

Gift exchange and the separation of ownership and control

May, 2007

Abstract

Numerous gift exchange experiments have found a positive wage-effort relationship. In (almost) all these experiments the employer both owns and controls the firm. This paper explores to what extent the separation of ownership and control affects the wage-effort relationship. We compare the standard bilateral gift exchange game between an owner-manager and a worker with two trilateral ones where the firm is owned by a shareholder and controlled by a manager. The wage-effort relationship is similar in all three situations. Most strikingly, workers reward higher wages with higher effort, even when the manager does not share in the firm's profits.

1 Introduction

The gift exchange game has been studied extensively in the laboratory. It represents the interaction between an employer and a worker in an environment where labor contracts are necessarily incomplete. In its most basic version the employer first decides on an unconditional wage transfer. After observing the wage that he will earn, the worker subsequently decides how much effort to supply. Effort increases the profits to the employer, but is also (increasingly) costly to the worker. Therefore, in case the worker is entirely selfish, he will not supply any effort at all, irrespective of the actual wage offered. Anticipating this entirely flat wage-effort schedule, a selfish employer just offers the minimum wage that is sufficient to induce the worker to accept the contract.

Experimental findings are in stark contrast to these theoretical predictions. Workers are typically willing to supply more effort when a higher wage is offered, yielding a significantly positive correlation between wages and effort commonly interpreted as reciprocal behavior (see e.g. Fehr and Gächter, 2000) for an overview). This finding appears rather robust and is found in various variations of the standard gift exchange game. For example, a number of recent studies still report a significantly positive correlation between wages and effort even with an interior optimal effort level instead of a corner solution (Engelmann and Ortmann, 2002), when high stake levels are considered (Fehr et al., 2002), with different productivity levels among employers (Hannan et al., 2002), and for subject pools different from the standard pool of undergraduate students (Hannan et al., 2002). The overall picture that emerges is that fairness and reciprocity considerations play a prominent role in labor relations.

However, all the robustness studies referred to above focus on *bilateral*, one-employer-one-worker, relationships (although sometimes embedded in a one-sided or double auction market).¹ This makes it problematic to directly extend the findings to the real world, because in the presence of social preferences – like fairness, efficiency norms and reciprocity considerations – the actual environment in which the employer-

¹Gift exchange appears robust both within and across different market institutions; see Fehr and Falk (1997) for an overview.

worker relationship is embedded becomes important. In particular, one has to take into account that in reality workers are typically part of a large(r) organization with many employees and multiple hierarchical levels. Fairness and reciprocity considerations are likely to be dependent on this more complex environment, making it a priori inappropriate to solely focusing on simple bilateral relationships.

This paper relates to earlier work in which it is investigated how robust the results of bilateral gift-exchange experiments are when the labor relationship is placed in a more complex and more realistic social context. In Maximiano et al. (2006) the gift exchange game is extended to a situation in which the employer employs several workers (see also Charness and Kuhn, 2006; Abeler et al., 2006). Somewhat surprisingly, there it is found that gift-exchange is very robust to increases in the size of the workforce. Their results also suggest that intention-based reciprocity is an important driving force behind gift exchange beyond outcome oriented social preferences.

Recent field experiments also explore the external validity of gift-exchange experimental results but the focus is instead on the representativeness of the laboratory environment. Findings from these field experiments are mixed. In a bilateral gift exchange game between a charitable organization and its potential donors, Falk (2007) finds that a small gift (like a postcard) significantly raised charitable donations. Gneezy and List (2006) test the gift exchange hypothesis in an actual spot labor market and explore the duration of the task - an important dimension of labor interactions not yet considered in laboratory experiments. They find that paying more than market-clearing wages has a short-lived effect on workers' effort, both in a library data entry task and in a door-to-door fund raising task.² However, placing the labor relation within a firm instead of in a spot labor market, Bellemare and Shearer (2007) obtain somewhat different results. They conducted a gift exchange field experiment within a tree-planting firm, where workers received a one-day surprise bonus. Their results suggest that the role of reciprocity is important within this firm, because workers

²Similar results are obtained by Kube et al. (2006) who replicate (and extend to negative reciprocity) part of Gneezy and List's experiment.

increased average daily productivity by 10 percent. Moreover, high tenure workers continue to reciprocate several days after the gift.

In this paper we explore gift exchange (in the laboratory) within multi-level hierarchies. In reality many firms have a complex hierarchical structure in which ownership and control are separated. The owners of a firm (shareholders) typically delegate authority to managers who are in charge of executive decisions. Part of the manager's job is to determine the firm's compensation structure. Within the context of the gift-exchange game this implies that the owners are not directly responsible for setting a worker's wage. It also implies that the manager who is in charge of wage determination, is not full residual claimant. She does not bear the (full) wage costs but also does not get the (full) benefits generated by the worker's reciprocal attitude. Separation of ownership and control may therefore have important implications for the observed relationship between effort and wages. For example, the worker may not choose a high effort in response to a high wage anymore, because the manager responsible for setting the higher wage does not benefit (sufficiently) from the worker's increased effort herself. Alternatively, the wage-effort relationship may remain largely unaffected, because the worker mainly wants to reward those who pay for the higher wage (i.e. the owners). Our experiment aims at investigating these issues.

We consider four different treatments, in three of which the firm decides on the wage itself. Our first treatment corresponds with the standard bilateral gift exchange game and serves as baseline. In this treatment ownership and control are vested in a single person, viz. an owner-manager. Two other treatments consider a situation in which the firm is owned by a shareholder but controlled by a manager. In both these trilateral treatments the manager chooses the worker's wage, i.e. has full control. They differ with respect to how the manager is rewarded though. In one trilateral treatment she receives a fixed wage and ownership and control are fully separated. The shareholder is then full residual claimant of the firm's profits. In the second trilateral treatment the manager is rewarded on the basis of performance pay, earning 25% of the firm's profits. Here ownership is shared between the shareholder and the manager and

the separation between ownership and control is more diffuse. In our fourth treatment the firm has no control and nature (i.e. a random device) decides on the actual wage offer (like in Charness, 2004). This bilateral owner-worker treatment is included to explore whether control per se is important for the gift exchange relationship observed.

Our main finding is perhaps somewhat surprising: the wage-effort relationship does not differ among the three treatments in which the wage is endogenously chosen by (a member of) the firm. Gift exchange is thus also observed in a more complex hierarchical structure and appears to be robust to the firm's separation of ownership and control. This finding is in line with both the hypothesis that workers only care about those who pay the wage (i.e. owners) and the hypothesis that they act reciprocal towards the firm as a whole and not towards one of its members in particular. The significantly less steep wage-effort relationship we observe for the treatment in which the wage is exogenously set by nature, however, suggests that the latter is a more likely explanation.

This paper proceeds as follows. In Section 2 we first describe the basic setup that we consider in our experiment and we subsequently derive the theoretical predictions under various behavioral assumptions. These behavioral predictions are derived from selfish preferences, outcome oriented social preferences (viz. inequality-aversion and quasi-maximin) and intention-based reciprocity. Section 3 presents the details of the experimental design. In Section 4 we report and discuss our main findings. Section 5 concludes.

2 Setup and behavioral predictions

2.1 Basic setup of the experimental game

The setting we consider concerns a firm which employs a single worker. Participation of the worker is secured, the only issue is to motivate the worker to put in (high) effort. Effort, however, is non-contractible. The firm can just offer an unconditional wage $w \geq 0$. After the wage offer has been made, the worker chooses how much effort

$e \geq e_{\min}$ to provide. Effort is increasingly costly, with $c(e)$ denoting the costs of effort in monetary terms. The marginal value product of effort is fixed per unit of effort and equals $v = 40$. The firm’s net profits equal:

$$\Pi = 40 \cdot e - w + 360 \quad (1)$$

Here the fixed term of 360 can be interpreted as the net profit the firm makes when the worker only supplies the enforceable effort level e_{\min} . The monetary payoffs for the worker m_W are given by:

$$m_W = w - c(e) + 20 \quad (2)$$

The fixed amount of 20 can be seen as the compensation the worker receives for supplying the enforceable effort level e_{\min} .

In the experiment the firm’s wage offer w has to be a multiple of 5 and between 0 and 100. The worker’s effort choice is restricted to integer values between 1 ($= e_{\min}$) and 10. The costs of effort are reflected in Table 1. With these parameters both the firm’s net profits Π and the worker’s monetary payoffs m_W are necessarily positive.

Table 1: The cost of effort function

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

Our main treatment variable concerns the composition of the firm. We compare two different hierarchical structures. The first one is the “Bilateral Condition” (BC) and provides the baseline. In this case the firm consists of a single *owner-manager* who decides on the worker’s wage herself. This baseline condition is similar to the common bilateral gift exchange game between an employer and a worker. The second hierarchical structure we consider is labelled the “Trilateral Condition” (TC). Here ownership and control are separated, resulting in a three-player game. In particular, the owner of the firm is now a shareholder who is not responsible for choosing the

worker's wage anymore; this responsibility has been delegated to a manager who has been hired to control the firm. The shareholder thus takes no decision at all. By keeping the firm's total net profits Π constant, we ensure that conditions BC and TC do not differ in terms of efficiency.

Separation of ownership and control implies that the one who decides does not bear the full costs and does not receive the full proceeds generated by her decisions. The characteristic feature of our trilateral condition, therefore, is that the manager is not full residual claimant. Instead, the shareholder claims most of the firm's net profits. To investigate the impact of this distinction between the one who decides and the one who obtains the (net) proceeds in the starkest possible way, we consider one treatment in which the manager does not share in the firm's net profits at all. In this TC0 treatment the manager is just paid a fixed wage (which is normalized to zero). Of course, in reality CEOs and other top-level executives in charge of controlling a firm are typically paid on the basis of performance pay. They thus share in the firm's net profits at least to some extent. The distinction between the one who decides and the one who gets the proceeds is then less extreme. Our TC25 treatment, in which the manager obtains 25% of the firm's net profits and the shareholder gets the remaining 75%, considers such a situation.

In all three treatments described so far the actual wage decision is taken by (a member of) the firm. One potential reason why gift exchange might be robust to the separation of ownership and control is that workers simply do not care about control at all. They only care about the fact that the wage is paid by the firm. In order to explore this we consider another bilateral treatment in which the owner of the firm does not choose the worker's wage herself. Instead, the wage is randomly determined by the experimenter, with each possible wage being equally likely. Because in this treatment the wage is exogenously given to the parties, we label it as BCexo. Table 2 depicts the players' monetary payoffs for the four different treatments that we consider.

Table 2: Players' monetary payoffs in the four treatments

Treatment	nr. of players	Shareholder	Manager	Worker
BC	$n = 2$	-	$40e - w + 360$	$w - c(e) + 20$
TC0	$n = 3$	$40e - w + 360$	0	$w - c(e) + 20$
TC25	$n = 3$	$30e - 0.75w + 270$	$10e - 0.25w + 90$	$w - c(e) + 20$
BCexo	$n = 2$	$40e - w + 360$	-	$w - c(e) + 20$

The above discussion reveals that, from a more conceptual point of view, an important difference between our treatments is the way in which the wage offer is generated. The wage can either be chosen by an interested party who has a stake in the worker's decision (BC, TC25), chosen by a 'third' party whose payoffs do not depend on the worker's choice (TC0), or determined by a random process (BCexo). Inspired by the ultimatum game experiments of Blount (1995), Charness (2004) studies exactly these three variations for the bilateral gift exchange game (see also Charness, 2000). His 'Employer' ('Random') treatment is comparable to our BC (BCexo) treatment. However, although his 'Third Party' treatment shares with our TC0 treatment that the wage is chosen by a party with no direct stake in the worker's decision, it differs in an important respect. In Charness (2004) the experimenter performs the role of 'third party' whereas in our case it is a subject with the role of manager. Unlike Charness, our treatment TC0 (and also TC25) thus explicitly considers a trilateral interaction among three subjects.³

2.2 Behavioral predictions and hypotheses

Our main interest lies in worker's behavior. When workers are entirely selfish and only interested in maximizing their own monetary payoffs, they will choose the minimum effort level $e^* = e_{\min} = 1$ irrespective of the wage offered. In that case effort will be minimal for all wages in all four treatments. A large number of experimental studies,

³A second, related variation between our treatments is whether an inactive, but interested dummy player (the shareholder in TC0 and TC25) is present or not. Such inactive dummy players have been incorporated in ultimatum game experiments by Güth and van Damme (1998), Bereby-Meyer and Niederle (2005), and Fershtman and Gneezy (2001).

however, have shown that a substantial fraction of workers make their effort choices according to considerations going beyond material self-interest. A common finding is a positive correlation between effort and wages, providing support for Akerlof (1982)'s efficiency wage theory. His *fair-wage hypothesis* predicts that firms will offer higher than market clearing wages, expecting that workers will work harder in return. Workers then compare the wage received with a norm they consider fair and choose whether to increase their effort or not. This results in a positive wage effort relationship.

The gift exchange relationship between workers and firms may originate from a variety of social norms. Recent theories of social preferences make this explicit and can be roughly classified into two kinds (cf. Camerer, 2003). The first type of theories are *outcome oriented* and assume that workers only care about the final distribution of payoffs. Workers may for instance compare their own payoffs with the average payoffs of others within the firm (Bolton and Ockenfels, 2000) or prefer an equitable distribution of material resources (Fehr and Schmidt, 1999). Alternatively, they may care about social welfare in general and take efficiency into account as well (Charness and Rabin, 2002; Engelmann and Strobel, 2004). A second type of social preference theories assume that people also care about how the distribution of payoffs came about. Workers may then have a preference for rewarding kind intentions of the firm and punishing unkind ones, independent of the final payoff distribution that results (Dufwenberg and Kirchsteiger, 2004). In that case a higher wage is rewarded with a higher level of effort, not (necessarily) because it leads to a more desirable distribution of payoffs but rather because the worker likes to do a favor in return. These theories of *intention-based reciprocity* thus stress the reciprocal nature of the gift-exchange relationship.

We derive the behavioral predictions for both types of theories. Under outcome-oriented social preferences treatments BC and BCexo can be treated as the same. First assume that workers are averse to inequality and have preferences like in Fehr and Schmidt (1999).⁴ In that case a worker is willing to increase (decrease) the other

⁴In particular, for a general n -player game with monetary payoffs $m = (m_1, \dots, m_n)$, player i 's utility is given by (for $i = 1, \dots, n$): $U_i(m) = m_i - \frac{1}{k-1}\alpha_i \sum_{j \neq i} \max\{m_j - m_i, 0\} - \frac{1}{k-1}\beta_i \sum_{j \neq i} \max\{m_i - m_j, 0\}$, with $0 \leq \beta_i < 1$ and $\beta_i \leq \alpha_i$.

firm members' material payoffs if these are below (above) the worker's own material payoffs. Given our parameter choices for Π and m_W , the owner-manager/owner is always ahead of the worker in both bilateral treatments. The worker therefore chooses $e^* = 1$ in response to any possible wage. In treatment TC0 the manager always earns less than the worker whereas the shareholder always earns more. Choosing higher than minimum effort does not affect the manager's earnings and only increases the shareholder's payoffs, thereby increasing inequality. Therefore, $e^* = 1$. Finally, in treatment TC25 the worker is always behind of the shareholder, but for relatively high wages ahead of the manager. Although choosing $e > 1$ for high wages decreases the payoff difference between manager and worker, it exacerbates the worker-shareholder inequality to a larger extent. Because the worker cares more about being behind than about being ahead, he again chooses $e^* = 1$ after any wage offered. All these predictions do not change if we instead assume that workers compare themselves with the *average* of the other firm members' payoffs, like Bolton and Ockenfels (2000) do.

The predictions obtained from inequality-aversion coincide with the ones based on selfish preferences. The driving force here is the large difference between the fixed amount of 360 appearing in the firm's payoff function (1) and the one of 20 the worker receives. In case these amounts are more equal, as is the case in most gift exchange experiments, at least in the two bilateral treatments a positive wage-effort relationship is predicted (see e.g. Figure A1 in Maximiano et al., 2006).

Selfishness/inequality-aversion hypothesis: Effort will be minimal for all wages in all treatments.

Next assume that workers care about overall surplus as well. In particular, let them be motivated by *quasi-maximin preferences* as introduced in Charness and Rabin (2002).⁵ Because in the bilateral treatments the worker is always the worst-off player,

⁵For a general n -player game with monetary payoffs $m = (m_1, \dots, m_n)$, player i 's utility is then given by (for $i = 1, \dots, n$): $U_i(m) = (1 - \lambda) \cdot m_i + \lambda \cdot [\delta \cdot \text{Min}\{m_1, \dots, m_n\} + (1 - \delta) \cdot (m_1 + \dots + m_n)]$. Here $\lambda \in [0, 1]$ gives the weight that player i attaches to social welfare, with the latter a weighted (by $\delta \in [0, 1]$) combination of the payoff of the worst-off player and overall surplus.

efficiency is the single reason that may induce him to exert more than the minimum effort level. Effectively, the worker's utility function can be written as $m_W + \rho_{BC} \cdot \Pi$.⁶ The larger his concern for efficiency, the larger parameter ρ_{BC} is. By putting in one extra unit of effort, the payoffs of the worker decrease by $c'(e) \in \{1, 2, 3\}$ whereas firm net profits increase with 40. The equilibrium effort level thus depends on how $c'(e)/40$ compares to ρ_{BC} . The important thing to note is that this effort level is *independent* of the wage offered; the predicted wage-effort relationship is entirely flat.

The same prediction applies for treatment TC0. Quasi-maximin preferences then correspond to a worker's utility of $m_W + \rho_{TC0} \cdot \Pi$. Because in this treatment the manager is the worst-off player, a higher effort level does not decrease the payoffs of the worst-off player anymore. This effectively increases the worker's concern for efficiency: $\rho_{TC0} \geq \rho_{BC}$.⁷ The worker will therefore choose an effort level that weakly exceeds the one in the bilateral treatments, i.e. $e_{TC0} \geq e_{BC}$. But again, the predicted wage-effort relationship is entirely flat. Finally, in treatment TC25 either the worker or the manager is worst-off, depending on the wage and effort levels chosen. For all wages $w \leq 60$ the worker necessarily earns the least and the situation is equivalent to treatment BC. Hence the wage-effort relationship is flat at $e = e_{BC}$ for all $w \leq 60$. In case $w > 60$ the manager earns the least when the worker would put in low effort whereas for high effort levels the worker gets the lowest payoff. The cutoff level of effort that divides the two cases is increasing in the wage. If the worker has a high concern for the worst off player (and thus a relatively low concern for efficiency), he just wants to exert effort up to this cutoff level. An increasing wage-effort relationship may therefore result for relatively high wages. But, for the larger part ($w \leq 60$) the wage-effort relationship will also necessarily be flat in treatment TC25.⁸

⁶In terms of the specification of Charness and Rabin (2002) we have $\rho_{BC} = \lambda \cdot (1 - \delta)$. Recall that under outcome oriented social preferences treatments BC and BCexo can be treated as equivalent.

⁷It holds that $\rho_{TC0} = \frac{\rho_{BC}}{1 - \lambda\delta}$.

⁸For the predicted effort level e_{TC25} it holds that $e_{BC} \leq e_{TC25} \leq \max\{e_{BC}, 5\}$. To see this, note that the worker prefers effort level e_{BC} given that he is the worst-off player. In case the manager is worst-off, effort incentives are (weakly) higher. This yields the lower bound. Similarly so, the upper bound follows because for $e \geq 5$ the worker is necessarily worst-off (i.e. independent of the wage). He will therefore never provide more than his optimal effort given that he is worst-off, i.e. e_{BC} , when this optimal level falls within this range. With $e_{BC} \leq e_{TC25} \leq \max\{e_{BC}, 5\}$ it follows that when $e_{BC} \geq 5$, also in treatment TC25 the wage-effort relationship is entirely flat at $e_{TC25} = e_{BC}$ for all wages. This

Efficiency hypothesis (quasi-maximin preferences): The wage-effort relationship is entirely flat in all treatments, except (possibly) for wages above 60 in treatment TC25.

From the above it follows that under outcome-oriented social preferences wages and effort are likely to be uncorrelated. The two leading models within this class, viz. inequality-aversion and quasi-maximin preferences, both predict an (almost) entirely flat wage-effort relationship. Therefore, if we observe a positive correlation in any of our four treatments, (intention based) reciprocity is likely to play an important role. When workers have a preference for behaving reciprocally, the *slope* of the wage-effort relationship may differ across treatments though. We next derive some hypotheses that make these differences explicit.

In the baseline BC-treatment the owner-manager both chooses the wage and captures the entire net profits Π . A higher wage benefits the worker and thus can be considered kind, also because the owner-manager has to pay the wage from her own pockets. Similarly, a low wage can be seen as unkind. A reciprocal worker then prefers to reward high wages and punish unkind ones. Because the monetary payoffs of the owner-manager are proportionally increasing in the worker's effort, the worker is expected to choose higher effort levels in response to higher wages. Hence a positively sloped wage-effort relationship results.

The owner-manager in the baseline treatment personifies both the one who controls the wage and the one who pays for the wage. Therefore, no matter which (combination) of these two roles the worker prefers to reward, predictions are the same. This does not apply in the other three treatments. First consider treatment TC0 where ownership and control are fully separated. The worker may then still consider higher wages as being more kind (although possibly to a smaller extent because the manager does not pay the wage herself), but he has no opportunity to reward the person responsible for choosing the higher wage. This holds because the manager's payoffs are

is the case whenever $\lambda(1 - \delta) \geq 1/5$, i.e. whenever the worker cares sufficiently about efficiency.

independent of the worker's effort choice. We may therefore expect that the worker chooses the same (minimum) effort level regardless of the wage. This is actually what the intention-based reciprocity model of Dufwenberg and Kirchsteiger (2004) predicts. The worker will feel neutral towards the shareholder, as she has no choice to make. And because the worker can neither be kind nor unkind to the manager, his effort choice will be guided by monetary payoffs m_W alone. As a result, the wage-effort relationship will be flat and thus *less* steep than in treatment BC.

Alternatively, the worker may just want to reward (or punish) the person who bears the actual costs of the (un)kind wage offer. If that is the main motivational force, we predict no differences between the slopes of the wage-effort relationship in treatments TC0 and BC. This also holds when the worker perceives the firm as being his employer, independently of its composition. In that case he will feel reciprocally towards the firm as a whole and not towards one of its members, i.e. the one who controls it and/or the one who owns it, in particular.

Predictions for treatment TC25 are related to the ones above. In the model of Dufwenberg and Kirchsteiger (2004) the kindness of a particular wage offer is judged relative to what the firm could offer the worker in principle. The perceived kindness of a given wage is thus the same as in treatments BC and TC0. The effort level the worker chooses in response then depends on the effectiveness of effort as a reciprocation device. The more effective effort is in this regard, the higher actual effort will be. When the worker prefers to reward only the one who controls the wage, we predict that the slope of the wage-effort relationship falls in between the one for BC and the one for TC0 (where it is predicted to be flat). However, in case the worker prefers to reward those who pay for the wage, we expect the same positive slope as in treatments BC and TC0. This follows because in all treatments effort is equally effective in rewarding the group of contributors to the wage. Also when the worker feels reciprocally towards the firm, the predicted slope of the wage-effort relationship is the same in treatments BC, TC0 and TC25.

Finally, consider treatment BCexo in which the wage is selected at random by the experimenter. In this case the controlling role of the firm is entirely absent. Trivially then, a flat wage-effort relationship is predicted when the worker is mainly motivated by reciprocity towards the one who controls the wage. This prediction also follows from the model of Dufwenberg and Kirchsteiger (2004), the reason being the lack of intentions that characterizes the BCexo treatment. If, however, the worker solely feels reciprocally towards the one who pays for the wage, the slope of the wage-effort relationship will be the same as in the other three treatments. A final possibility is that both control and ownership are important for the worker's reciprocal feelings, although he does not care exactly how these two roles are parcelled out within the firm. In that case the slope will be less steep in treatment BCexo as compared to the other three treatments.

The above predictions based on reciprocity can be summarized as follows (see also Figure 1).⁹

Reciprocity hypothesis: Effort levels are increasing in the wage offered in treatment BC. This is consistent with a motivation to reward and/or punish the one who controls the wage, with a motivation to reward and/or punish the one who pays for the wage, or a combination of both. These underlying motives predict differences among treatments though, as follows:

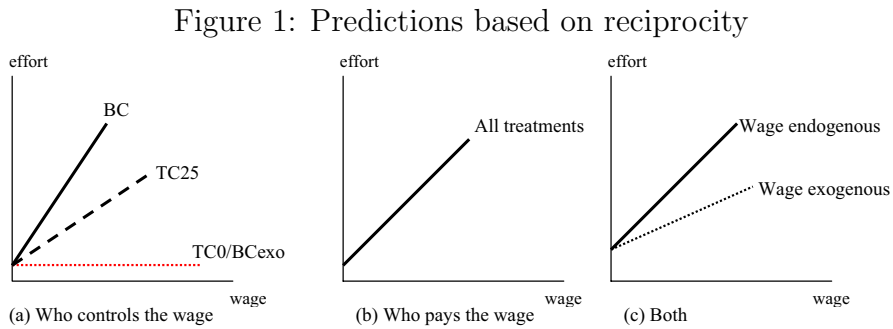
Reciprocity towards who controls the wage: The slope of the wage-effort relationship in treatment TC25 will be less steep than in treatment BC. The

⁹Although we have separated the two motives in the main text, in practice workers may be motivated by both outcome-oriented social preferences and reciprocity. In the theory of Falk and Fischbacher (2006) the (un)kindness a worker experiences from a given wage offer depends on both the intentions and on the payoff (in)equality that results. Their model predicts that $e^* = 1$ for all possible wages in treatments BC, TC0 and BCexo. In the two bilateral treatments this follows because the worker always earns less than the owner-(manager), in treatment TC0 it follows because the worker cannot affect the manager's earnings. Finally, also in treatment TC25 the wage-effort relationship is flat at $e^* = 1$ whenever the manager earns more than the worker, which necessarily happens when $w \leq 60$. For $w > 60$ we may observe a positively sloped wage-effort relationship. The full model of Charness and Rabin (2002) also combines outcome-oriented social preferences and (negative) reciprocity; individuals are motivated by social welfare but withdraw this concern when others behave unfair.

wage-effort relationship will be flat in treatments TC0 and BCexo.

Reciprocity towards who pays for the wage: The slope of the wage-effort relationship will be the same in all four treatments.

Reciprocity towards the firm as a whole: The slope of the wage-effort relationship will be the same in treatments BC, TC0 and TC25, and less steep in treatment BCexo.



3 Experimental design and procedures

3.1 Experimental design

Our experiment is inspired by the bilateral gift exchange game of Fehr et al. (1998). The specifics of our design are more in line with Maximiano et al. (2006) though. Like in that paper we let subjects play the gift exchange game only once, rather than repeatedly in a row. One-shot interaction has the advantage that subjects' choices cannot be guided by intertemporal considerations. For example, if a worker displays a reciprocal attitude, this cannot be explained by a desire for reputation formation or other repeated game considerations. A second important feature is that we make use of the strategy method. Each worker has to make a contingent effort choice, indicating his/her response to every possible wage offer. Because there are 21 possible wage levels, each worker has to make 21 (contingent) effort choices.

The main advantage of using the strategy method is that it allows us to elicit the wage-effort relationship of every individual worker. Besides that, given that in our experiment subjects interact only once, the use of the strategy method generates a relatively large amount of data at low cost. Theoretically the use of the strategy method should not matter. In practice, however, behavior may be affected because workers have to think about all possible wage offers and not just about the firm's actual wage offer. The literature is still inconclusive about the impact of the elicitation method. A number of experimental studies finds that the strategy method does not affect subjects' behavior; e.g. Cason and Mui (1998) for a dictator game, Brandts and Charness (2000) for both a chicken and a prisoners' dilemma game, Oxoby and McLeish (2004) for an ultimatum game, Sonnemans (2000) and Bosch-Domènech and Silvestre (2005) for an individual decision making setting and Falk and Kosfeld (2006) for a simplified principal-agent game. This last experiment is closest to ours with respect to the game studied. Some other studies do find different results; e.g. Güth et al. (2001) for a binary-offer ultimatum game and Brosig et al. (2003) for two different sequential games. Taking all these results together it seems fair to conclude that the impact of using the strategy method is rather limited. Moreover, even if the strategy method influences behavior, there is no reason to expect that this influence will differ among our treatments. Treatment comparisons are thus likely to be unaffected.¹⁰

Compared to Fehr et al. (1998) and Maximiano et al. (2006) the payoff functions that we use here differ in two notable ways. First, the marginal returns to effort are four times higher, i.e. we have $v = 40$ rather than $v = 10$. This secures that even in treatment TC25 effort is a sufficiently effective instrument to reward the one who controls the wage (i.e. the manager). Second, in all our treatments the worker's payoffs are always substantially below the firm's net profits. This allows a better identification of the motives underlying workers' reciprocal behavior. As discussed in Subsection 2.2, both inequality-aversion and quasi maximin cannot explain a positively sloped wage-effort relationship when the worker is always behind. If the two fixed amounts of 360

¹⁰Like Charness and Dufwenberg (2006, p. 1585), we are not aware of any study in which a treatment effect is found when using the strategy method but not when using the direct response method.

and 20 appearing in Π and m_W respectively were chosen more equally, these underlying motives would be more difficult to separate. This holds because in that case also under inequality-aversion and quasi-maximin preferences a positive wage-effort relationship may result. For the intention-based reciprocity predictions the inequality between the two fixed endowments is irrelevant. Here, the focus is on the kindness underlying a particular wage offer rather than on the distribution of outcomes. Moreover, as we assume that the worker judges the kindness of a particular wage offer relative to what the firm could give him in principle, the amount of the firm's initial endowment is immaterial.¹¹

3.2 Experimental procedures

The experiment was run at the University of Amsterdam in November 2005 and February 2006. Overall 293 subjects participated, most of them were undergraduate students in Economics and Business (58%). The average age of participants was 22 years and 42% of them were female. Overall we conducted 14 sessions. Three sessions considered the BC treatment (62 participants in total), another three sessions the BCexo treatment (60 subjects), four sessions the TC0 treatment (87 subjects) and another four sessions the TC25 treatment (84 subjects). Subjects earned on average 18.7 euros (including a show-up fee of 12 euros) in less than one hour. Earnings varied substantially, with the minimum earnings equal to 12 euros and a maximum of 35.3 euros.

The experiment was computerized using the z-Tree software developed by Fischbacher (1999). Subjects started with general on-screen instructions. They also received a summary of the instructions on paper (see Appendix A). To ensure that subjects understood the experiment, in particular how their payoffs were calculated,

¹¹Alternatively, if the worker judges the kindness of a particular wage offer relative to what the firm could get in principle, the reciprocity hypothesis would remain the same for all treatments except for TC25 in the case the worker cares about who controls the wage. Here, and compared to the BC treatment, the effectiveness of effort as a reciprocal device is lower but the kindness of a particular wage offer is higher, i.e. $wage/\Pi_{BC}^{manager} < wage/\Pi_{TC25}^{manager}$. So, in this case it would be possible to observe the same wage-effort relationship in both BC and TC25.

all subjects had to answer a number of control questions correctly before the experiment started.¹² At the start of the experiment subjects learned their roles. In the bilateral sessions half of the subjects performed the role of owner-manager cq. owner, the remaining half were assigned the role of worker. In the trilateral sessions, one-third of the subjects were given the role of shareholder, one-third the role of manager and the remaining one-third were assigned the role of worker. In all treatments subjects were randomly and anonymously matched into firms with either two (BC and BCexo) or three (TC0 and TC25) firm members.

When making actual choices, workers had to fill in a wage-effort table without knowing the actual wage. In particular, for all 21 possible wage levels (multiples of 5 between 0 and 100) they had to indicate their effort choice, an integer between 1 and 10. In the BCexo treatment, the owner did not take any decision. The wage was randomly drawn individually by the experimenter in front of each worker. This was done after all workers had made their choices. In the BC-treatment the owner-manager had to set the wage without knowing her/his worker's effort choices. In the trilateral treatments the manager did so. Note that in the trilateral treatments shareholders did not take any decision at all. After everyone had made a decision (if any), all subjects were informed about the choices of the other members in the firm. They also learned their own and the other firm members' payoffs. No information was given about choices and payoffs in firms to which the subject did not belong.

At the end of the experiment the number of points a subject had earned were converted into euros, with 30 points corresponding to 1 euro. Each participant learned his/her own earnings in euros, filled in a short background questionnaire and was individually and privately paid.

¹²Like in Maximiano et al. (2006) subjects generated the numerical examples for the control questions themselves, by first making hypothetical choices for both the (owner-)manager's and the worker's role (in treatment BCexo they did so only for the worker's role).

4 Results

In this section we report the main results of our experiment. Because we are mainly interested in workers' willingness to reciprocate, we first present workers' effort choices and use these to evaluate the hypotheses presented in Subsection 2.2. After that we turn to (owner-) manager behavior and overall earnings.

4.1 Worker behavior

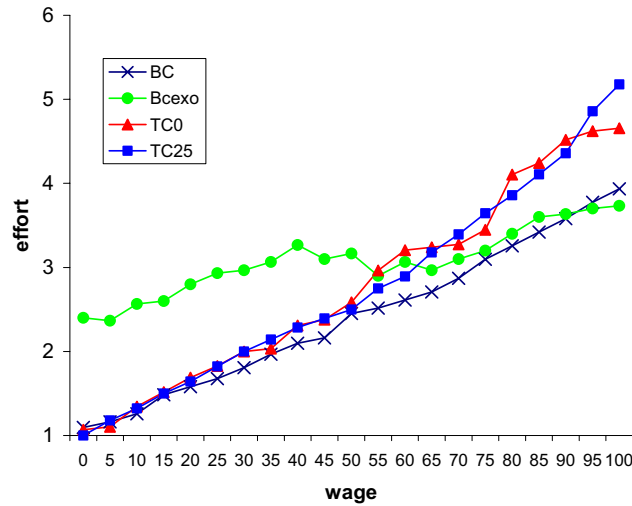
In all treatments the average effort of the workers over all possible wages is higher than the minimum effort level. In particular, it equals 2.41 in treatment BC, 2.77 in TC0, 2.76 in TC25 and 3.07 in treatment BCexo. This contradicts the predictions based on standard theory and inequality-aversion that the worker will always provide the minimum effort regardless of the wage received. Moreover, also the efficiency hypothesis based on quasi-maximin preferences is rejected. This follows from Figure 2 below, which displays the average effort by wage for each treatment separately. The higher the wage offered, the higher the average effort level that is chosen in response. This positively sloped wage-effort relationship holds no matter the firm's composition and how profits are distributed among firm members. Specifically, the Spearman rank correlation between effort and wages is significant at the 1% level and around 0.29 in the BC treatment, 0.35 in TC0, 0.33 in TC25 and 0.15 in the BCexo treatment. The following result summarizes these findings.

Result 1. In all treatments mean effort levels are above minimum and increasing in the wage.

Figure 2 also suggests that the wage-effort relationship is less steep when wages are determined exogenously than when wages are chosen by a member of the firm. Moreover, workers' behavior appears very similar across treatments in which wages are endogenously chosen. To further explore these two observations we regress workers'

effort on the wage level, three treatment dummies and four interaction terms. To account for the panel structure of our data (we have 21 data points per individual worker), we estimate a random effects model. The results are presented in Table 3. The intercepts in the BC treatment and in both trilateral treatments are equal and figure around one. The intercept in the BCexo treatment is significantly higher and around 2.5. Moreover, also the average effort provided for low wages is higher in treatment BCexo than in the other three treatments; (Mann-Whitney) ranksum tests for each wage separately reveal that average effort is significantly higher (at the 10% level) in treatment BCexo for wages of 10 and below.¹³ This finding indicates that workers punish wages that are intentionally low. It is in line with previous experiments that investigated the role of intentions for reciprocal behavior (cf. Charness, 2004; Falk et al., 2000).

Figure 2: Average effort by wage



¹³This holds for all bilateral comparisons between BCexo and BC, TC0 and TC25, respectively, except for the difference in average effort levels between BCexo and TC25 when $w = 10$.

Table 3: Random effects linear regression with effort as dependent variable

<i>Independent variable</i>	<i>coefficient</i>	<i>robust std. err.</i>
Constant	0.9891	0.1904***
Wage	0.0283	0.0070***
TC0	-0.1123	0.3160
TC25	-0.0276	0.2494
BCexo	1.4939	0.5800***
TC0*Wage	0.0095	0.0102
TC25*Wage	0.0042	0.0101
BCexo*Wage	-0.0165	0.0082**
TC25*Wage*I _{wage>60}	0.0055	0.0035

*** Indicates significance at the 1%-level. ** Indicates significance at the 5%-level. There are $N = 2478$ observations in total, with 118 clusters (workers) and 21 observations per cluster. The Wald statistic equals $\chi^2 = 195.52$; $p = 0.000$.

With respect to the estimated slope of the wage-effort relationship the following observations can be made. The highly significant wage coefficient reinforces the reciprocity hypothesis and rejects that only inequality-aversion and/or efficiency considerations are at play. The interaction terms reveal that the slopes in both trilateral treatments are steeper than the slope in treatment BC. Increasing the wage by 10 points leads to an estimated increase in effort of 0.32 in treatment TC25 and 0.38 in treatment TC0, as compared to around 0.28 in treatment BC. However, these differences are not statistically significant. The estimated slope of 0.12 for the BCexo treatment is significantly lower though (but still significantly different from zero).¹⁴ Given the efficiency hypothesis for treatment TC25 (see Section 2.2) we also include the interaction term $TC25*Wage*I_{\{wage>60\}}$ to check whether the slope in treatment TC25 is steeper for wages higher than 60. The estimates report a positive but insignificant coefficient.

In addition, we can perform non-parametric tests. First, we consider a two-sample ranksum test for each wage separately. It shows that differences in

¹⁴Somewhat different results are obtained when we restrict the analysis to wages higher than 50. For this subsample the estimated slope does not differ across all four treatments.

average effort levels between treatments BC and TC0, treatments BC and TC25, and treatments TC0 and TC25 are insignificant for all wage levels. In contrast, differences between BCexo and each of the other treatments are significant for very low wages. Second, in order to check whether the findings at the aggregate level are replicated at the individual level, we fit a linear effort function $e_i = a + b * w_i$ for each individual subject. We subsequently use a two-sample ranksum test to compare differences in individual slopes across treatments.¹⁵ Ranking individual estimates of b 's from low to high we do not find statistical evidence that the average rank in the BC treatment is higher than in each of the other treatments. We also do not find statistical evidence that supports the hypothesis that the average rank in treatment TC25 is higher than in treatment TC0. However, the average rank in TC25 is significantly higher than in BCexo at $p=0.01$ (one-tailed test) and it can be rejected that the average rank in TC0 equals the average rank in BCexo at $p=0.01$ (two-tailed test).¹⁶ Comparing all the estimates in the three treatments in which the wage is endogenously chosen with the estimates in treatment BCexo, we see that the average rank in BCexo is significantly lower at $p=0.01$ (one-sided test). Table B2 in Appendix B reports the average fitted wage coefficients and the test results.

Result 2. The wage-effort relationship is equally steep in all three treatments in which the wage is endogenously chosen.

Result 3. The wage-effort relationship is steeper when the wage is chosen by a member of the firm than when it is exogenously determined.

¹⁵ Charness (2004) performs the same test. He estimates the wage-effort slope for each worker using worker's actual effort choices and corresponding wage offers for a total of 10 periods. Using actual choices might overestimate (underestimate) the impact of wages if very low (high) wages are only rarely observed and if workers behave reciprocal. Additionally, his inferred individual slopes are not independent due to interaction effects. So, our analysis benefits from using the strategy method by which we can better elicit the wage-effort relationship for every individual.

¹⁶We use a one-tailed test whenever the reciprocity hypothesis predicts the direction of the difference in slopes (see Figure 1), otherwise we use a two-tailed test.

To further investigate worker behavior in more detail we analyze individual worker's choices and classify all workers into four categories. In particular, we use the estimated wage coefficient b together with the observed monotonicity of the relationship to classify individual workers. The four different categories are defined as follows:

Selfish/inequality-averse behavior: This group consists of those workers who provide the minimum effort level $e = 1$ regardless of the wage.

Efficiency oriented behavior: This class contains those subjects with an entirely flat wage-effort relationship at some higher than minimum effort $e > 1$, and also those subjects in treatment TC25 whose wage-effort relationship is flat up to a wage of 60 and non-decreasing after 60.

Reciprocal behavior: These subjects increase their effort with wages in a monotonic way. Workers within this category have a significantly positive wage coefficient; $b = \partial e / \partial w > 0$.¹⁷

Erratic behavior: These subjects do not make monotonic effort choices. This class also contains those subjects with “random” behavior in the BCexo treatment (e.g. to mimic the randomness in the wage).

Note that the first two categories together cover quasi-maximin preferences. Workers who display efficiency oriented behavior just have a concern for efficiency that is sufficiently strong to make a difference with selfish/inequality-averse behavior (cf. Subsection 2.2).

Figure 3 shows the distribution of subjects over the four categories by treatment. The figure also indicates the average effort provided by the different worker types. In

¹⁷Here subjects in treatment TC25 whose wage-effort is non-decreasing for wages above 60 are excluded, because they are already included in the class of efficiency oriented behavior. In the experiment we only have a single subject who behaves this way; s/he chooses $e = 1$ for $w \leq 90$ and $e = 7$ for $w > 90$.

all treatments a large fraction of workers behaves selfish or according to inequality-averse preferences. However, this does not explain our Result 2 that the wage-effort relationship is equally steep in treatments BC, TC0 and TC25. Estimating a random effects model like in Table 3 for the subsample of reciprocal types only, the wage-effort relationship is still equally steep in the treatments with an endogenously chosen wage. Also the distribution of worker types does not differ across these treatments. Although the percentage of selfish/inequality-averse workers is highest in treatment BC, differences with treatments TC0 and TC25 are insignificant. A chi-square test reveals that there are no differences in the frequency distribution of worker types between treatments BC and TC0, between BC and TC25, and between treatments TC0 and TC25.


The type distribution in the BCexo treatment is somewhat different though. In particular, there are fewer subjects acting reciprocally in this treatment, as compared to the three endogenous treatments.¹⁸ Similar results are obtained from comparing the distribution of individual slopes across treatments by means of a Kolmogorov-Smirnov test. There is no difference in the distribution of individual slopes between each pair of treatments in which the wage was endogenously chosen. In contrast, the distribution of individual slopes is significantly different in treatment BCexo (see Figure 1 and Table 2, in Appendix B). Apart from fewer reciprocal subjects, the average contribution of these type of workers is also lower in the BCexo treatment. These two observations together explain why the wage-effort relationship is significantly less steep when the wage is exogenously determined (Result 3). Finally, in this treatment we observe some workers that display efficiency oriented behavior. The typical worker within this class chooses a very high effort level of around $e = 9$, independent of the wage drawn by nature. This can be explained by the very favorable ratio of low marginal costs (3 at most) to

¹⁸This follows from making bilateral comparisons between treatment BCexo and treatments BC, TC0 and TC25, respectively, by means of a binomial test. Differences in the fraction of reciprocal workers between BCexo and the two trilateral treatments are significant at the 5% level, differences between BCexo and BC are not. The latter may result from the fact that the binomial test is not very powerful with low numbers. Comparing the complete type distribution over the four categories by means of a chi-square test, significant (10% level) differences are found only between treatments BCexo and TC0.

high marginal returns (40), which allows for large efficiency gains. Moreover, the fact that efficiency-oriented behavior is almost exclusively observed in the BCexo treatment may indicate that some workers may have a preference for efficiency but withdraw this concern for intentionally low wages.¹⁹

The main findings concerning the distribution of worker types are summarized in Results 4, 5 and 6 below.

Figure 3: Classification of workers by treatment and average contribution per worker type



Result 4. In all treatments there is a large proportion (around 50%) of workers that behave selfish or inequality-averse.

Result 5. The distribution of worker types does not differ across treatments in which the wage is endogenously chosen by a member of the firm.

Result 6. There are fewer workers behaving reciprocally in the BCexo treatment and their average effort contributions is lower as compared to the endogenous treatments, although differences are not always significant.

¹⁹Therefore, in the three endogenous treatments we cannot exclude that some reciprocal workers may behave according to the full model of Charness and Rabin (2002). However, we may exclude efficiency-oriented behavior from reciprocal behavior in treatment BCexo. Here it seems less likely that a worker will withdraw his efficiency concern to punish non-intentionally low wages.

Taken together, the aggregate results support the reciprocity hypothesis; on average we find a positively sloped wage-effort relationship. However, we also observe large variations in individual worker's behavior. In all treatments, around half of the workers provide the minimum effort regardless of the wage. When the wage is endogenously chosen, the remaining half behaves reciprocally. Comparing these endogenous treatments with other (bilateral) gift exchange experiments, the percentage of reciprocal workers that we find here is somewhat lower. For example, Maximiano et al. (2006) obtain a fraction of 64% of reciprocal workers and the overall wage-effort relationship found there is also steeper than the one we observe here.²⁰ The explanation for this is straightforward. Because of the highly asymmetric fixed amounts in Π and m_W in the present experiment, also inequality-averse subjects should display a flat wage-effort relationship (cf. Section 2). In most previous experiments the initial endowments are much more symmetric and therefore inequality-averse subjects are predicted to choose a positive wage-effort relationship. This explains why we find a significantly lower slope and also around 10% more workers that always choose the minimum effort. At the same time, we still find that the slope is substantial and significantly different from zero. Like in Maximiano et al. (2006), these findings suggest that reciprocity is an important driving force behind gift exchange, beyond (exclusively) outcome-oriented social preferences.²¹

Looking at whom the worker likes to reciprocate, the results for treatment TC0 clearly reject the hypothesis that he is only reciprocal towards the one who controls the wage. The one who pays for it certainly plays a role. Our finding that the wage-effort relationship does not differ among the three endogenous treatments (Result 2) is in fact in line with the worker only caring about the one who pays for the wage. But it also

²⁰Here we use the bilateral 1-1 treatment in Maximiano et al. (2006) as comparison. In this treatment the owner-manager's earnings equal $10e - w + 90$ whereas the worker's earnings are $w - c(e) + 90$. The wage coefficient found for that treatment equals 0,0556 compared to 0.0283 here for the bilateral treatment (cf. Table 3). We acknowledge though that this is not a perfect comparison, given that also the returns on effort differ besides the fixed amounts.

²¹See Falk et al. (2005) and Falk et al. (2000) for evidence that reciprocity is quantitatively the more important motivational force.

supports the hypothesis that the worker is reciprocal towards the firm as a whole and not particularly sensitive to its exact composition, i.e. to how ownership and control are divided within the firm. The finding that the wage-effort relationship is significant less steep in the BCexo treatment (Result 3) in fact indicates that this is the more plausible explanation. Control per se is thus important for the reciprocal relationship (cf. Charness, 2004 and Falk et al., 2000). At the same time, the significantly positive slope we observe in treatment BCexo reinforces the observation that the worker also prefers to reciprocate the one who pays for the wage. This holds because, unlike in the earlier experiments of Charness (2004), the positive wage-effort relationship in our BCexo treatment cannot be attributed to distributive concerns like inequality-aversion or quasi-maximin.

Overall we conclude that the worker feels reciprocally towards the firm; both the fact that the firm controls the wage and pays for it are important for the gift exchange relationship. Exactly how ownership and control are subdivided among (potentially) various firm members appears unimportant though.

4.2 Wage offers and earnings

Although our main interest lies in workers' effort choices, we briefly look at (owner-) managers' behavior as well. Standard theory predicts that (owner-) managers should pay the minimum wage in treatments BC and TC25. However, if the (owner-) manager herself is inequality-averse or expects to gain or lose from a positive or negative reciprocal reaction of the worker, she may pay more than the minimum wage. Because in treatment TC0 the manager does not benefit herself from the worker's effort, a higher wage there may reveal the manager's preferences for equality and/or efficiency. But in this treatment any wage offer is consistent with selfish preferences as well. Table 4 presents the average wages (besides some other measures of the wage distribution) for all treatments.

Table 4: Wage distribution measures and Mann-Whitney tests

	BC (n=31)	TC0 (n=29)	TC25 (n=28)	BCexo (n=30)
Average	61.45	67.24	49.29	64.5
Percentile 25	35	35	17.5	45
Median	70	90	52.5	65
Percentile75	100	100	77.5	90
% offers=0	19.35	13.79	17.86	0
% offers=100	29.03	44.83	14.29	13.33
BC vs. TC0	0.3858 [†]			
BC vs. TC25	0.1630 [†]			
TC0 vs. TC25	0.0469 [†]			

[†] Remark: p-value of a ranksum test comparing wage offers between the two treatments.

Because wages in BCexo are drawn at random we focus on the endogenous treatments. Only a small percentage of (owner-) managers pays the minimum wage of zero and the average wage is substantially higher in all cases. In the BC treatment the mean wage equals 61.5 points. In treatment TC0 the manager offers on average around 6 points more and in treatment TC25 on average 12 points less. These differences between treatments increase when we look at the median wage. Especially in treatment TC0 the median wage is very high and equal to 90. Moreover, a larger percentage of managers (44.8%) offers the maximum wage of 100 when they do not bear (part of) the wage costs. Despite the observed differences between treatments, ranksum tests fail to find significant differences in wages between the BC treatment and the two trilateral ones. But wages are significantly higher (at the 5% level) in treatment TC0 as compared to treatment TC25. These results carry over to the proportions of minimum and maximum wage offers; we find a significant difference (at the 5% level) for the proportion of maximum wage offers between treatments TC0 and TC25. (Higher wages in TC0, however, do not come at the expense of firm's profits; see the discussion of Tables 5 and 6 below.)

Result 7. Wages in the endogenous treatments are on average substantially higher than minimal. Wages are highest in treatment TC0 and lowest in treatment TC25, although differences among treatments are typically insignificant.

We next turn to the actual earnings subjects obtain in the experiment. Table 5 reports the subjects' average earnings by treatment and by role, excluding the show up fee of 12 euros every participant received for showing up on time. Workers earn slightly more in TC0, which is driven by the higher wages they receive in this treatment. The firm's overall net profits are highest in treatment TC25. Here the wage costs are the lowest and the effort provided the highest. Despite the higher wage costs in TC0, the firm's overall profits are higher in TC0 than in treatment BC.

The disparity between the workers' and the firms' overall earnings is mainly driven by the difference in the fixed amount of 20 for the worker and 360 for the firm. Ignoring these, the average additional earnings the firm makes are smaller than those of the worker in treatment BC. In the two trilateral treatments the additional earnings are higher for the firm than for the worker.

Table 5: Average earnings by treatment and role (in experimental points)

		BC	TC0	TC25	BCexo
Workers	Endowment	20	20	20	20
	Overall earnings	78.71	82.97	65.54	80.5
Managers	Endowment	360	0	90	
	Overall earnings	403.06	0	109.11	
Shareholders (Owners in BCexo)	Endowment		360	270	360
	Overall earnings		432.07	327.32	427.5

Remark: Earnings in points. The conversion rate is 30 points = 1 euro. Apart from these points, subjects received a show up fee of 12 euros.

Table 6 reports what firms could have earned by choosing the 'optimal' wage given the (unknown) actual wage-effort schedules of the workers. The profit maximizing wage equals 100 in treatments BC and TC25, 90 in treatment TC0 and zero in BCexo.

Comparing Tables 4 and 6 it can be seen that actual earnings are higher than minimal in all treatments. Also, firm’s potential earnings are higher when ownership and control are separated, but the difference between the actual and maximal earnings is lower in the BC treatment.

Table 6: Optimal wage and potential earnings

	Optimal wage	max Π	min Π	$\Pi_{\{w=0\}}$
BC	100	417.42	401.45	403.87
TC25	90	450.69	399.14	402.76
TC0	100	467.14	383.57	400.00
BCexo	0 [†]	456.00	409.33	456.00

Remark: † Optimal random wage.

5 Conclusion

In this paper we address the following question: does the separation of ownership and control between shareholders and managers that characterizes modern corporations influence workers’ reciprocal attitudes? We do so by investigating the robustness of experimental gift exchange in a three-tier hierarchy, where a manager controls the worker’s wage but is not full residual claimant. She therefore does not get the full benefits generated by the worker’s reciprocal reaction (if at all). Our experimental design covers four treatments. The standard bilateral gift exchange game, with a single owner-manager and a worker, serves as baseline (BC). This treatment is compared with two other treatments in which the firm is owned by a shareholder but controlled by a manager. In these trilateral treatments the manager chooses the worker’s wage. The manager herself either receives a fixed wage such that the shareholder is full residual claimant of the firm’s profits (treatment TC0), or she is paid on the basis of performance pay receiving 25% of firm’s profits (treatment TC25). In our final fourth treatment (BCexo) the firm has no control and the wage is chosen by a random device. This treatment is again bilateral in nature, as there is only an owner but no manager.

To exclude that efficiency acts as a confounding factor when comparing different treatments, the firm's profit function is kept constant across treatments. It is chosen such that (like in reality) the firm is always much richer than the worker. Compared to previous experiments we also increased the marginal return to effort. Our parameter choices allow us to exclude inequality-aversion and efficiency concerns (quasi-maximin preferences) as possible explanations for a positively sloped wage-effort relationship. Other characteristic features of our design are that we employ the strategy method and that only a single period is played.

We find that workers on average exhibit a reciprocal attitude in all treatments. The higher the wage paid, the higher the effort level the firm receives on average. Somewhat surprisingly, this also holds in treatment TC0 where the manager in charge of choosing the wage does not benefit from the worker's reciprocal behavior at all. The wage-effort relationship does not differ between the baseline treatment and the two trilateral treatments. In our experiment gift exchange thus appears to be robust to the separation of ownership and control. Comparison with the results for the treatment in which the wage is exogenously given indicates that both ownership and control are important for the reciprocal relationship. This suggests that the worker acts reciprocally towards the firm as a whole and not to one of its members in particular.

Another important finding is that we observe a large heterogeneity in workers' reciprocal attitudes. In the treatments where the wage is endogenously chosen the proportion of reciprocal workers figures around 50%. The remaining half of the workers always provides the minimum effort. Given our parameterization, the latter behavior is in line with both selfish and inequality-averse preferences. In the treatment where the wage is exogenously given, the fraction of reciprocal types is lower. Here we also observe that around 12% of the workers are strongly motivated by efficiency concerns. These individual differences are important to keep in mind when interpreting the average wage-effort relationship observed.

Overall our findings tentatively suggest that a worker's motivation to reciprocate does not vanish (or does not even diminish) in firms with diffuse ownership stock.

Workers may reciprocate a kind wage offer with higher effort and/or an unkind wage offer with lower effort, even if it implies that a large part (or all) of the benefits goes to an anonymous, impersonal group of wealthy shareholders. To further investigate whether this is indeed the case, the next step is to explore the external validity of our results in the field.

References

- Abeler, J., Altmann, S., Kube, S., and Wibral, M. (2006). Reciprocity and payment schemes: When equality is unfair. IZA discussion paper. No. 2500.
- Akerlof, G. (1982). Labor contracts as partial gift-exchange. *Quarterly Journal of Economics*, (97):543–69.
- Bellemare, C. and Shearer, B. (2007). Gift exchange within a firm: evidence from a field experiment. IZA discussion paper. No. 2696.
- Bereby-Meyer, Y. and Niederle, M. (2005). Fairness in bargaining. *Journal of Economic Behavior and Organization*, (56):173–186.
- Blount, S. (1995). When social outcomes aren't fair: The effect of causal attributions on preferences. *Organizational Behavior and Human Decision Process*, (63):131–144.
- Bolton, G. and Ockenfels, A. (2000). ERC: A theory of equity, reciprocity and competition. *American Economic Review*, (90):166–193.
- Bosch-Domènech, A. and Silvestre, J. (2005). Elicitation in the lab. does embedding matter? Working Paper.
- Brandts, J. and Charness, G. (2000). Hot vs. cold: Sequential responses and preference stability in experimental games. *Experimental Economics*, 2:227–238.
- Brosig, J., Weimann, J., and Yang, C. (2003). The hot versus cold effect in a simple bargaining experiment. *Experimental Economics*, (6):75–90.

- Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- Cason, T. and Mui, V.-L. (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 Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, (74):1579–1601.
- Charness, G. and Kuhn, P. (2006). Does pay inequality affect worker effort? experimental evidence. Working Paper, University of California.
- Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, (117):817–869.
- Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior*, (47):268–298.
- Engelmann, D. and Ortmann, A. (2002). The robustness of gift exchange: A reconsideration. Working Paper.
- Engelmann, D. and Strobel, M. (2004). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review*, (94):857–869.
- Falk, A. (2007). Gift exchange in the field. *Econometrica*, forthcoming.
- Falk, A., Fehr, E., and Fischbacher, U. (2000). Testing theories of fairness - Intentions matter. Working Paper No. 63, University of Zurich.

- Falk, A., Fehr, E., and Fischbacher, U. (2005). Driving forces behind informal sanctions. *Econometrica*, 73:2017–2030.
- Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, 54:293–315.
- Falk, A. and Kosfeld, M. (2006). The hidden costs of control. *American Economic Review*, (96):1611–1630.
- Fehr, E. and Falk, A. (1997). Reciprocity in experimental markets. Forthcoming in: C.R. Plott and V.L. Smith (Eds.), Handbook of experimental economics results.
- Fehr, E., Fischbacher, U., and Tougareva, E. (2002). Do high stakes and competition undermine fairness? Evidence from Russia. Working Paper no. 120, University of Zurich.
- Fehr, E. and Gächter, S. (2000). Fairness and retaliation: The economics of reciprocity. *Journal of Economic Perspectives*, (14):159–181.
- Fehr, E., Kirchler, E., Weichbold, A., and Gächter, S. (1998). When social norms overpower competition: Gift exchange in experimental labor markets. *Journal of Labor Economics*, (16):324–351.
- Fehr, E. and Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, (114):817–868.
- Fershtman, C. and Gneezy, U. (2001). Strategic delegation: An experiment. *Rand Journal of Economics*, (32):352–368.
- Fischbacher, U. (1999). z-tree. Zurich toolbox for ready-made economic experiments. *Experimental Economics*, forthcoming.
- Gneezy, U. and List, J. (2006). Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments. *Econometrica*, (74):1365–1384.

- Güth, W., Huck, S., and Müller, W. (2001). The relevance of equal splits in ultimatum games. *Games and Economic Behavior*, (37):161–169.
- Güth, W. and van Damme, E. (1998). Information, strategic behavior and fairness in ultimatum bargaining: An experimental study. *Journal of Mathematical Psychology*, (42):227–247.
- Hannan, R., Kagel, J., and Moser, D. (2002). Partial gift exchange in an experimental labor market: Impact of subject population differences. *Journal of Labor Economics*, (20):923–951.
- Kube, S., Maréchal, M., and Puppe, C. (2006). Putting reciprocity to work: Positive versus negative responses in the field. University of St. Gallen Discussion paper.
- Maximiano, S., Sloof, R., and Sonnemans, J. (2006). Gift exchange in a multi-worker firm. *Economic Journal*, forthcoming.
- Oxoby, R. and McLeish, K. (2004). Sequential decision strategy vector methods in ultimatum bargaining: Evidence on the strength of other-regarding behavior. *Economic Letters*, 84:300–305.
- Sonnemans, J. (2000). Decisions and strategies in a sequential search experiment. *Journal of Economic Psychology*, 21:91–102.

Appendices

A. Summary of the instructions for treatment TC25

In this experiment you are taking part in a study of the labor market. There are three types of participants: shareholders, managers and workers. One third of the participants will be assigned the role of shareholder; one third the role of manager and the remaining one third 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 two other participants. All groups will be composed of three participants: one shareholder, one manager and one worker. (Each group represents a firm.) 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 In this stage you (may) have to make a decision without knowing the choices of those with whom you are matched. If you are a shareholder, you actually do not have to make a decision at all. If you are a manager, you have to set the wage of the worker. This wage should be a multiple of 5 and in between 0 and 100. Each manager is allowed to set only one wage.

If you are a worker, you have to decide which effort level you want to provide for each possible wage set by the manager. 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 the others. If you are a shareholder, you will learn the wage chosen by the manager and the corresponding effort level of the worker. If you are a manager, you will learn the worker's effort choice for the wage that you offer. And if you are a worker, you will learn the wage chosen by your manager. You will also be informed about your period earnings and those of the two participants with whom you are matched. These period earnings (in number of points) are calculated as follows:

$$\text{Shareholder's period earnings} = \frac{3}{4} * (40 * \text{Effort level of worker} - \text{Wage offered} + 360)$$

$$\text{Manager's period earnings} = \frac{1}{4} * (40 * \text{Effort level of worker} - \text{Wage offered} + 360)$$

$$\text{Worker's period earnings} = \text{Wage} - \text{Cost of effort provided} + 20$$

Note that the shareholder earns 75% of the firm’s total profits, the manager earns 25%.

At the end of the experiment the period earnings will be converted into euros at the rate of 30 points = 1 euro. In addition to the period earnings earned in the experiment you will receive a show up fee of 12 euros.

B. Tables and figures

Table 1: Estimated coefficients and rank-sum test tests (p-values)

	coef. estimates	b_{BC}	b_{TC25}	b_{TC0}	b_{BCexo}
b_{BC}	0.0283	-	0.862 ^a	0.842 ^a	0.109 ^a
b_{TC25}	0.0390	-	-	0.513 ^a	0.013 ^a
b_{TC0}	0.0378	-	-	-	0.014 ^b
b_{BCexo}	0.0118	-	-	-	-
b_{end}	0.0348	-	-	-	0.009 ^a

^aOne-tailed test: H1: the average rank of “row treatment” is higher than the average rank of “column treatment”; ^bTwo-tailed test: H1: average rank of “row treatment” differs from the average rank of “column treatment”.

Figure 1: Kernel density plot of individual wage-effort slopes by treatment

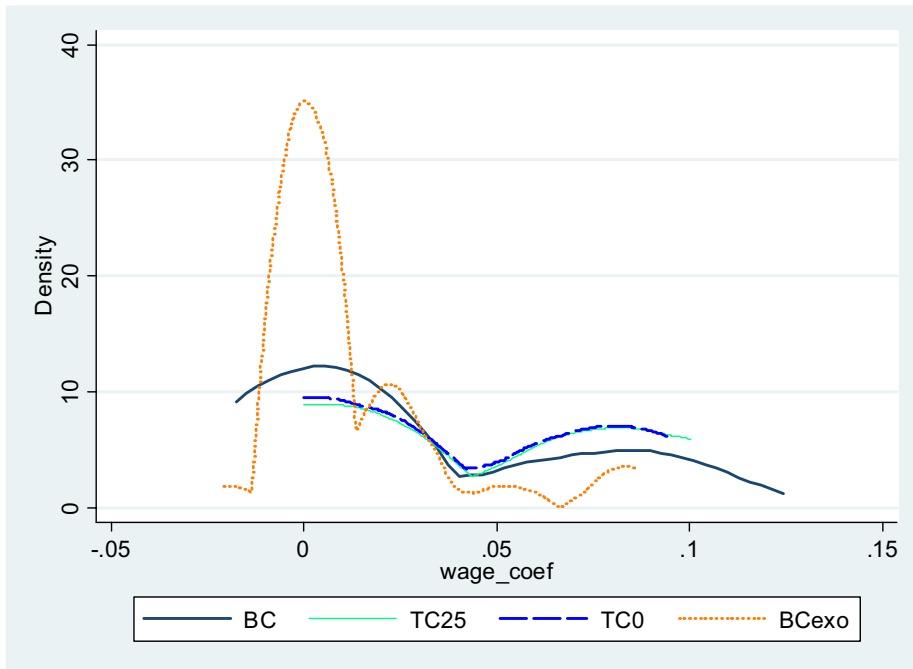


Table 2: Distribution of individual slopes, Kolmogorov-Smirnov test (p -values)

	b_{BC}	b_{TC25}	b_{TC0}	b_{BCexo}
b_{BC}	-	0.391	0.493	0.048
b_{TC25}	-	-	0.904	0.001
b_{TC0}	-	-	-	0.000
b_{BCexo}	-	-	-	-
b_{end}	-	-	-	0.020