either negative, positive or both – plays an important role in gift exchange. Outcome oriented social preferences per se are insufficient to explain gift exchange behavior as we observe it.

4.2 Employer behavior and earnings

Although workers' effort choices are the main focus of this study, it is interesting to compare the wages offered by the employers in the two treatments. Standard theory predicts that employers should pay the minimum wage in both treatments. Inequality-averse employers or employers that expect favorable responses (and thus efficiency gains) may pay more than the minimum wage. In treatment 1-4 the initial endowment and the potential earnings of the employer are four times as large as in treatment 1-1. Inequality-averse employers will therefore offer higher wages in treatment 1-4 than in treatment 1-1 (assuming that workers behave similar in the two treatments). In general, the 'optimal' wage level for the employer depends on the actual wage-effort schedules of the workers. Averaging over these actual schedules we observe that in both treatments the employers' potential earnings are monotonically decreasing in the wage offered. Even though a higher average effort is provided for higher wages, the higher returns do not outweigh the additional wage costs.[16] So, on average it is not profitable for employers to offer higher than minimum wages.

The average wage offered in treatment in 1-4 is higher (41.67, n=12) than the mean wage in treatment 1-1 (33.21, n=42). However, this difference is far from statistically significant (two sided Mann-Whitney test; p=0.221). Moreover, the percentage of employers that offers the minimum wage is slightly higher in treatment

[16] The same conclusion follows from using the estimated wage coefficients in Table 2. A 10-points increase in the wage then leads to a decrease in net profits of -4.5 points in treatment 1-1 and of 4 times –5.4 points (= -21.6 points) in treatment 1-4.

1-4 (16.7%) than in treatment 1-1 (14.3%). We thus do not find strong evidence that the employer offers higher wages when she has more workers.

We next turn to the actual earnings subjects obtain in the experiment. Table 4 reports the subjects' average earnings by treatment and by role, excluding the show up fee of €5 every participant received for showing up on time. Workers earn a little bit more in the 1-4 treatment than in the 1-1 treatment. This is in line with the marginally higher wages that are observed in the 1-4 treatment, together with the somewhat lower effort levels chosen (cf. Result 2).

Table 4. Average earnings by treatment (excluding show up fee of €5)

		Treatment 1-1	*Treatment 1-4*
Workers	Endowment	90	90
	Profit	29.4	38.48
	Total	119.4 (€13.26)	128.48 ((€14.28)
Employers	Endowment	90	360
	Profit	-0.36	-47.5
	Total	89.64 (€9.96)	312.5 (€34.72)

Remark: Earnings in points. The conversion rate is 9 points = 1 euro.

Employers earned much more in the 1-4 treatment than in the 1-1 treatment. The main cause here is that in the 1-4 treatment the initial endowment is four times as large as in the 1-1 treatment. Ignoring the initial endowment, the average profit the employer makes is actually much smaller in the presence of multiple workers (they then make a loss on average). Again, the driving force for this observation is twofold. In treatment 1-4 workers provide less effort on average and employers also pay (marginally) higher wages than in treatment 1-1. As a result, employers with four workers earn less than four times the amount employers with only one worker earn. In practice treatment 1-4 is thus not simply a four times upscaling of treatment 1-1.

Comparing average earnings across roles, in the 1-1 treatment workers on average earn more than employers do. This is a common finding in gift-exchange experiments. In treatment 1-4, however, employers earn substantially more than workers. Also in this respect the 1-4 treatment seems a better representation of actual worker-employer relationships.

5. Robustness

The somewhat surprising finding that gift-exchange is robust to enlarging the size of the workforce raises some questions that deserve further analysis. An important feature of our design is the inequality between employers and workers created by the different initial endowments they received in the 1-4 treatment. Whereas in the 1-1 treatment both the worker and the employer received the same initial endowment, in the 1-4 treatment employers received four times as much. This difference drives the predictions of the inequality-aversion model (cf. Subsection 2.2).

An important question therefore is how subjects perceived the initial endowments. As can be seen from the summary of the instructions handed out to the subjects in the 1-4 treatment (cf. Appendix B), the payoff functions did not include the initial endowment. Despite the fact that subjects were carefully informed on the initial endowment for both roles, they may have ignored this in their considerations. It is also possible that they believed that the higher endowment of the employer was needed to pay for the wages (although this is only the case for very high wages). If the initial endowments are ignored, the inequality-aversion model predicts the same results for both treatments (in line with what we observe).[17]

To address this issue we ran an additional 1-4 treatment where we made very explicit the differences in initial endowments between roles. In this 1-4E treatment the payoff functions explicitly included the initial endowment (see the summary of the instructions in Appendix B). Subjects used these functions both for the calculations they were ask to make in the control questions and during the experiment. Hence the (large) inequality created by the initial endowments was made very salient to them. For the 1-4E treatment we conducted 3 extra sessions, with 60 participants in total.

[17] We are grateful to a referee for this suggestion, which inspired the analysis in this section.

Like in the other two treatments participation was restricted to one session. Moreover, we ensured that participants did not have participated in any of the previous treatments. We decided not to run additional sessions for the 1-1 treatment, because there employers and workers get exactly the same initial endowment and there is no inequality that needs to be stressed.

Figure 1 in Section 4.1 depicts the average effort by wage for all treatments. There appears to be no difference between the treatments 1-4 and 1-4E; Mann-Whitney ranksum tests for each wage separately show no statistical significance at the 5% level. The same conclusion is obtained when we incorporate the data from treatment 1-4E in the regression of Table 2 (cf. Section 4.1). The intercept for the 1-4E treatment appears insignificant (as it is for treatment 1-4) whereas the interaction term with the wage is significant and equal to –0.0078. Hence, the difference in slopes between treatment 1-1 and treatment 1-4E is of the same magnitude as the difference in slopes between treatments 1-1 and 1-4.

Another feature of treatment 1-4E is that we also measured workers' expectations about the effort levels of their fellow workers. When discussing the theoretical predictions in Subsection 2.2 we assumed that the workers would expect the other workers to behave similar. In order to check this assumption we asked workers to fill in another wage-effort table, but now with the average effort they expected the other three workers in the firm had provided. Workers filled in this table *after* they had made their own effort choices but *before* they were informed about the wage they received. In that way their own actual effort choices could not be affected by asking for their expectations and the expectations themselves could not be based on observed outcomes (i.e. on updated information). We also asked employers (right after they had chosen the wage) to fill in a wage-effort table with the average effort they expected from their own employees.[18]

On average workers expected the other workers to provide the same, or a little more, effort than they did so themselves. The difference between own effort and expectations is significant only for some of the higher wages: 35, 55, 60, 80, 85 and 90 (two-sided signrank test). This provides some evidence for the suggestion made in Subsection 4.1 that workers may be willing to free ride on their co-workers' reciprocal behavior. Overall, however, the stated expectations reveal that, as a first

[18] We did not reward the expectations by a proper scoring rule, because we wanted the financial incentives in treatment 1-4E to be exactly the same as in treatment 1-4.

approximation, the assumption that workers expect their fellow workers to behave similar is not unreasonable. In Appendix A we use the measured expectations to derive more accurate predictions under both inequality-aversion and quasi-maximin preferences. The general conclusion that follows from this analysis is that the intuitive predictions of Subsection 2.2 need not be adapted.

Employers made on average correct predictions about the average effort for each possible wage (we find no significant differences at the 5% level using a two-sided Mann-Whitney test). Recall from Table 4 that employers on average made a loss. The reported expectations imply that the employers who did offer a higher than minimum wage, did (on average) expect this loss.

Summing up, the analysis of the extra 1-4E treatment strengthens our previous results. In treatment 1-4E workers were made fully aware of the (large) inequality differences between them and their employers. Still they did not substantially change their attitudes towards gift-exchange, as inequality-aversion models would predict.[19] Moreover, if we take into account workers' actual expectations about their co-worker's effort, we also strengthen our previous observation that intention-based reciprocity is an important driving force behind gift exchange.

Obviously other dimensions of robustness remain. The relatively large endowments that we use (needed to cover potential losses of employers) may have made marginal incentives less salient than in previous studies. Yet this is partly counterbalanced by the fact that we let subjects play the game only once rather than repeatedly in a row. The similarity of the results for the 1-1 treatment with previous bilateral gift exchange experiments in fact suggests that our findings will be quite robust to lowering endowment levels.

Finally, as discussed in Subsection 3.1 a major difference between our gift exchange experiment and previous ones is that we make use of the strategy method. At the aggregate level this has the advantage that we obtain a sufficient number of effort choice observations for the full range of feasible wages. At the individual level it allows us to observe the wage-effort relationship of each worker. Although a number of experimental studies suggest that employing the strategy method is innocuous, other studies question this conclusion. One reassuring fact here is again that our results for the 1-1 treatment are consistent with those obtained by others

[19] Anderson et al. (2005) similarly find that inequality per se (induced by different show-up fees) does not significantly affect trusting behavior and trustworthiness in the well-known trust game.

without using the strategy method. Hence the strategy method per se does not make workers more (or less) reciprocal. A priori there is thus also no reason to expect this to happen in the 1-4 treatment, so it seems rather implausible that the strategy method itself is the driving force behind our finding that gift exchange is robust to expanding the number of workers per employer.

Nevertheless, it seems reasonable to conjecture that the strategy method induces workers to make effort choices that are more monotonic. Compare, for example, the wage-effort relationship we observe in treatment 1-1 (cf. Figure 1) with the one obtained by Fehr et al (1993) who did not employ the strategy method. In the latter case average observed effort is less monotonic in the wage offered (see their Figure I). Also the percentage of erratic workers we find in our two treatments is rather low (cf. Table 3). Having workers react directly to actual wage offers would likely lead to more erratic behavior and a less monotonic graph than the one seen in Figure 1. Exactly how robust our findings are to the elicitation method used is thus an open question.

6. Conclusion

In this paper we explore whether experimental gift exchange is robust to enlarging the size of the workforce. To that purpose we compare a baseline one-employer – one-worker treatment with another treatment in which each employer has four workers. Although our design is different from earlier studies, most notably in the use of the strategy method and the play of just a single round, the results for our 1-1 treatment are in line with previous findings (cf. Fehr and Gächter, 2000). In particular, we also find that workers on average provide more than minimum effort and that around 60% of them have an upward sloping wage-effort relationship. In the baseline treatment workers earn more than employers, as is typically obtained in bilateral gift exchange games. The results of the 1-4 treatment show that the findings on gift exchange are quite robust. Effort levels are only marginally lower than in treatment 1-1, despite the fact that the employer now earns much more than the worker does. This in turn suggests that intention-based reciprocity is an important driving force behind gift

exchange; like in Charness (2004, p. 669), outcome-oriented social preferences alone are insufficient to explain gift exchange observed in our experiment.

Our results are somewhat surprising. A priori we expected that workers would be much less willing to provide higher than minimum effort in treatment 1-4 than in treatment 1-1, because of the large payoff inequality between the employer and her workers in the former treatment. This inequality is in large part due to the asymmetric initial endowments subjects with different roles got. To exclude that subjects simply may have ignored these endowments, we ran another one employer - four workers treatment in which the differences in endowments between employer and worker is emphasized. Results are identical to our original 1-4 treatment. The robustness of our findings with respect to the elicitation method remains an open question though.

In contrast to other recent experiments that focus on horizontal fairness between workers, this paper is mainly concerned with fairness and reciprocity considerations between agents at different levels in the hierarchy. Given that the gift exchange model seems to survive the more realistic 'multi-worker' design, it is also worthwhile to investigate robustness in other, related directions. A first extension that comes to mind is towards a multi-level hierarchy. In larger firms the hierarchical structure is typically quite complex. Apart from shareholders there are usually multiple managerial layers. The person who is responsible for a worker's wage is typically not full residual claimant and thus does not get the full benefits generated by the worker's reciprocal reaction (if at all). Under the assumptions of selfishness and rationality this does not matter and a focus on a simple principal-agent relationship is justified. This is not the case anymore when alternative motives come into play.

Another issue that seems worth investigating in the context of the gift exchange model is the endogenous selection of employers. In a typical experiment (like ours) participants are students who want to earn some money. They are randomly assigned to act either as a worker or as an employer. A positive effort level is not very costly for a worker, while the benefits for the employer can be huge. A student in the worker's role may therefore be inclined to help the fellow student playing the employer (who behaves nice), also because the roles could easily have been the other way around. This kind of Kantian or Rawlesian reasoning is less likely in the real world, because being an employer or a worker is to a large extent determined by background variables and/or personal choice and not by lot.

References

Anderson, L.R., Mellor, J.M. and J. Milyo, 2005, An experimental study of the effects of inequality and relative deprivation on trusting behavior, Working paper nr. 14, College of William and Mary, Department of Economics.

Akerlof, G.A., 1982, Labor contracts as partial gift-exchange, *Quarterly Journal of Economics* 97, 543-69.

Baker, G., Jensen, M.C. and K.J. Murphy, 1988, Compensation and incentives: Practice vs. theory, *The Journal of Finance* 43, 593-616.

Bolton, G. and A. Ockenfels, 2000, ERC: A theory of equity, reciprocity and competition, *American Economic Review* 90, 166-193.

Bosch-Domènech, and J. Silvestre, 2005, Elicitation in the lab: Does embedding matter?, Working Paper.

Brandts, J. and G. Charness, 2000, Hot vs. cold: Sequential responses and preference stability in experimental games, *Experimental Economics* 2, 227-238.

Brandts, J. and G. Charness, 2004, Do labour market conditions affect gift exchange? Some experimental evidence, *Economic Journal* 114, 684-708.

Brosig, J., Weimann, J. and C.-L. Yang, 2003, The hot versus cold effect in a simple bargaining experiment, *Experimental Economics 6,* 75-90.

Camerer, C.F., 2003, *Behavorial game theory: Experiments in strategic interaction*, Princeton University Press, Princeton, New Jersey.

Cason, T.N. and V.-L. Mui, 1998, Social influence in the sequential dictator game, *Journal of Mathematical Psychology* 42, 248-265.

Charness, G., 2000, Responsibility and effort in an experimental labor market, *Journal of Economic Behavior and Organization* 42, 375-384.

Charness, G., 2004, Attribution and reciprocity in an experimental labor market, *Journal of Labor Economics* 22, 665-688.

Charness, G. and P. Kuhn, 2004, Do workers' wages matter? Theory and evidence on wage secrecy, wage compression and effort, Working Paper, University of California.

Charness, G. and M. Rabin, 2002, Understanding social preferences with simple tests, *Quarterly Journal of Economics* 117, 817-869.

Charness, G., Frechette, G.R. and J.H. Kagel, 2004, How robust is laboratory gift exchange?, *Experimental Economics* 7, 189-205.

Cox, J., 2004, How to identify trust and reciprocity, *Games and Economic Behavior* 46, 260-281.

Dufwenberg, M. and G. Kirchsteiger, 2004, A theory of sequential reciprocity, *Games and Economic Behavior* 47, 268-298.

Engelmann, D. and A. Ortmann, 2002, The Robustness of gift exchange: A reconsideration, Working Paper.

Engelmann, D. and M. Strobel, 2004, Inequality aversion, efficiency, and maximin preferences in simple distribution experiments, *American Economic Review* 94, 857-869.

Falk, A., Fehr, E., and U. Fischbacher, 2000, Testing theories of fairness – Intentions matter, Working Paper, University of Zurich.

Falk, A. and U. Fischbacher, 2006, A theory of reciprocity, *Games and Economic Behavior* 54, 293-315.

Falk, A. and M. Kosfeld, 2005, The hidden costs of control, *American Economic Review*, forthcoming.

Fehr, E. and A. Falk, 1999, Wage rigidity in a competitive incomplete contract market. *Journal of Political Economy* 107, 106-34

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

Fehr, E. and K. Schmidt, 1999, A Theory of Fairness, Competition, and Cooperation *Quarterly Journal of Economics* 114, 817-868.

Fehr, E, Fischbacher, U. and E. Tougareva, 2002, Do high stakes and competition undermine fairness? Evidence from Russia, Working Paper no. 120, University of Zurich.

Fehr, E., Gächter, S. and G. Kirchsteiger, 1996a, Reciprocal fairness and noncompensating wage differentials, *Journal of Institutional and Theoretical Economics* 152, 608-640.

Fehr, E., Kirchler, E., Weichbold, A. and S. Gächter, 1998a, When social norms overpower competition: Gift exchange in experimental labor markets, *Journal of Labor Economics* 16, 324-351.

Fehr, E., Kirchsteiger, G. and A. Riedl, 1993, Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics* 108, 437-59.

Fehr, E., Kirchsteiger, G. and A. Riedl, 1996b, Involuntary unemployment and non-compensating wage differentials in an experimental labour market. *Economic Journal* 106, 106-21.

Fehr, E., Kirchsteiger, G. and A. Riedl, 1998b, Gift exchange and reciprocity in competitive experimental markets. *European Economic* Review 42, 1-34.

Fischbacher, U., 1999, z-Tree. Zurich toolbox for readymade economic experiments, Working Paper No. 21, University of Zurich.

Güth, W., Huck, S. and W. Müller, 2001a, The relevance of equal splits in ultimatum games, *Games and Economic Behavior* 37, 161-169.

Güth, W., Königstein, M., Kovács, J. and E. Zala-Mezõ, 2001b, Fairness within firms: the case of one principal and many agents, *Schmalenbach Business Review* 53, 82-101.

Hannan, R.L., Kagel, J.H. and D.V. Moser, 2002, Partial gift exchange in an experimental labor market: Impact of subject population differences, productivity differences and effort requests on behavior, *Journal of Labor Economics* 20, 923-951.

Loewenstein, G., 1999, Experimental economics from the vantage-point of behavioural economics, *Economic Journal* 109, F25-F34.

McCabe, K., Rigdon, M. and V. Smith, 2003, Positive reciprocity and intentions in trust games, *Journal of Economic Behavior and Organization* 52, 267-275.

Meidinger, C., Rullière, J. and M. Villeval, 2000, Free-riding and fairness in principal-multi-agent relationships: Experimental evidence, *Rivista Internazionale di Scienze Sociali* 153, 409-432.

Meidinger, C., Rullière, J. and M. Villeval, 2003, Does team-based compensation give rise to problems when agents vary in their ability?, *Experimental Economics* 6, 253-272.

Oxoby, R. and K. McLeish, 2004, Sequential decision strategy vector methods in ultimatum bargaining: Evidence on the strength of other-regarding behavior, *Economic Letters* 84, 399-405.

Schram, A., 2005, Artificiality: The tension between internal and external validity in economic experiments, *Journal of Economic Methodology* 12, 225-237.

Warglien, M. and A. Rossi, 2001, An experimental investigation of fairness and reciprocal behavior in a triangular principal-multiagent relationship, Discussion paper, University of Trento.

Appendix A: Formal analysis of social preferences

In this appendix we provide a formal derivation of (some of) the predictions discussed in Subsection 2.2. First assume that agents are motivated by inequality-aversion as in Fehr and Schmidt (1999).[20] For a general (k+1)-player game with monetary payoffs $m=(m_0,..,m_k)$, player i's utility is then given by (for i=0,1,..,k):

$$U_i(m) = m_i - \frac{1}{k}\alpha_i \sum_{j \neq i} \max\{m_j - m_i, 0\} - \frac{1}{k}\beta_i \sum_{j \neq i} \max\{m_i - m_j, 0\} \qquad \text{(A1)}$$

where parameter α_i measures the extent to which player i dislikes disadvantageous inequality and β_i the extent to which s/he dislikes earning more than others. Like Fehr and Schmidt we impose the restrictions $\alpha_i \geq \beta_i$ and $0 \leq \beta_i < 1$. In the 1-1 treatment we have k=1 and the utility of the worker (player 1) reduces to:

$$U_1(m) = \begin{cases} (1+\alpha_1)\cdot m_1 - \alpha_1 \cdot m_0 & \text{if} \quad m_1 \leq m_0 \\ (1-\beta_1)\cdot m_1 + \beta_1 \cdot m_0 & \text{if} \quad m_1 > m_0 \end{cases} \qquad \text{(A2)}$$

Given a particular wage w, choosing more than minimal effort e>1 decreases the worker's payoffs $m_1 = w - c(e) + \omega$ and increases the employer's payoffs $m_0 = ve - w + \omega$. Clearly, with preferences like in (A2) a worker is never willing to do so when he earns less than the employer. He is only willing to choose more than minimal effort when w is sufficiently high, such that he is ahead of the employer, and β_1 is large enough. The value of β_1 is thus decisive for the effort choice in response to a particular wage (the value of α_1 only has a marginal influence).[21] Figure A1 shows the model's predictions for β_1 equal to 0, 0.25 and 0.6 respectively (and $\alpha_1=1$).[22] Workers will always exert minimum effort when the wage is low and those workers with $\beta_1>0$ will exert larger effort when wages are higher.

[20] See Bolton and Ockenfels for a related specification.

[21] When m_0 and m_1 are continuous variables the predictions do not depend on the exact size of α_1 at all. However, in the experiment effort levels and wages are restricted to particular discrete values (cf. Subsection 2.1), and so are m_0 and m_1. This introduces a marginal influence of α_1 on the predictions.

[22] The chosen values of β follow the type distribution inferred by Fehr and Schmidt (1999, Table 3 on p. 844) from a large body of ultimatum game data.

Figure A1: Predictions of the Fehr-Schmidt model for different betas (and α_1=1).

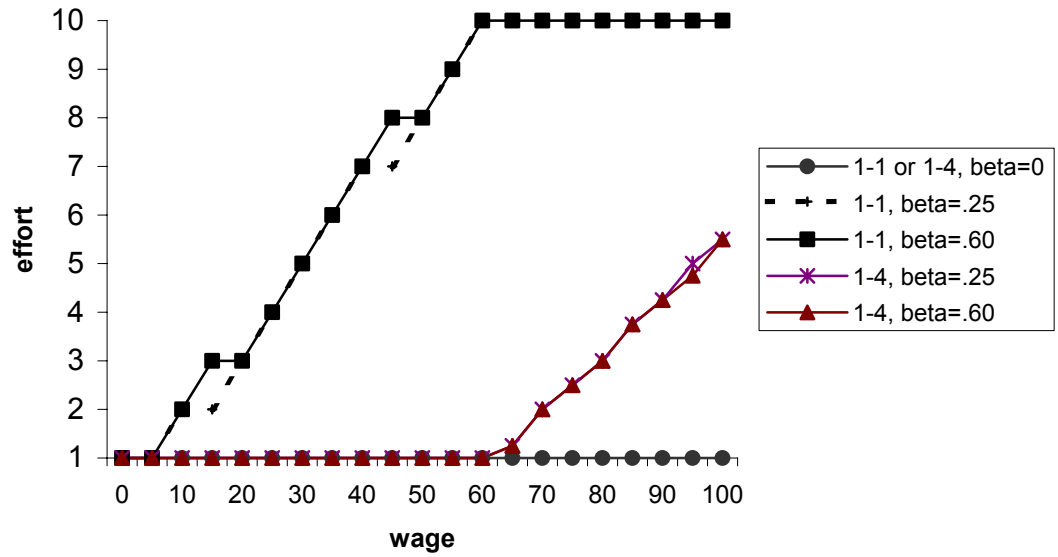

Next consider the 1-4 treatment. In order to make predictions we now have to make additional assumptions about how individual workers perceive the game. First assume that: (A) a worker is only interested in her own earnings and those of the employer. Workers' preferences are then given by expression (A2). In order to consider the overall earnings of her employer m_0, the worker now has to predict the behavior of his co-workers. Because all workers are in exactly the same position it seems natural to assume that: (B) each worker assumes that his co-workers have the same preferences and will behave similar. For any given wage we thus focus on the symmetric equilibria of a game among homogenous workers. Because of the discreteness of workers' effort choices, such symmetric equilibria do not always exist. In these cases we consider the asymmetric equilibria with the lowest variance in workers' effort choices and we take the average effort as equilibrium prediction.

Under assumptions A and B the analysis is similar as for the 1-1 treatment, with the exception that the employer's payoffs are multiplied by four. She therefore earns less than her workers only if she pays a relatively high wage. An inequality-averse worker will therefore choose higher than minimum effort only for wages of 65 and above (see Figure A1). As a result, for such a worker there is a big difference in the wage-effort relationship between treatment 1-1 and treatment 1-4.

If we relax assumptions A and B, these predictions change only a little. First we relax assumption A and maintain B: each worker now perceives the situation as a 5-person game and also cares about his fellow workers, but still assumes that all of

them behave similar. Workers' earnings will then be equal and inequality-aversion only matters for the comparison between a worker himself and his employer. Effectively, an individual worker's preferences are given by (A2) when we multiply α_1 and β_1 by ¼ (=1/k). Effort levels will be marginally lower than those depicted in Figure A1 and the difference with treatment 1-1 will remain at least as large.

To relax also assumption B we make use of the workers' actual expectations as measured in treatment 1-4E (cf. Section 5).[23] We first derive the predictions for each individual worker separately based on his/her stated expectations and then take the average over all individual predictions. The following picture emerges. When a worker perceives the game as a two-player game, he will provide the minimum effort for all wages below 80. For higher wages effort may be higher, but it never exceeds 4. If a worker perceives the game as a 5-player game, he provides the minimum effort for w≤75 and e=2 otherwise. Compared to the predictions in Figure A1, a worker is thus expected to supply (slightly) less effort. He just relies on the effort of his co-workers to diminish the inequality between himself and the employer.

It is important to realize that Figure A1 depicts the predictions for an *individual* worker of a given type. Predictions concerning average observed behavior clearly depend on the assumed type distribution within the population. With a mixture of worker types the horizontal parts in Figure A1 will get smoothed out when we consider workers' average effort. The main point that should be taken from the figure is that under inequality-aversion workers in the 1-1 treatment will make a larger effort if wages are higher but effort will be lower or even minimal in the 1-4 treatment.

Next assume that agents care about total surplus as well and have quasi-maximin preferences like in Charness and Rabin (2002). They assume homogenous agents and in the reciprocity-free version of their model player i's utility is given by:

$$U_i(m) = (1-\lambda)\cdot m_i + \lambda\cdot\left[\delta\cdot\min\{m_0,...,m_i,...,m_k\} + (1-\delta)\cdot(m_0+...+m_i+...+m_k)\right] \quad \text{(A3)}$$

In this expression parameter $\lambda\in[0,1]$ measures the extent to which player *i* cares about social welfare versus his own self-interest. Parameter $\delta\in[0,1]$ reflects the

[23] An alternative approach would be to assume, in line with Fehr and Schmidt (1999), that β's of 0, 0.25 and 0.6 are equally likely and that the workers know this. It can be shown that this does not affect the predictions. The intuition is that providing effort is relatively cheap and workers' earnings are therefore very close to each other. The theoretical predictions are then almost completely determined by the income inequalities between employer and worker.

concern for helping the worst-off player versus maximizing total surplus (efficiency). In the 1-1 treatment the worker's utility reduces to:

$$U_1(m) = \begin{cases} m_1 + \lambda(1-\delta)\cdot m_0 & \text{if} \quad m_1 \le m_0 \\ (1-\lambda\delta)\cdot m_1 + \lambda \cdot m_0 & \text{if} \quad m_1 > m_0 \end{cases} \qquad \text{(A4)}$$

The important difference with inequality-averse preferences reflected in (A3) is that the worker now also prefers higher earnings for the employer even when he himself is behind. Whether he is willing to sacrifice the marginal costs of effort $c'(e) \in \{1,2,3\}$ to give the employer an extra return of $v=10$ depends on the relative weight he puts on the employer's monetary payoffs m_0. From (A4) it follows that this relative weight is higher when the worker is ahead than when he is behind. In order to derive precise predictions we assume that $\lambda \geq 0.3$. This assumption is consistent with the estimates Charness and Rabin (2002) obtain for their model.[24] Given $c'(e) \leq 3$ a worker who is ahead is then always willing to increase his effort. After differences in payoffs are leveled he may still supply more effort, depending on the marginal costs and the worker's effective concern for efficiency $\rho \equiv \lambda(1-\delta)$. Because there are three values of $c'(e)$ in our experiment, there are four relevant cases to consider. When $\rho < 1/10$, referred to as case I, the concern for efficiency is effectively empty and only the concern for the worst off player remains. This in particular implies that a worker does not want to supply any effort when he is behind. In the other three cases a worker chooses more than minimum effort even if he earns less than the employer. The weight ρ workers attach to efficiency then strongly influences their effort choice.

Figure A2 below displays the predictions in the four different cases. In case I the worker does not really care about efficiency and the predictions closely track those of the inequality-aversion model reflected in Figure A1 (with $\beta \geq 0.25$).[25] When the concern for efficiency increases, the lowest effort that the worker supplies shifts up. Note that not the complete wage-effort relationship shifts up, but only the effort levels that were lower than the new minimum effort level. Hence, the higher the concern for

[24] Charness and Rabin (2002) estimate the quasi-maximin model using the data of a series of experimental games. They obtain an estimated λ of around 0.4 in various different specifications.

[25] Minor differences occur because in the experiment wages and effort can take particular discrete values only, so even within the four relevant cases there can be (very) small differences in predictions for different values of λ and δ. This does not occur when wages and effort can vary continuously.

efficiency, the flatter the wage-effort schedule becomes. This more general prediction continues to hold when we relax our assumption that $\lambda \geq 0.3$.

Figure A2: Predictions of the quasi-maximin model for the 1-1 treatment (with $\lambda \geq 0.3$ and $\rho \equiv \lambda(1-\delta)$).

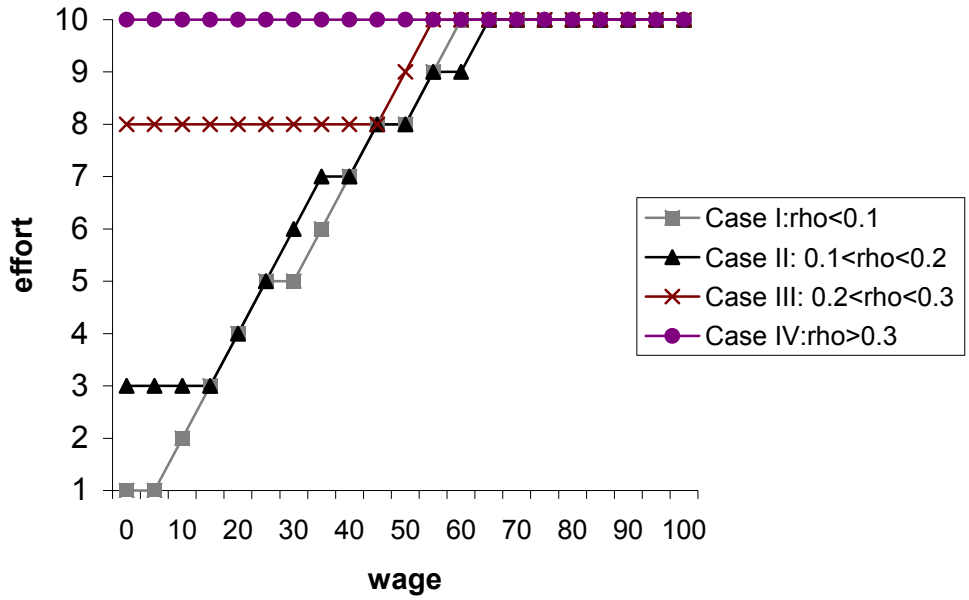

In the 1-4 treatment we again have to make assumptions about how individual workers perceive the game. As before we first assume that (A) each worker only cares about his own earnings and those of the employer (hence workers' preferences are given by (A4)), and (B) that each worker assumes that his co-workers behave similar. Under these assumptions Figure A3 results. Comparing the predictions with those for the 1-1 treatment we observe that a higher concern for efficiency makes the predictions more similar across treatments, but also necessarily increases the intercept of the wage-effort relationship at the same time.

If we relax assumption A the predictions for the 1-4 treatment do not change. This holds because a worker's effort decision does not affect the monetary payoffs of his fellow workers and, given that each worker behaves the same, it holds that $m_1=m_2=m_3=m_4$. The general utility function (A3) then reduces to (A4) plus a fixed constant. Although Charness and Rabin (2002) assume homogeneous preferences, we may also relax assumption B by using the stated expectations of individual workers in treatment 1-4E. Compared to Figure A3, workers are then predicted to choose slightly

lower effort levels. In case I, for example, a worker will provide the minimum effort level for all wages below 75 and the highest effort level he is willing to choose equals 4. In case II workers chooses e=3 as long as the wage is below 80. For higher wages they supply more effort, but never more than 5. The predictions for the remaining two cases are exactly equal to those in Figure A3.

When we consider average effort in a heterogeneous worker population the graphs in Figures A2 and A3 will get smoothed out. The essential prediction here is that when subjects care (sufficiently) about efficiency, differences between treatments 1-1 and 1-4 are reduced as compared to the predictions under inequality-aversion. But, effort supplied for low wages necessarily increases at the same time.

Figure A3 Predictions of the quasi-maximin model for the 1-4 treatment (with $\lambda \geq 0.3$ and $\rho \equiv \lambda(1-\delta)$).

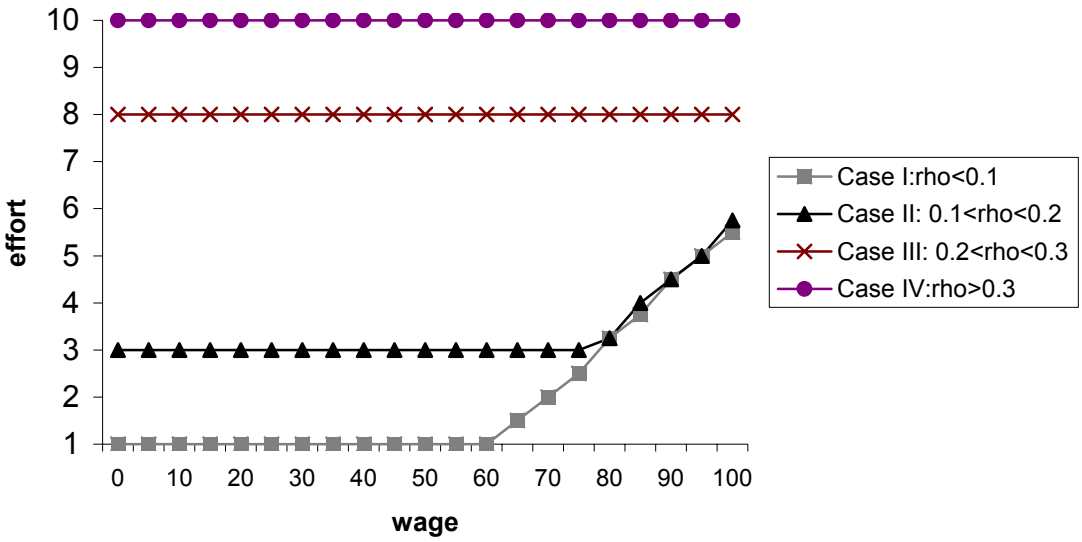

The formal analysis of the intention-based reciprocity models is quite complex and will not be done here. Dufwenberg and Kirchsteiger (2004) provide a full equilibrium analysis of the sequential Prisoners' dilemma game for their specification of reciprocity preferences. They show that when reciprocity concerns are sufficiently strong, cooperation can be sustained. Because the gift-exchange game corresponds to a sequential PD with multiple actions, this prediction translates to higher effort returned for higher wages; see also Charness (2004, p. 669) for an intuitive discussion. The irrelevance of the size of the workforce directly follows from how

Dufwenberg and Kirchsteiger measure the kindness of intentions. A given wage is just compared with the range of feasible wages, which is the same in both treatments. Falk and Fischbacher (2006, Subsection 4.2) formally analyze a gift exchange game similar to our 1-1 treatment within the context of their model. They show how positive reciprocity can create an upward sloping wage-effort relationship. The implications for gift exchange of the negative reciprocity element in the model of Charness and Rabin (2002) are informally discussed in Charness (2004, Section IV). Intuitively, when the employer "misbehaves" by choosing an intentionally low wage, the worker will withdraw his concern for social welfare. This pulls the effort levels provided for low wages in Figures A2 and A3 downwards.

Appendix B: Summary of the instructions for the 1-4 [1-4E] treatment

In this experiment you are taking part in a study of the labor market. There are two types of participants: employers and workers. 20% of the participants will be assigned the role of employer; the other 80% will be workers. You will be randomly assigned one of these roles. Which role you have, you will hear at the start of the experiment. Your role will not change during the experiment.

The experiment consists of one period only. In this period you will be randomly and anonymously matched with other subjects. Each worker is matched with one employer and each employer has four workers. You will not know with whom you are matched. During the experiment you will earn money based on the choices you and the participants with whom you are matched make. These earnings are calculated in points.

The single period has two stages. These stages have the following setup:

Stage 1 You have to make a decision <u>without</u> knowing the choices of those with whom you are matched. If you are an employer, you have to set the wage of your workers. This wage should be a multiple of 5 and in between 0 and 100. Each employer is allowed to set only one wage. The wage is equal for all four workers.

If you are a worker, you have to decide which effort level you want to provide for each possible wage set by the employer. There are 21 possible wages (ranging from 0 till 100), so each worker has to make 21 effort choices. These effort levels should be integers and in between 1 and 10. Effort is costly for the worker, and the costs (in number of points) belonging to a particular level of effort are reflected in the following cost schedule:

Cost schedule of feasible effort levels

Effort	1	2	3	4	5	6	7	8	9	10
Cost	**0**	**1**	**2**	**4**	**6**	**8**	**10**	**12**	**15**	**18**

Stage 2 In this stage you will be informed about the choices of others. If you are a worker, you will learn the wage chosen by the employer to whom you are matched. If you are an employer, you will learn the effort choices of your four workers for the wage that you offer. You will also be informed about your period earnings. If you are an employer you will also learn your workers' period earnings and if you are a worker you will also learn your employer's period earnings. These period earnings (in number of points) are calculated as follows:

Employer's period earnings = 10*(Effort level 1 + Effort level 2 + Effort level 3 + Effort level 4) – 4*Wage offered

Worker's period earnings = Wage – Cost of effort provided

At the start of the experiment you will receive an endowment that depends on the role that you have in the experiment. If you are an employer you will receive 360 points and if you are a worker you will receive 90 points. The total number of points that you earn in the experiment equals the sum of your initial endowment and your period earnings. At the end of the experiment these points will be converted into euros at the rate of: 9 points = 1 euro. In addition to the earnings earned in the experiment you will receive 5 euros.

[*In the 1-4E treatment the last part starting with the period earnings is replaced by*:

Employer's period earnings = 10*(Effort level 1 + Effort level 2 + Effort level 3 + Effort level 4) – 4*Wage offered + 360

Worker's period earnings = Wage – Cost of effort provided + 90

At the end of the experiment the period earnings will be converted into euros at the rate of: 9 points = 1 euro. In addition to the period earnings earned in the experiment you will receive 5 euros.]