Economica (2007) **74**, 259–297
doi:10.1111/j.1468-0335.2006.00537.x

Promotion Rules and Skill Acquisition: An Experimental Study

By HESSEL OOSTERBEEK,† RANDOLPH SLOOF† and JOEP SONNEMANS†
†*University of Amsterdam and Tinbergen Institute*

Final version received 2 December 2005.

Standard economic theory identifies a trade-off between up-or-stay and up-or-out promotion rules. Up-or-stay never wastes the skills of those not promoted but may provide insufficient incentives to invest in skills. Up-or-out can always induce investment in skill acquisition but may waste the skills of those not promoted. The paper reports an experiment designed to study this trade-off. Under up-or-out, parties behave almost exactly as theory predicts. But under up-or-stay (and stay-or-stay), results differ markedly from theoretical predictions. In that case workers invest rather frequently, although the prediction is that they would not. These deviations can be explained by various reciprocity mechanisms.

> Promotions serve two roles in an organization. First, they help assign people to the roles where they can best contribute to the organization's performance and success. Second, promotions serve as incentives and rewards. These conceptually distinct roles are sometimes in conflict... (Milgrom and Roberts 1992, p. 364)

INTRODUCTION

In his survey article on incentives in organizations, Gibbons (1998) identifies a trade-off between up-or-stay promotion rules as formalized by Prendergast (1993), and up-or-out promotion rules as described by Kahn and Huberman (1988):

> The up-or-stay rule induces the worker to invest in skills only if jobs at different levels are sufficiently different in how they utilize skill, whereas the up-or-out rule can induce investment in training even if training yields identical productivity increases in all jobs. On the other hand, the up-or-stay rule never wastes the acquired skills of those not promoted, whereas an up-or-out rule has this problem in many settings. (Gibbons 1998, p. 127)

To illustrate, consider the following example in which a worker and a firm interact during two periods. In the first period the worker has the opportunity to make a relationship-specific investment at a cost of 25. This investment increases the probability that the worker is of high productivity in the second period. Without investment this probability equals $\frac{1}{4}$; with investment it is increased to $\frac{3}{4}$. Before the second period starts, the firm has to assign the worker to one of two job levels: easy or difficult. Alternatively, the firm may decide to dismiss the worker. After the investment decision, but before the firm makes its assignment decision, the worker's actual productivity level is revealed to both parties. If the worker is of low productivity, he produces 100 when the firm assigns him to an easy job and 0 if assigned to a difficult job. If the worker turns out to be of high productivity, he produces 175 if he is assigned to the easy job and 220 in the difficult job. Independent of the productivity level within the firm, the worker's outside productivity equals 0. (For ease of presentation, these productivity levels are summarized in Table 1.) In this situation efficiency requires that low-productivity workers are assigned to the easy jobs and high-productivity workers to the difficult jobs, and that the worker invests, since the expected joint gain of $(\frac{3}{4}-\frac{1}{4})$ (220–100) = 60 exceeds the investment cost of 25.

Assume an environment in which parties cannot contract upon the worker's investment decision. The firm can, however, attach wages to different job levels and can commit to these wages. Assume further that the firm specifies a wage contract in which the worker earns 110 in either job. Table 2 then gives the gross payoffs for the two parties depending on the worker's level of productivity and the firm's assignment decision.

When the worker is of high productivity, the firm will promote him to the difficult job. In that job the worker produces 220, earning a wage of 110, and both parties have a payoff equal to 110 (ignoring the costs of possible investment). When the worker is of low productivity, the firm is best off when the worker is dismissed (since 0 exceeds -10 and -110). A worker who anticipates this assignment rule will decide to invest. Investment increases his probability of being of high productivity from $\frac{1}{4}$ to $\frac{3}{4}$ to, and a high productivity worker earns 110 more than a low productivity worker. Hence the expected wage increase generated by the investment is 55, which exceeds the investment cost of 25. This situation resembles the up-or-out rule: the worker is induced to invest, but at the same time skills are wasted when low-productivity workers are laid off rather than kept in the easy job.

Next, consider the situation in which the wage contract specifies a wage of 110 for the difficult job and 70 for the easy job. Table 3 gives the resulting gross payoffs. With these wage levels, the firm will assign a low-productivity worker to the easy job and a high-productivity worker to the difficult job. Given this assignment rule, it is in the worker's best interest not to invest. Investing results

TABLE 1
SURPLUS GENERATED BY WORKER TYPE

	Low productivity	High productivity
Difficult	0	220
Easy	100	175
Out	0	0

TABLE 2
PAYOFFS WHEN WAGES IN BOTH JOBS EQUAL 110 (UP-OR-OUT)

	Low productivity		High productivity	
	Firm	Worker	Firm	Worker
Difficult	–110	110	110	110
Easy	–10	110	65	110
Out	0	0	0	0

TABLE 3
PAYOFFS WHEN WAGE IN DIFFICULT JOB EQUALS 110 AND IN EASY JOB 70 (UP-OR-STAY)

	Low productivity		High productivity	
	Firm	Worker	Firm	Worker
Difficult	–110	110	110	110
Easy	30	70	105	70
Out	0	0	0	0

TABLE 4
PAYOFFS WHEN WAGE IN DIFFICULT JOB EQUALS 110 and in Easy Job 50 (STAY-OR-STAY)

	Low productivity		High productivity	
	Firm	Worker	Firm	Worker
Difficult	–110	110	110	110
Easy	50	50	125	50
Out	0	0	0	0

in an expected gain of $\frac{1}{2}(110–70) = 20$, which falls short of the investment cost of 25. This situation resembles the up-or-stay rule: assignment is efficient, but the worker lacks the proper incentives to invest.

Finally, consider the situation in which the wage contract stipulates a wage of 110 for the difficult job and 50 for the easy job. The resulting gross payoffs for the two parties conditional on the worker's productivity level and the firm's assignment decision are given in Table 4. With these wage levels the firm will always offer the worker the easy job, independent of the worker's productivity level. This contract is therefore dubbed the stay-or-stay rule. The worker who anticipates that he will never be promoted abstains from investment, as the return equals zero. Consequently, this contract is worse than the other two, because now both the investment decision and the assignment rule are inefficient. When the worker does, however, believe that he will be promoted when he is of high productivity, he will make the investment, because $\frac{1}{2}(110–50) = 30$ exceeds the investment costs of 25.

None of the three wage contracts introduced above achieves the first best. Under the up-or-out rule, efficient investment is achieved but the assignment decision is inefficient. Under the up-or-stay rule, assignment is efficient but the investment decision is inefficient. Finally, under the stay-or-stay rule, both the assignment decision and the investment decision are inefficient.

These outcomes are determined by the choice of the wage levels. For the given values of the exogenous variables, however, it is impossible to specify fixed wages for the two jobs such that both the investment decision and the assignment rule are efficient (see Section I). This means that in this example there is necessarily a trade-off between efficient assignment and efficient investment. This trade-off parallels the choice between wage contracts that *de facto* stipulate an up-or-out rule and an up-or-stay rule.[1] The stay-or-stay rule performs worse than the other two because both decisions deviate from the efficient choices.

Whether the three promotion rules work as theory predicts, and whether there does indeed exist a trade-off between the two types of efficiency, is an empirical issue. Given the very nature of the issues involved, however, it seems difficult to gather field data to investigate this. First, the models are motivated by the fact that skill acquisition is not verifiable by a third party, such as a court. But if a court cannot verify the acquisition of skills, why should the empirical researcher be able to do so? Second, even if it is possible for a researcher to observe the skill level of a worker in a specific job, it would be much harder to determine what the skill level of that same worker would have been in the job to which he was not assigned, or what he would have produced if employed elsewhere.

With these kinds of data problems hampering the test of a theory, laboratory experiments offer an attractive alternative. In experiments, almost everything is under the control of the researcher. More specifically, the worker's investment decision is observed and the researcher sets the payoffs of different alternatives.

This paper reports on a series of experiments designed to consider investment inefficiency and job assignment inefficiency under the three wage contracts. Our results lend support to the predicted (comparative statics) differences between the three contracts: (i) investment rates are significantly higher under up-or-out than under up-or-stay and stay-or-stay; (ii) inefficient dismissals are significantly more likely under up-or-out than under the other two contracts; and (iii) (efficient) promotions are significantly less likely under stay-or-stay. With respect to point predictions, up-or-out contracts perform by and large as theory predicts; that is, there is almost no investment inefficiency, and job assignments are typically inefficient. Up-or-stay contracts perform better than predicted; the investment inefficiency is much smaller than theory predicts, while assignment efficiency is only slightly below the predicted level. Stay-or-stay contracts also perform better than predicted; here too the investment inefficiency is much smaller than predicted. Moreover, many high-productivity workers are promoted to the difficult job, such that assignment inefficiency is also smaller than predicted.

The deviations from the theoretical predictions under up-or-stay and stay-or-stay can be explained by various reciprocity mechanisms and are consistent with the notion that reciprocity can serve as an informal contract enforcement device (cf. Fehr *et al.* 1997). The term 'reciprocity' refers to the notion that individuals are prepared to forgo some money in order to punish behaviour that is considered 'unkind' and to reward behaviour that is considered 'kind'. Under the up-or-stay rule, a firm may consider no investment by the worker as an unkind action and may therefore punish this by assigning the worker to a lower job level. A worker who anticipates this type of negative reciprocity may decide to invest, thereby achieving the first-best outcome. Similarly, under the stay-or-stay rule the firm may consider investment by the worker as a kind action, which can be rewarded by promoting high-productivity workers to the difficult job. Alternatively, a high-productivity worker can interpret assignment to the easy job as an unkind decision, which can be punished by not accepting this job offer. Anticipating this reciprocal reaction, the firm may prefer to promote the high-productivity worker after all. To distinguish empirically between the two different reciprocity mechanisms that may operate under stay-or-stay, we conducted an extra treatment for this contract in which the firm does not observe the worker's investment choice. When this is the case, it becomes more difficult for the firm to reward investment. From this extra treatment, we conclude that anticipated negative reciprocity is the stronger mechanism under stay-or-stay.

The remainder of this paper is organized as follows. The next section describes the model in more detail, showing more generally the conditions under which parties inescapably face a trade-off between investment and assignment efficiency. This should convince readers that such a trade-off occurs not only in a constructed example like the one presented above. Section II describes the design of the experiment and the theoretical predictions. Section III presents and discusses our results. Section IV summarizes the findings from the additional treatment with unobservable investment, and Section V concludes.

I. The Model

The model introduced in the previous section in the form of a numerical example contains eight exogenous variables. These are:

1. y_D^H: the value of a high productivity worker in the difficult job
2. y_E^H: the value of a high productivity worker in the easy job
3. y_D^L: the value of a low productivity worker in the difficult job
4. y_E^L: the value of a low productivity worker in the easy job
5. p_I: the probability that a worker who made the investment (i.e. a trained worker) is of high productivity;
6. p_N: the probability that a worker who did not invest (untrained worker) is of high productivity
7. c: the cost of training
8. r: the worker's outside value

Timing is as follows. First, the firm chooses the values of the wage levels w_D for the difficult job and w_E for the easy job. We assume that $w_D \geqslant w_E$. Next, the worker decides whether to invest and spend c, or not to invest and spend 0. The firm observes this decision. Subsequently nature determines the worker's actual productivity level, which is revealed to both parties. If the worker invested the probability that the worker is of high productivity equals p_I, otherwise it equals p_N. After having observed the worker's productivity, the firm offers the worker one of the two inside positions or lays her off. Finally, the worker decides whether to accept the firm's job offer (if any) or to quit.

Table 5 generalizes Tables 2–4 from the Introduction. Following Prendergast (1993), we assume that

$$(1) \qquad y_D^L < y_E^L < y_E^H < y_D^H.$$

TABLE 5
GROSS PAYOFFS FOR FIRM AND WORKER

	Low productivity		High productivity	
	Firm	Worker	Firm	Worker
Difficult	$y_D^L - w_D$	w_D	$y_D^H - w_D$	w_D
Easy	$y_E^L - w_E$	w_E	$y_E^H - w_E$	w_E
Out	0	r	0	r

Following Kahn and Huberman (1988), we also assume that

$$(2) \qquad y_E^L > r.$$

These assumptions imply that assignment is efficient when low-productivity workers are assigned to the easy job and high-productivity workers to the difficult job. Furthermore, we assume that training is efficient, which is the case if

$$(3) \qquad (p_I - p_N)\left(y_D^H - y_E^L\right) > c.$$

The first-best is achieved when the parties' private incentives are aligned with the efficient outcomes. The firm's assignment rule will be efficient if $y_E^L - w_E > 0$ and $y_D^H - w_D > y_E^H - w_E$. The worker invests and stays with the firm if $(p_I - p_N)(w_D - w_E) > c$ and $w_E > r$.

It is now straightforward to show that, if

$$(4) \qquad (p_I - p_N)\left(y_D^H - y_E^H\right) < c,$$

then there does not exist a fixed wage contract that achieves the first-best solution: the assumption that training is efficient only implies inequality (3), and since $y_E^H > y_E^L$ inequality (4) may hold. In that case, there is either investment inefficiency or assignment inefficiency. And if the wages are set at a suboptimal level, as under the stay-or-stay contract, it may even be the case that both types of inefficiency emerge.

For the remainder of the paper, we assume that the exogenous conditions are such that there is indeed a necessary trade-off between investment efficiency and assignment efficiency; that is, we assume that inequality (4) holds. We also continue to assume that conditions (1), (2) and (3) hold. Under these assumptions, a fixed wage contract is an up-or-out contract when

$$(5) \qquad y_E^L < w_E \leqslant w_D < y_D^H$$

and

$$(6) \qquad w_D - w_E < y_D^H - y_E^H.$$

The first inequality in (5) ensures that the firm will fire a low-productivity worker because his wage exceeds his highest possible value. The last inequality ensures that a high-productivity worker will be kept. Condition (6) then guarantees that this type of worker will be offered the difficult job rather than the easy one. Moreover, the worker will make the investment when $\mathrm{w_D}$ is chosen such that $(p_I - p_N)(w_D - r) > c$, which is always possible because of (2) and (3). We therefore end up with a situation in which the investment decision is efficient, but assignment is not.[2]

If instead the two wage levels are set such that condition (6) holds together with

$$(7) \qquad r < w_E < y_E^L,$$

then we effectively have an up-or-stay contract. Condition (7) guarantees that the firm will offer a low-productivity worker the easy job (and that the worker will accept this offer). As before, condition (6) ensures that the high-productivity worker is promoted to the difficult job. We thereby obtain efficient assignment. In this case, however, the worker has insufficient incentives to invest. If condition (6) is combined with inequality (4), we have that $(p_I - p_N)(w_D - w_E) < c$, implying that the worker's costs of investing exceed his expected benefits from investing.

Finally, if condition (7) is supplemented by condition

$$(8) \qquad w_D - w_E > y_D^H - y_E^H,$$

we have a stay-or-stay contract. (Note that (8) is just the opposite of (6).) Owing to condition (8), the firm will not promote a high-productivity worker to the difficult job, because the increase in productivity from doing so falls short of the associated wage raise. The worker who anticipates this non-promotion policy will not incur the cost of investment even if the wage differential is large enough to cover these costs, i.e. even if

$$(9) \qquad (p_I - p_N)(w_D - w_E) > c$$

A stay-or-stay contract that satisfies condition (9) is an interesting case to consider, because it provides an investment incentive to workers who believe that they will be promoted if they are of high productivity. This is of particular relevance under reciprocity motivations, which will be discussed in Section II.

While the values of w_D and w_E covered by the three contracts do not cover all possible combinations of w_D and w_E, it can be shown that other values cannot improve upon the up-or-out and up-or-stay contracts in terms of efficiency.[3] An improvement over the stay-or-stay contract is not relevant, because this contract is already predicted to perform worse than the other two contracts.

II. Experimental Design and Predictions

Choice of parameters and standard predictions

In the experiment we use the parameter values that appeared in the numerical examples in the Introduction. That is, for the exogenous variables we chose the following values:

$$y_D^L = 0,\ y_E^L = 100,\ y_E^H = 175,\ y_D^H = 220,\ p_I = 3/4,\ p_N = 1/4,$$
$$c = 25 \text{ and } r = 0$$

These values satisfy the exogenous conditions (1)–(4) specified in the previous section.

Rather than let the firm choose freely the values of w_D and w_E, we restrict this choice to three fixed-wage packages—the first represents an up-or-out contract; the second depicts an up-or-stay contract; whereas the third portrays the stay-or-stay contract. In all three cases the wage level in the difficult job is set equal to $w_D = 110$. This implies that when a high-productivity worker is assigned to this job the worker and the firm split the gross surplus equally. The essential difference between the three promotion rules is the wage in the easy job w_E (cf. Tables 2–4 in the Introduction).

According to conditions (5) and (6), the up-or-out rule arises when the wage level in the easy job exceeds $y_E^L = 100$ and (weakly) falls short of $w_D = 110$. To represent this rule, we chose $w_E = 110$ ($= w_D$). As we showed in the Introduction, the firm is predicted to assign high-productivity workers to the difficult job and to dismiss low-productivity workers. The backward induction prediction is then that this induces workers to invest, because $(p_I - p_N)\ (w_D - r) = \frac{1}{2}(110 - 0) = 55$ exceeds the cost of investment c ($= 25$).

For the up-or-stay contract, we need by condition (7) that $0 = r < w_E < y_E^L = 100$. Moreover, condition (6) requires that $w_E > w_D - y_D^H + y_E^H = 65$. We chose to set w_E equal to 70 for this case. With this wage level for the easy job, the firm will assign high-productivity workers to the difficult job and low-productivity workers to the easy job. Anticipating this assignment policy of the firm, workers are predicted not to invest, since $(p_I - p_N)$ $(w_D - w_E) = \frac{1}{2}(110 - 70) = 20$ falls short of the investment cost of 25.

Finally, for the stay-or-stay contract, we chose w_E equal to 50. With this value $w_D - w_E = 60$, exceeding $y_D^H - y_E^H = 45$ and thereby satisfying condition (8). Moreover, for this value $(p_I - p_N)\ (w_D - w_E) = 30$, thereby exceeding the cost of investment ($c = 25$), and representing the more interesting version of the stay-or-stay contract that fulfils condition (9).

The (subgame-perfect) equilibrium outcomes and parties' net payoffs for each of the three contracts are summarized in Table 6. For all contracts, the prediction is that workers will accept a job if offered. (Acceptance or rejection of an 'out' offer is irrelevant.)

When the firm has to choose between contracts, in equilibrium it prefers up-or-out to stay-or-stay, which in turn it prefers to up-or-stay. This is because, for the given parameter values, the expected payoff for the employer is larger with the up-or-out contract ($82\frac{1}{2}$) than with the up-or-stay (50) or stay-or-stay ($68\frac{3}{4}$) contracts.[4] Consequently, if firms (and workers) all made optimal

TABLE 6
THE CONTRACTS CONSIDERED AND THEIR THEORETICAL PREDICTIONS

	Up-or-out	Up-or-stay	Stay-or-stay	First-best
Investment	Yes	No	No	Yes
Assignment high-productivity workers	Difficult	Difficult	Easy	Difficult
Assignment low-productivity workers	Out	Easy	Easy	Easy
Net payoff to firm	$82\frac{1}{2}$	50	$68\frac{3}{4}$	165 (firm and worker combined)
Net payoff to worker	$57\frac{1}{2}$	80	50	

choices, we would never observe how subjects played under the other two contracts. Therefore our experimental design also covers treatments in which the firm cannot choose between different contracts (see below). In these treatments, the game starts with the worker making his investment decision.

Predictions based on reciprocity

The above predictions are based on the assumption that people are motivated solely by their own monetary payoffs, and they match the predictions in the relevant theoretical literature. However, a large body of earlier experimental studies strongly suggests that subjects are also motivated by considerations as altruism, equity, fairness and reciprocity. In the context of contract enforcement, the experimental results of Fehr *et al.* (1997) are particularly relevant (see also Gächter and Fehr 1998). Their results suggest that reciprocity might matter, and may lead to different predictions. *Reciprocity* entails that one is willing to forgo some money in order to punish behaviour that is considered as unfair and to reward behaviour that is considered as fair. In the words of Fehr *et al.* (1997, p. 839), '[T]he essential feature of reciprocity motives is thus a willingness to pay for responding fairly (unfairly) to a behaviour that is perceived as fair (unfair).' Owing to this willingness to pay to enforce fair behaviour, reciprocity might be capable of mitigating a contract enforcement problem. In that way, reciprocity may be efficiency enhancing.

This may also be a relevant mechanism for the contracts studied in this paper. First, consider the up-or-out rule. The firm is predicted to dismiss a low-productivity worker, irrespective of the worker's investment decision. But assume now that the firm considers an investment as a fair action that deserves a reward. Assigning a low-productivity worker who invested to the easy job rather than dismissing him can give this reward. The firm then spends 10 to give the worker a reward of 110 (see Table 2). In this case, *positive reciprocity* supports the first-best outcome. Of course, in practice subjects are heterogeneous and not all of them will behave in this way (cf. Fehr and Schmidt 2002). Therefore, we do not expect to observe full efficiency under the up-or-out contract, even when a significant fraction of firms is willing to reciprocate. First best is attained only when all firms behave in a reciprocal manner.[5]

Second, consider the up-or-stay rule. The firm is predicted to assign a low-productivity worker to the easy job, again irrespective of the worker's investment decision. But assume now that the firm considers non-investment to be an unfair action that deserves some punishment. The firm can punish a low-productivity worker who did not invest by dismissing him, instead of offering him the easy job. This costs the firm 30 and imposes a loss of 70 on the worker. A worker who anticipates this will make the investment. Now the *anticipation of negative reciprocity* supports the first-best outcome. In fact, the mechanism here is more powerful, because a 'small amount' of reciprocity is sufficient to attain full efficiency; when the probability of punishing a low-productivity worker for not investing exceeds 9.5%, the worker is better off if he makes the investment.[6] So only a small fraction of reciprocal firms is required to induce the worker to invest under this reciprocity mechanism.

Reciprocity can also play a role under the stay-or-stay rule. Here standard theory predicts that the worker does not invest and is assigned to the easy job independently of his productivity. However, if the firm thinks that investment is a kind action by the worker, it can reward him by assigning him to the difficult job when he turns out to be of high productivity. This reduces the firm's benefit by 15 (from 125 to 110) and increases the worker's wage by 60 (from 50 to 110). In this situation, the worker is better off by investing when the probability of assigning a high-productivity worker who invested to the difficult job exceeds 55.5%.[7]

As an alternative (negative) reciprocity mechanism under the stay-or-stay contract, high-productivity workers may regard no promotion as an unfair action. They have an opportunity to punish the firm for this by not accepting the easy job. This punishment reduces the firm's payoff by 125 and comes at a cost of 50 to the worker. This mechanism induces a worker to invest if a sufficiently large fraction of firms believe that a sufficiently large fraction of high-productivity workers are prepared to retaliate. In particular, a fraction of at least 12% reciprocal workers is now needed to attain full efficiency.[8]

Under the stay-or-stay rule, we thus have two different reciprocity mechanisms that support investment in skill acquisition. To discriminate between these two mechanisms, we designed an additional treatment in which firms do not observe workers' investment decisions (cf. Section IV). Observability of investment is required for firms to reward workers' investment (positive reciprocity), but is not required for high-productivity workers to punish firms for not promoting them (negative reciprocity).

We argued above in an intuitive way why reciprocity is potentially highly relevant, as it directly relates to the efficiency of particular types of contract in practice. This discussion is not based on a formal alternative theory, however, and so has the potential disadvantage that different notions (e.g. positive and negative reciprocity) are used in a somewhat ad hoc way. It is therefore worthwhile to investigate whether recent models of social preferences can capture the different reciprocity mechanisms distinguished above in a more systematic way. These models assume that people do not care solely about their own monetary payoffs, and the models can be roughly classified into two kinds (see Fehr and Schmidt 2002, and Camerer 2003, for comprehensive overviews).

First, *distributional theories* are outcome-based and assume that players care only about the final distribution of payoffs. Two leading models within this class assume that utility is piecewise linear in own and other players' monetary payoffs. In our setup, the firm's utility is then described by

$$(10)\qquad U_F(m_F, m_W) = \begin{cases} m_F - \alpha(m_W - m_F) & \text{if } m_F < m_W \\ m_F - \beta(m_F - m_W) & \text{if } m_F \geqslant m_W \end{cases}$$

Here α and β are parameters of the model reflecting the marginal rate of substitution between the firm's own material payoffs m_F and the worker's material payoffs m_W. The *inequality aversion* model of Fehr and Schmidt (1999) assumes that people dislike income inequality, especially when they are getting less than others. This amounts to assuming that $0 \leqslant \beta \leqslant \alpha$ and $\beta \leqslant 1$. *Quasi-maximin preferences* introduced by Charness and Rabin (2002) allow that people care about joint payoffs and also about the payoffs of the worst off player. These arise for $1 > \beta > -\alpha > 0$.

In principle, distributional preferences can capture our intuitive reciprocity reasonings in an indirect way, because the marginal rate of substitution (MRS) between the firm's and the worker's payoffs varies according to whether the firm is ahead or behind, and thus depends on whether or not the worker has borne investment costs. Take for instance the up-or-out rule, and let the worker be of low productivity. Assigning him to the easy job means that the worker earns more (110 minus investment costs of possibly 25) than the firm (− 10) does. Inequality aversion then predicts that the worker is dismissed. But under quasi-maximin preferences firms care about efficiency and may therefore be willing to keep the worker, even when he earns more. Moreover, because the firm also has a special concern for those who are behind, the firm is more likely to retain the worker when the worker did invest than when he did not. This is because firing the worker becomes less attractive when the worker has already borne investment costs. When

$$-\frac{1}{12}<\alpha<\frac{25\beta-10}{95}$$

the firm is willing to offer the low-productivity worker the easy job only when he has made the investment. In a similar vein, quasi-maximin preferences can yield the negative reciprocity mechanism under up-or-stay, whereas inequality-aversion cannot.[9] Preferences like those in (10) cannot capture the positive reciprocity mechanism under stay-or-stay, however. When $\beta>\frac{1}{5}$, a high-productivity worker is promoted to the difficult job, irrespective of whether or not he invested. The intuition here is that the firm is always ahead under this contract, and hence the MRS between its own and the worker's payoff does not vary with the actual investment costs borne.

The second type of social preference theories assume that people also care about *how* the distribution of payoffs came about. They model reciprocity in terms of a player's belief about the other players' intentions that drove their choices. Players have a preference for rewarding kind or fair intentions and punishing unkind or unfair ones. Hence preferences depend on beliefs about intended kindness. The models of Rabin (1993), Dufwenberg and Kirchsteiger (2004) and Falk and Fischbacher (2006) follow this approach and consider equilibria in action *and* (higher-order) beliefs. Because the equilibrium analysis is very complicated and because typically multiple equilibria exist, these models seem intractable in many applications (cf. Fehr and Schmidt 2002). Cox *et al.* (2006) therefore propose a simplification that avoids the use of beliefs. They assume distributional preferences as in (10), but make the MRS between the firm's own and others' payoffs dependent on the behaviour of others (besides other things such as status). This implies that parameters α and β in (10) become functions of whether or not the worker invested (and also of the contract that applies). However, their theory does not specify exactly how this dependence looks like. Cox *et al.* argue that the worker's actual choice should be compared with an action that is 'neutral in some appropriate sense' (p. 9), and that this reference action should be inferred from the actual data rather than specified *a priori* (although one problem is that the reference action cannot always be identified). This is certainly a valid approach, but makes the model less suitable for predictive purposes.[10]

In sum, distributional preferences *per se* cannot easily capture the various reciprocity mechanisms described (and for which we find empirical support—see Section III) in a systematic way. The intention-based reciprocity model of Cox *et al*. (2006) seems to be able to do this (in a simple manner), but it shares the problem with our intuitive arguments in that in many circumstances it is unclear whether an action will be considered as fair or unfair. For instance, should it be considered fair under an up-or-out contract when the worker invests, even though this action is in his own best interest? Alternatively, is it unfair if the worker does not invest, or is that just stupid because not investing reduces not only the firm's expected payoff, but also that of the worker? The Cox *et al*. model allows for both possibilities, and thus different reciprocity predictions are possible based on different assumptions of what is fair and what is unfair. Overall, we conclude that, while the general idea is rather straightforward, the lack of a single, clear-cut and commonly accepted definition of reciprocity (and fairness) makes predictions based on this notion somewhat arbitrary.

Treatments and sessions

Overall 280 subjects participated in the experiment. The experiment took place in a computer laboratory in which cubicles separated subjects. Subjects could communicate only by means of a computer network, and they did not know with whom they were connected. The subject pool was the undergraduate student population of the University of Amsterdam. About two-thirds of them were students in economics. They earned on average €20 in about $1\frac{1}{2}$ hours.[11] In each session there were 20 participants: 10 subjects were assigned the role of employer, the remaining 10 were assigned the role of worker. Participants kept the same role during the whole session. The roles were communicated only after the complete instructions were read and understood. This was done in order to make subjects familiar with the strategic elements of the experiment. Role-independent instructions have the advantage of giving subjects better opportunities to see things from the viewpoint of both roles and thus to understand the strategic situation better. This may enhance their awareness of the decision problem at hand.[12] Obviously, this feature of our design is not very realistic, because in reality workers and firms will typically know in advance which role they play.

The design of the experiment covers eight treatments grouped in 14 sessions. Table 7 maps sessions to treatments. Five treatments deal with the situation in which the worker's investment is observable (cf. sessions 1–12). Three of these treatments consider the three contracts in isolation. In these cases the stage in which the firm determines the wage levels is skipped. The game thus starts with the investment decision of the worker. In two other treatments the employer chooses between up-or-out and up-or-stay, and between up-or-out and stay-or-stay, respectively. The remaining three treatments relate to the private information situation in which the firm does not observe the worker's investment choice. The first of these treatments considers the up-or-out contract in isolation and the second the stay-or-stay contract in isolation, while the third includes the firm's choice between these two contracts.

TABLE 7
OVERVIEW OF THE EXPERIMENTAL SESSIONS

Session	Treatments
1 + 2	10 rounds up-or-out
	10 rounds up-or-stay
	10 rounds up-or-out versus up-or-stay
3 + 4	10 rounds up-or-stay
	10 rounds up-or-out
	10 rounds up-or-out versus up-or-stay
5 + 6	10 rounds up-or-out
	10 rounds stay-or-stay
	10 rounds up-or-out versus stay-or-stay
7 + 8	10 rounds stay-or-stay
	10 rounds up-or-out
	10 rounds up-or-out versus stay-or-stay
9 + 10	30 rounds up-or-out versus up-or-stay
11 + 12	30 rounds up-or-out versus stay-or-stay
13	10 rounds up-or-out; investment unobserved
	10 rounds stay-or-stay; investment unobserved
	10 rounds up-or-out versus stay-or-stay; investment unobserved
14	10 rounds stay-or-stay; investment unobserved
	10 rounds up-or-out; investment unobserved
	10 rounds up-or-out versus stay-or-stay; investment unobserved

In sessions 1 and 2, subjects played the up-or-out contract in isolation 10 times, followed by the up-or-stay contract in isolation 10 times, and finally the game including the choice between these two contracts 10 times. Sessions 3 and 4 are identical except that the order of the isolated up-or-out and up-or-stay contracts is reversed. Sessions 5–8 are exact copies of sessions 1–4, but now the up-or-stay contract is replaced by the stay-or-stay contract. Thus, in the first eight sessions the same subjects are involved in three different treatments. The motivation for this is twofold. First, it allows us to test differences in investment decisions between contracts on a within-subject basis. This reduces selectivity bias in the subject pool when comparing different treatments, and allows the use of more powerful sign-rank tests instead of rank sum tests. Second, this setup gives subjects sufficient experience with the contracts available before they actually have to choose between them.

In sessions 1–8 we used a rotating scheme with two groups of ten subjects. During ten rounds, each subject from one group (employers) met each subject from the other group (workers) once. This scheme preserves the one-shot nature of the game (cf. Kamecke 1997); that is, although subjects play the same game repeatedly, they play a series of one-shot games with different opponents. Subjects thus do not play a repeated game, thereby ruling out reputational considerations within a treatment. Between treatments, the rotating scheme was changed in order to establish that subjects were not matched in the same order to other subjects.[13] Subjects were informed about these features of the rotating scheme.

The firm's choice of contract in the last ten rounds of sessions 1–8 is likely to depend on the experiences during the first 20 rounds. In order to control for this, sessions 9–12 were held. These sessions consider the single treatment with an endogenous choice between contracts. Sessions 9 and 10 consider the choice between the up-or-out and up-or-stay contracts, while sessions 11 and 12 deal with the choice between the up-or-out and stay-or-stay contracts. With 30 rounds, 10 workers and 10 employers, it was not possible to ensure that subjects were not matched to each other more than once, but the rotation scheme used in these sessions guarantees that the same matchings were kept at a minimum. Once again, subjects were informed about these characteristics of the rotation scheme.

Finally, sessions 13 and 14 deal with the private information treatments. These sessions mimic sessions 1–4, the only difference being that the firm does not observe the worker's investment choice.

The experiment was computerized. Subjects started with on-screen instructions. All subjects had to answer some questions correctly before the experiment started; for example, they had to calculate the earnings of subjects for some hypothetical situations. Subjects also received a summary of the instructions on paper. (A translated copy of this handout is reproduced as Appendix A.) The instructions and the experiment were phrased neutrally; words like 'opponent', 'game' and 'player' were avoided. At the start of the first game, all subjects received a message that informed them about their role, i.e. worker or employer. After the subjects had played 30 games they filled out a short questionnaire. At the end of the experiment, the earned experimental points were exchanged for money at a rate of 1 point = 1 eurocent. Subjects were paid individually and discretely.

III. RESULTS

Summary of the findings

The main results of the experiment are summarized in Table 8, which gives for each contract the mean investment rate, the share of high productivity workers that are efficiently assigned (to the difficult job) and the share of low-productivity workers that are efficiently assigned (to the easy job). The backwards induction predictions are given in parentheses.

The results in Table 8 are pooled over sessions and over exogenous and endogenous treatments with observed investment choices. (The results from the private information sessions 13 and 14 are presented separately in Section IV.) Although the theoretical prediction is that the endogenous treatments produce information only about the up-or-out contract, as that is the firm's preferred contract, this turns out not to be the case. In 52% of cases, firms prefer the stay-or-stay contract when they have the choice between the up-or-out and the stay-or-stay contract. In 39% of cases firms prefer the up-or-stay contract to the up-or-out contract when presented with this choice.

Appendix B provides more detailed information about investment behaviour by session and endogeneity of contract. With few exceptions, there are no significant differences in investment rates between endogenous and exogenous treatments or between sessions within a treatment.[14] Appendix C provides

TABLE 8
SUMMARY OF RESULTS

	Share of efficient decisions			Statistical significance (*p*-values)		
	Up-or-out (1)	Up-or-stay (2)	Stay-or-stay (3)	Up-or-out versus up-or-stay (4)	Up-or-out versus stay-or-stay (5)	Up-or-stay versus stay-or-stay (6)
Investment	0.85 (1)	0.49 (0)	0.44 (0)	0.0000	0.0000	0.4658
				[0.0313]	[0.0313]	[0.6992]
Assignment						
High → Difficult	0.99 (1)	0.95 (1)	0.67 (0)	0.0000	0.0000	0.0000
				[0.0625]	[0.0313]	[0.0044]
Low → Easy	0.19 (0)	0.92 (1)	0.97 (1)	0.0000	0.0000	0.0000
				[0.0313]	[0.0313]	[0.2402]

Remark: For investment decisions the *p*-values without brackets are based on subject-level data (i.e. individual mean investment rates); for assignment outcomes these *p*-values are based on observation level data (i.e. using each interaction as one observation). The *p*-values in square brackets are based on session level averages. All tests are two-sided; those in columns (4) and (5) are based on within-subjects comparisons using a signed rank test, those in column (6) are based on between subjects comparisons using a ranksum test.

detailed information about assignment outcomes by session and endogeneity of contract. In this case significant differences between sessions and by endogeneity of contract occur frequently. In sessions 5 and 7, in particular, subjects behaved differently. But although the different sessions produced different results, the general picture that emerges in Table 8 remains the same when we base our inferences on results from separate sessions (see below).[15]

The results for the up-or-out contract are closest to the game-theoretic predictions. The investment rate of 0.85 is fairly close to unity: (almost) all high-productivity workers are assigned to the difficult job, and low-productivity workers are typically not retained in the easy job but instead are dismissed.[16] For the up-or-stay contract, assignment coincides almost perfectly with the predicted assignment rule, but the mean investment rate is 0.49 while the prediction is zero investment. Finally, for the stay-or-stay contract, deviations from game-theoretic predictions are the largest. The actual investment rate is 0.44 while theory predicts a zero investment rate. Moreover, 67% of the high productivity workers are promoted to the difficult job whereas it is predicted that such workers would be assigned to the easy job.[17]

The three columns on the right-hand side of Table 8 report the *p*-values of statistical tests comparing the three different promotion rules. For each individual worker, we calculate his mean investment rate separately for each promotion rule he is confronted with. An individual worker is confronted either with both the up-or-out and the up-or-stay rule, or with the up-or-out and the stay-or-stay rule. These two bilateral comparisons can thus be made on a within-subjects basis, and we then use a Wilcoxon signed rank test for matched pairs. The comparison of the up-or-stay with the stay-or-stay rule can be made only between subjects; hence in that case a rank sum test is employed. It appears that individual investment rates are significantly higher under the up-or-out rule than under the two other contracts (which do not differ in this respect). In line with theoretical predictions, up-or-out does indeed increase workers' propensity to invest. Because it can be argued that subjects' decisions within a session are interdependent and should not therefore be treated as independent, we also conducted statistical tests using average values per session as the unit of observation; *p*-values based on these—rather conservative—tests are reported in Table 8 in square brackets. They lead to exactly the same conclusion.[18]

Assignments are the outcome of two decisions: the employer's job offer and the worker's acceptance decision. For that reason they cannot be compared on an individual subjects basis. The assignment patterns under the different promotion rules are therefore statistically compared using each interaction as one observation (observation level data). We also report *p*-values from rank sum and sign rank tests based on session level means. Promotion of (high-productivity) workers is significantly more likely under the up-or-out and the up-or-stay contract than under the stay-or-stay contract.[19] At the same time, dismissal of low-productivity workers occurs significantly more often under the up-or-out rule than under the other two rules. Taking the investment and worker assignment results together, the comparative statics predictions across the three promotion rules are borne out by the data.

In terms of overall efficiency, the up-or-out contract performs slightly worse than predicted. There is some extra efficiency loss because workers sometimes do not invest. On the other hand, there is a small efficiency gain because 19% of the low-productivity workers are not laid off as predicted, but instead are assigned to the easy job. Overall, the realized surplus is on average around $130\frac{1}{2}$ under the up-or-out contract, while 140 was predicted under this contract (with 165 as the maximum social surplus).

For the up-or-stay contract there are some additional efficiency losses, because actual assignments do not exactly mimic the predicted efficient assignment rule. The associated losses are, however, small in comparison with the efficiency gain resulting from the 0.49 investment rate. In sum, the realized social surplus on average amounts to $136\frac{1}{2}$ compared with the prediction of 130.

Finally, for the stay-or-stay contract there are efficiency gains on two accounts: the 0.44 investment rate, and the 67% of high-productivity workers who are assigned to the difficult job. On the other side of the balance sheet, a small loss is due to the 3% of low-productivity workers not assigned to the easy job. Moreover there is some additional loss, which cannot be read from Table 8, attributable to high-productivity workers who end up outside the firm rather than, as predicted, in the easy job. Taken together, the realized surplus is on average almost 122, while only $118\frac{3}{4}$ was predicted.

In the next three subsections we investigate in more detail the subjects' behaviour under each of the three contracts in order to find out which mechanisms are at work. This investigation will reveal that under the up-or-stay and stay-or-stay contracts reciprocity enhances efficiency, while under the up-or-out contract there is less scope for reciprocity considerations. In the last subsection we analyse employers' choice of contract in more detail.

Behaviour under the up-or-out contract

Averaged over treatments and sessions, the mean of individual investment rates under the up-or-out contract equals 0.85. The exact values vary somewhat between sessions and treatments, but in all sessions and treatments a vast majority of workers invest.

Figure 1 provides information about the frequency distribution of individual investment rates. Under the up-or-out contract, somewhat less than one-half of all workers have an investment rate exactly equal to the game theoretical prediction of 1, and somewhat more than one-half of them have an investment rate that exceeds 0.9. From these figures, we conclude that for a majority of workers the normal thing to do under the up-or-out contract is to invest.

Assignment of workers to the different job options appears to be similar for the endogenous and exogenous treatments (cf. Appendix C). Table 9 reveals that the majority of assignments fit the equilibrium predictions. When the worker is of high productivity, he is almost always assigned the difficult job, irrespective of whether he invested or not. If the worker did not invest and turned out to be of low productivity, he is typically dismissed, as predicted.

The single contingency where theory predicts less well concerns the case where the worker is of low productivity despite his having invested. Although the majority of observations (72%) is still in line with the theoretical prediction of dismissal, in 27% of the cases the worker is offered the easy job. This assignment behaviour can be explained by positive reciprocity: the firm is

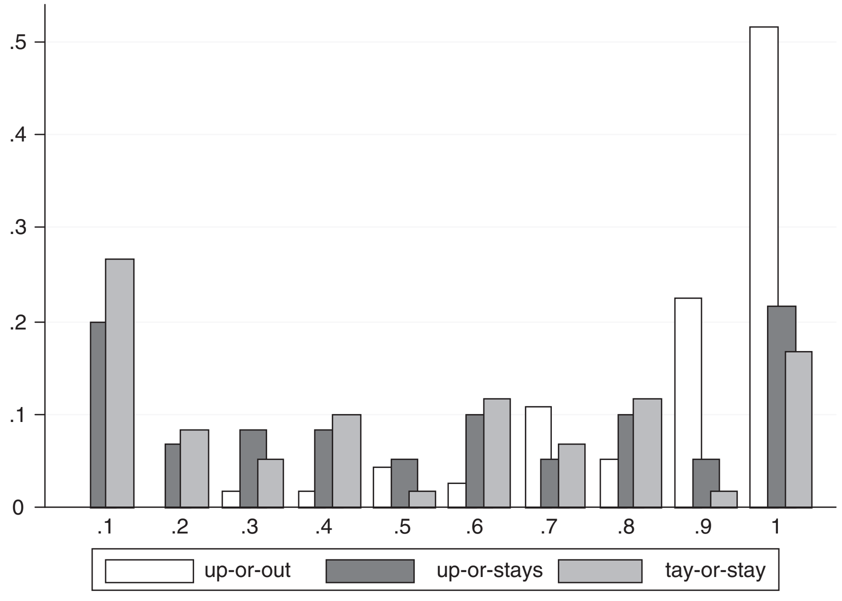

FIGURE 1. Frequencies of individual investment rates under the three contracts pooled for all sessions.

TABLE 9
WORKER ASSIGNMENT UNDER UP-OR-OUT

	No investment		Investment	
	Low	High	Low	High
Difficult	1% (2)	**97%** (67)	0% (1)	**100%** (1156)
Easy	2% (4)	0% (0)	27% (101)	0% (4)
Out	**97%** (187)	3% (2)	**72%** (267)	0% (1)

Note: Endogenous and exogenous treatments are pooled. Numbers of cases are in parentheses. Bold-faced numbers represent assignments that are in line with theoretical predictions.

sometimes willing to forgo some money in order to reward the worker for his (unsuccessful) investment. But, given that this channel to reciprocate is very cheap to the firm (it costs 10 points to give a reward of 110 points), the result that it is not observed in 72% of the cases suggests that this mechanism is not very strong. This does not necessarily imply that subjects are not motivated by positive reciprocity. It may also be that in this setting most firms consider investment not as a fair action, but as the neutral reference action. In any case, the impact of this type of reciprocity on efficiency is fairly minor, because efficiency gains increase only proportionally with the fraction of firms that are willing to give the reward.[20]

Given the actual patterns of job offers, workers' expected payoffs from investment equal almost 65 points, whereas no investment yields them an expected 29 points. Workers thus do indeed obtain strong incentives to invest when the up-or-out rule applies. Firms have no reason to deviate from this rule, because job offers are (almost) always accepted.

Behaviour under the up-or-stay contract

Averaged over sessions and treatments, the mean individual investment rate under the up-or-stay contract equals 0.49. The actual mean investment rates fluctuate slightly between sessions and treatments, but only one of the differences is significant (see Appendix B). Although mean individual investment rates are very stable over the different treatments and sessions, individual investment rates are dispersed; see Figure 1. The frequency distribution is bimodal, with a small peak near the game-theoretical prediction of zero and a slightly higher peak at the opposite end.

Under this contract, theory predicts that high-productivity workers will be offered promotion, while low-productivity workers will be offered the easy job. No worker will be dismissed and, similarly, quits are not predicted to occur. Table 10 presents the realized assignment patterns. (Quits are unimportant here; see Appendix C.)

In two of the four relevant contingencies, standard theory predicts less than perfectly. It concerns the two contingencies (low and high) in which the worker did not invest. These assignments can be interpreted as forms of negative reciprocity, where the firm punishes the worker for not making the investment. When non-investing workers appear to be of low productivity, in 9% of cases they are dismissed instead of being offered the easy job. In 14% of the cases, workers who are of high productivity despite having not invested are offered the easy job or are even dismissed. As predicted, the employer's job offer *de facto* determines assignment, because job offers are almost always accepted.

While only a small fraction of the workers who did not invest are actually punished, this fraction is large enough to motivate workers' investment behaviour. Given the actual job offer patterns, the expected payoff for the worker from investment is about 74 points, while the expected payoff from not investing is also about 74 points. The 'fictitious play' strategy against the firms' aggregate observed job offer behaviour is thus indifferent between investment and no investment. Hence investment is not irrational at all when the up-or-stay contract applies, because the worker may correctly anticipate the negative reciprocal response of (a small fraction of) employers when confronted with no investment. More generally, our results suggest that only a 'small' amount of reciprocal behaviour has 'large' (efficiency) consequences under the up-or-stay contract (cf. the previous subsection).

TABLE 10
WORKER ASSIGNMENT UNDER UP-OR-STAY

	No investment		Investment	
	Low	High	Low	High
Difficult	0% (0)	**86%** (77)	0% (0)	**98%** (280)
Easy	**91%** (279)	13% (12)	**97%** (103)	2% (6)
Out	9% (29)	1% (1)	3% (3)	0% (0)

Note: Endogenous and exogenous treatments are pooled. Numbers of cases are in parentheses. Bold-faced numbers represent assignments that are in line with theoretical predictions.

Behaviour under the stay-or-stay contract

Averaged over sessions and treatments, the mean individual investment rate under the stay-or-stay contract equals 0.44. There appear to be some significant differences between sessions, but the observed dispersion is *not* systematically related to different treatments (endogenous or exogenous—see Appendix B). Similar to up-or-stay, we find a large dispersion in investment rates across subjects, and again the frequency distribution is bimodal (see Figure 1): it has a peak at the game-theoretical prediction of zero, and a somewhat smaller peak at an investment rate of 1. High investment rates under the stay-or-stay contract suggest that workers expect that they will be promoted when they are of high productivity. Below we show that this expectation can be supported by reciprocity.

Standard theory predicts that workers are always offered the easy job, independent of their investment and productivity. No dismissals or quits are predicted to occur. In the experiment, the actual pattern of assignments appears to depend on whether the stay-or-stay rule is exogenously given or endogenously chosen. Table 11 presents the observed assignments for the exogenously given case. Since we now observe variation in workers' quit behaviour when they are offered an easy job (cf. Tables A4(a) and A4(b)), we added an extra row to the table.

Actual assignments of low-productivity workers are almost exactly in line with predicted assignments. This is not true for high-productivity workers. A substantial fraction of these workers is assigned to the difficult job rather than, as predicted, to the easy job. This assignment pattern can be supported by two different reciprocity mechanisms. First, a high-productivity worker who is not promoted may feel mistreated. He may then punish the firm by not accepting the easy job. This costs the worker 50 points and costs the firm 125 points. As can be seen in the table, this happens fairly frequently: 33 out of 85 (39%) high-productivity workers who are offered the easy job quit. Firms that anticipate this reciprocal response may want to avoid it by granting the worker promotion. Promotion yields the firm 110 points, whereas offering the easy job in these circumstances yields $(1 - 0.39) \times 125 = 76\frac{1}{4}$ points in expected payoff terms. Employers who care only about their own payoffs thus already have a strong incentive to deviate from the predicted stay-or-stay rule.

TABLE 11

WORKER ASSIGNMENT UNDER EXOGENOUSLY GIVEN STAY-OR-STAY RULE

	No investment		Investment	
	Low	High	Low	High
Difficult	0% (0)	44% (26)	0% (0)	58% (73)
Easy and accept	**94% (159)**	**39% (23)**	**100% (46)**	**23% (29)**
Easy and quit	1% (1)	17% (10)	0% (0)	18% (23)
Out	6% (10)	0% (0)	0% (0)	0% (0)

Note: Numbers of cases are in parentheses. Bold-faced numbers represent assignments that are in line with theoretical predictions.

A second (positive) reciprocity mechanism arises when the firm considers an investment by the worker as a friendly action which justifies a reward in the form of a promotion. This costs the firm 15 points and yields the worker a benefit of 60 points. This form of reciprocity obviously applies only to workers who made the investment. The 44% promotion rate among the high-productivity workers who did not invest can be unambiguously interpreted as a result of anticipated negative reciprocity. The fact that the promotion rate of high-productivity workers who did make the investment is even higher (58%) may indicate the presence of positive reciprocity. Alternatively, it may suggest that the anticipated negative reciprocal reaction of high-productivity workers if they are not promoted is stronger when they have made the investment.[21] This observation, together with the sheer percentages of 44% and 58%, suggests that anticipated negative reciprocity is stronger in impact. Section IV elaborates on this point when we present and discuss the results from the private information condition.

Table 12 presents the assignment pattern for the endogenously chosen stay-or-stay rule. Its main difference from Table 11 is that high-productivity workers are offered the difficult job much more often. In particular, the endogenous choice of the stay-or-stay rule increases their promotional chances by about 20%, independent of whether or not they made the investment. An explanation might be that firms that offer the high-productivity worker the easy job when the stay-or-stay rule is given exogenously now choose the up-or-out rule in the endogenous treatment in order to avoid the negative reciprocal reaction of workers.

As before, we can calculate the expected benefits for the worker of investment and no investment under the assumption of 'fictitious play'. For the exogenous (endogenous) treatments, the expected payoff from investment equals 51.1 (resp., 61), and the expected payoff from not investing equals 54.35 (58.8).[22] These expected payoffs of investment and no investment are fairly similar, explaining why on average workers are indifferent between investment and no investment. Note that under the stay-or-stay contract the theoretically expected payoffs differ substantially between investment and no investment (25 *v*. 50 points). In sum, under the stay-or-stay contract, also, investment is not that irrational after all, because the worker may rationally anticipate the reciprocal response of the employer when he does invest.

TABLE 12

WORKER ASSIGNMENT UNDER ENDOGENOUSLY CHOSEN STAY-OR-STAY RULE

	No investment		Investment	
	Low	High	Low	High
Difficult	0% (0)	61% (49)	0% (0)	80% (177)
Easy and accept	**99% (237)**	**21% (17)**	**100% (75)**	**10% (22)**
Easy and quit	1% (3)	18% (14)	0% (0)	10% (23)
Out	0% (1)	0% (0)	0% (0)	0% (0)

Note: Numbers of cases are in parentheses. Bold-faced numbers represent assignments that are in line with theoretical predictions.

The choice of contract

To complete this section, we now turn to the analysis of the firm's choice of contract. Contract choices are made only in the endogenous treatments. As noted before, investment behaviour does not differ significantly between exogenous and endogenous treatments. For worker assignment, we find differences only under the stay-or-stay rule.

When the choice is between the up-or-out and up-or-stay contracts, the mean individual propensity to choose the up-or-out contract (averaged over sessions) equals 61%. When the choice is between up-or-out and stay-or-stay, this propensity equals 48%. In both cases standard theory predicts 100%. The low mean rates are caused mainly by a large dispersion in the individual choice rates. For both choices (up-or-stay versus up-or-out and stay-or-stay versus up-or-out) the frequency distributions appear to be bimodal, with peaks below 25% and above 75%. Hence for both type of choices the majority of subjects choose one or the other; very few subjects tend to choose the two contracts available to them about equally.

Which factors determine the choice of contract? Firms are likely to base their choices on expected actual payoffs under the different contracts. Actual payoffs may deviate from the theoretically predicted payoffs. Table 13 provides information concerning the average net payoffs to the firm and the worker.

In line with theoretical predictions, the firm is best off under the up-or-out contract, and the worker under the up-or-stay contract. The firm's payoffs, however, are only slightly higher under up-or-out than under stay-or-stay, and the difference is not significant at the 5%-level (p-value of signed rank test equals 0.084).[23] Hence firms should be indifferent between these two contracts. That the firm chooses the stay-or-stay contract in 52% of cases rather than the up-or-out contract therefore comes as no surprise. But the firm appears to earn significantly more under the up-or-out contract than under the up-or-stay contract (signed rank test, $p = 0.015$). Therefore, the finding that in 39% of the cases the firm chooses the up-or-stay contract rather than the up-or-out contract is more difficult to explain.

One potential explanation may be the different experiences that firms have gained while using the different contracts. That is, past experiences with the contracts in isolation obtained in the exogenous treatments may play a role in the actual contract choice in the endogenous treatments. This can be tested by

TABLE 13
NET PAYOFFS

	Firm			Worker		
	Up-or-out	Up-or-stay	Stay-or-stay	Up-or-out	Up-or-stay	Stay-or-stay
First-best	80	90	95	85	75	70
Theoretical prediction	82.5	50	68.75	57.5	80	50
Observed outcome	70.9	64.0	69.2	59.6	72.5	52.6

Note: Endogenous and exogenous treatments are pooled.

regressing the individual choice rates for the up-or-stay contract on the earnings differential between the up-or-stay and up-or-out contract experienced in the exogenous treatments (including a constant term).

In sessions 1–4, subjects first played 10 rounds of the exogenous up-or-out treatment and 10 rounds of the exogenous up-or-stay treatment. They then played 10 rounds of the endogenous game with a choice between these two contracts (cf. Table 7). Regressing the frequency of firms' choices of the up-or-stay contract on the difference of their earnings in the exogenous treatments shows that the earnings difference has a significantly positive effect on the frequency of the choice for the up-or-stay contract. This partly explains the anomaly that the up-or-stay contract is often chosen.

IV. Unobservable Investment Decisions and Reciprocity

The most significant deviations from the game-theoretic predictions are found for the stay-or-stay rule. The investment rate is much higher than predicted, and high-productivity workers are often promoted to the difficult job. Moreover, when given the choice between up-or-out and stay-or-stay, the latter is chosen in 52% of cases while 0% is predicted. In Section III we explained the observed assignment pattern under stay-or-stay by referring to two different reciprocity mechanisms. The first is based on anticipated negative reciprocity; the firm anticipates that high-productivity workers who are not promoted may feel mistreated when offered the easy job and will quit. This induces the firm to offer them promotion. The second concerns positive reciprocity. The firm may want to reward workers who invested successfully. Both types of reciprocity may induce the worker to invest, and thus may explain the high investment rates observed under stay-or-stay. In the previous section we suggested that the anticipation of negative reciprocity is likely to be the stronger force. In this section we substantiate this claim more rigorously by discussing the results of an additional experiment in which there is no room for positive reciprocity as described above.

In the original model of Kahn and Huberman (1988) only the worker observes the investment level chosen. In contrast, in our basic setup workers' investment decisions are publicly observed. As Malcomson (1997, p. 1946) correctly notes, this difference does not matter theoretically.[24] Subgame perfection namely implies that worker assignment depends not on (observing) the actual investment made, but on realized productivity only. Yet the positive reciprocity mechanism described above crucially depends on the investment decision being observable to the employer. If the firm does not observe this choice, it is impossible for the firm to reward or punish the worker's investment behaviour.[25] The only way in which reciprocity then can operate is if the firm anticipates a negative reciprocal reaction of a high-productivity worker if he is not promoted. Especially high-productivity workers who do not get promoted after they invested may want to reciprocate by leaving the firm.

In two additional sessions we studied the stay-or-stay rule for the case in which the worker's investment is unobservable to the employer. To keep the experimental setup as comparable as possible, these two sessions copied sessions 1–4 (with the single difference being the information on workers'

TABLE 14(a)
WORKER ASSIGNMENT UNDER EXOGENOUS STAY-OR-STAY, PRIVATE INFORMATION TREATMENT

	No investment		Investment	
	Low	High	Low	High
Difficult	0% (0)	78% (25)	0% (0)	59% (30)
Easy and accept	**99% (98)**	**19% (6)**	**100% (18)**	**18% (9)**
Easy and quit	0% (0)	3% (1)	0% (0)	24% (12)
Out	1% (1)	0% (0)	0% (0)	0% (0)

Note: Numbers of cases are in parentheses. Bold-faced numbers represent assignments that are in line with theoretical predictions.

TABLE 14(b)
WORKER ASSIGNMENT UNDER ENDOGENOUS STAY-OR-STAY, PRIVATE INFORMATION TREATMENT

	No investment		Investment	
	Low	High	Low	High
Difficult	0% (0)	93% (13)	0% (0)	97% (35)
Easy and accept	**100% (52)**	**0% (0)**	**100% (10)**	**0% (0)**
Easy and quit	0% (0)	7% (1)	0% (0)	3% (1)
Out	0% (0)	0% (0)	0% (0)	0% (0)

Note: Number of cases are in parentheses. Bold-faced numbers represent assignments that are in line with theoretical predictions.

investment). We therefore included a (private information) treatment for the up-or-out contract, for the stay-or-stay contract and for the up-or-out versus stay-or-stay contract.

We focus on behaviour under the stay-or-stay rule.[26] Tables 14(a) and (b) list the assignments in the private information treatment (for exogenous and endogenous separately). As experimental procedures were such that it was impossible for employers to observe the investment decisions of workers, any differences in firm behaviour between no investment and investment are due to chance. With one exception out of 179 cases, the low-productivity worker is offered the easy job within the firm. The assignment of high-productivity workers again depends on whether or not the stay-stay-stay contract is endogenously chosen. In the exogenous treatments the promotion rate of high productivity workers equals 66%; in the endogenous treatments it is 96% (where averages over no investment and investment are taken). These percentages are even higher than in the public information treatments (cf. Tables 11 and 12). This is a strong indication that anticipated negative reciprocity is the important motive for the promotion decision. This anticipation is justified by the observed rejection rates of the easy job offers by high-productivity workers: 25% (2 out of 8 for exogenous and endogenous together) reject the easy job after no investment and 59% (13 out of 22) do so after investment. These percentages make it rational for firms to offer high-productivity workers promotion.

The individual mean investment rates for stay-or-stay under private information do not differ significantly from those under public information (see Appendix B). The choice between up-or-out and stay-or-stay is also unaffected by the information condition.[27] The fact that results stay the same when we reduce the scope for positive reciprocity provides additional evidence that anticipated negative reciprocity is the dominant reciprocity mechanism under stay-or-stay. A general conclusion is thus that, in practice, firms should and will take account of the reciprocal reaction of high-productivity workers who are not promoted when designing their promotion rule.

V. Conclusion

This paper reports an experiment designed to study the trade-off between investment efficiency and job assignment efficiency under different promotion rules. One of the main findings is that the comparative statics predictions for the three different promotion rules are supported by the data: (i) investment rates are significantly higher under up-or-out than under up-or-stay and stay-or-stay; (ii) inefficient dismissals are significantly more likely under up-or-out; and (iii) (efficient) promotions are significantly less likely under stay-or-stay. At the same time, another key finding is that in practice the promotion rules perform differently from what standard theory predicts. More specifically, up-or-out rules perform worse than predicted, while up-or-stay and stay-or-stay rules perform better. For the parameters chosen in our experiment, this difference in performance changes the efficiency ranking of the different promotion rules.

The differences between theoretical predictions and empirical realizations can be explained by the fact that contracts differ in the extent to which they give scope for reciprocity. In an up-or-out contract, theory predicts that the worker will invest. If he does so, and the firm considers this as the neutral reference action, there is no reason for the firm to reward the investment. Consequently there is no incentive to improve upon the predicted up-or-out assignment of workers, and hence there is no mechanism supporting the assignment of low-productivity workers to the easy job within the firm. In the experiments we find that only in a relatively small number of cases is the low-productivity worker offered an easy job, even when the worker did invest. This is despite of the fact that offering the worker such a job rather than dismissing him is a fairly cheap action for the firm. Thus, although up-or-out policies are a particular harsh way of dealing with employees, especially when they have collected firm-specific skills (Baker *et al.* 1988; Prendergast 1993, p. 533), employers in the experiment are prepared to dismiss their employees in order to secure a relatively small gain.

The up-or-stay and stay-or-stay contracts allow for reciprocity in a way that is efficiency-enhancing. The theoretical prediction is that the worker will not invest. The two contracts provide different incentives to reciprocate. In the up-or-stay contract the worker may decide to invest because he anticipates that the firm will punish non-investment by offering a less attractive job than that corresponding with his productivity (anticipated negative reciprocity). This mechanism does indeed operate. Non-investing workers *are* sometimes punished, and a substantial fraction of the workers *do* invest. Under the stay-or-stay contract, both positive and anticipated negative reciprocity can play a role. The worker may want to invest if he anticipates that the firm will reward his investment by offering him a more attractive job (positive reciprocity). On the other hand, the firm may offer the difficult job to the high-productive worker because it realizes that offering him the easy job may lead the worker to quit (anticipated negative reciprocity). This would make investment profitable for the worker. Both reciprocity mechanisms appear to be at work here, thereby increasing investment levels and also improving assignment efficiency. According to the results of an additional treatment with private information about the investment made, however, anticipated negative reciprocity seems to be the more important mechanism under stay-or-stay.

From the experimental evidence, we conclude that labour contracts that allow for reciprocity may lead to efficiency levels above those predicted in the theoretical literature. This conclusion is in line with the results of Fehr *et al.* (1997) about the role of reciprocity in the enforcement of labour contracts. The difference between the contracts studied in that paper and the ones considered here is that Fehr *et al.* designed labour contracts with the explicit aim of giving scope for reciprocity, while we have analysed labour contracts that emerge directly from the relevant theoretical literature.

Clearly, many interesting issues remain. In our experiment subjects were anonymously paired and could not communicate with each other. In reality, interactions between firms and workers are non-anonymous, communication is possible, and workers and firms interact frequently. This may have an impact on investment incentives. For example, Ellingsen and Johannesson (2004) explicitly investigate the impact of communication in a holdup setting where underinvestment is predicted. In their experiment subjects can make non-binding promises and threats. Theoretically these cheap talk messages are immaterial, yet it appears that they do mitigate underinvestment. Interestingly, allowing the investor to make threats increases investment rates to a larger extent than does allowing for non-investors to make promises, even though promises appear more credible than threats. This result nicely fits our finding that under the stay-or-stay rule (where holdup is predicted) negative reciprocity based on implicit threats is a stronger motivator for investment than is positive reciprocity based on implicit promises. It would be interesting to investigate a situation in which employers can make a cheap talk announcement about their planned assignment rule, as well as choosing a particular wage contract. This may provide a better representation of reality, where workers typically know less about a firm's wage policy than about its general employment strategy (e.g. up-or-out or a no-layoff policy). Given the finding of Ellingsen and Johannesson that subjects may have a taste for keeping their word, it seems reasonable to expect that an employer's cheap talk promise of promotion may increase workers' investment rates.

Another instrument that may provide (more) commitment in practice is reputation. In reality, the firm is a 'long-run' player with typically many workers. Although potential employees may not have direct experience with a prospective employer, they can learn a lot from the experiences of former and current employees. The firm therefore has the opportunity to build up a reputation for being reciprocal. It then seems reasonable to expect that, under repeated rather than one-shot interaction, firms are less likely to behave opportunistically. Moreover, because many labour relationships are long-term, workers may also want to establish a reputation for trustworthy behaviour. This too could be investigated experimentally.

APPENDIX A: SUMMARY OF THE INSTRUCTIONS

The experiment started with online computer instructions. In the first part of the instructions the rules of the experimental game were explained to the subjects. Subsequently the subjects were asked to answer five questions. These questions were used to establish whether the rules of the game were understood; subjects could proceed only after they had filled in the correct answers. In the third and final part of the online instructions, subjects were made familiar with the windows that they would see on their computer screen during the experiment.

In addition to the online instructions, a summary sheet of these instructions was handed out to the subjects. A direct translation of one these summary sheets is given below to provide some information on exactly how the experiment was framed to the subjects. It concerns the summary sheet for sessions 9 and 10, where only the endogenous treatments were considered.[28] In the other sessions we first explained the simpler (exogenous-treatment) games. Instructions for the endogenous treatment were given only after the two exogenous treatments had been completed.

Summary of the instructions

Below you will find a brief summary of the instructions of today's experiment. The complete experiment consists of 30 rounds. At the beginning of each round all the subjects are paired in couples. The exact pairings were already determined before the start of the experiment. Each round, you are paired with another subject. Pairings are such that during the whole experiment you will encounter another subject only three times at a maximum. Moreover, you will never be paired with the same subject in two consecutive rounds. When, if ever, you meet the same subject again is unpredictable.

One of the subjects in a pair has role A, the other has role B. You will keep the same role all the time. What exactly your role is will be determined after the instructions have been concluded.

Each round consists of four stages. In each stage, one of the subjects within a pair takes a decision. In stages 1 and 3 the subject with role A takes a decision; in stages 2 and 4 the subject with role B takes the decision. During the four stages of one round you will remain coupled to the same subject. The four subsequent stages take the following form.

Stage 1: Subject A within a pair chooses between the UPPER and LOWER tables. A's choice is communicated to B. The particular table chosen co-determines the number of points the subjects will receive at the end of the round. In the explanation of stage 3 we will elaborate further on this.

Stage 2: Subject B within a pair chooses between DISK 1 and DISK 2. If subject B chooses disk 1, this will cost him/her 25 points. If subject B chooses disk 2, this does not cost anything. B's choice for a particular disk is communicated to A. Subsequently, the disk is spun round. When it comes to a stop, it will point at a colour—either blue or yellow. The colour indicated by the disk is communicated to both subjects within a pair. With disk 1 the probability that it will point at yellow is 75%, whereas the probability of its pointing at blue is 25%. With disk 2 these probabilities are reversed: there is a 75% possibility of obtaining blue and a 25% chance of obtaining yellow. The indicated

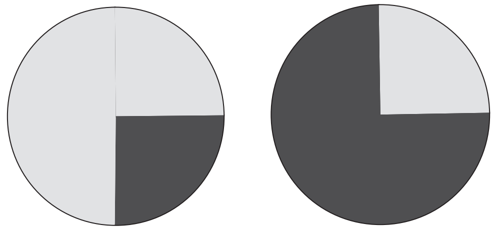

FIGURE A1. The two disks between which subject B can choose.

colour co-determines the number of points the subjects receive at the end of the round. Disks 1 and 2 are shown in Figure A1; the costs associated with each that subject B has to bear are given in parentheses.

Stage 3: Subject A makes a PROPOSAL to subject B. A may choose between three proposals, X, Y and Z. The tables printed on the additional sheet that is handed out reflect the number of points each subject earns according to the proposal. The number of points received depends on: the table (upper, lower) chosen in stage 1, the proposal (X, Y, Z) and the colour indicated by the disk (blue, yellow). Please note that with respect to subject B the tables reflect the GROSS EARNINGS. If subject B chooses disk 1 at stage 2, the earnings reflected in the tables for subject B have to be reduced by 25 points.

Stage 4: In this stage subject B decides whether to ACCEPT or REJECT the proposal made by subject A in stage 3. If the proposal is accepted, the earnings equal the corresponding number of points reflected in the table in question. If the proposal is rejected, the earnings for both subjects are zero. Once again, these figures represent gross earnings, which for subject B have to be reduced by 25 points in case he/she has chosen disk 1.

One POINT in the experiment corresponds with one EUROCENT in money. We will use the rounded exchange rate of the Euro, i.e. 100 eurocent (€1) equals Fl. 2.20. At the end of the experiment you will be paid in guilders, based on the total number of points you have earned.

APPENDIX B: MEAN INDIVIDUAL INVESTMENT RATES BY SESSION, TREATMENT, CONTRACT AND PERIOD

For each session, Table A1 reports the mean individual investment rate under the three contracts, for the exogenous (exo) and endogenous (endo) treatments separately.

Statistical tests are performed in the following way. For each individual worker, we calculated his mean investment rate separately for each treatment and contract he is confronted with. These individual mean investment rates were then compared across sessions, treatments, contracts and periods. We always made bilateral comparisons. When this comparison concerns different individual subjects, a Mann–Whitney rank sum test was employed, whereby the same individuals were compared using a Wilcoxon signed rank test for matched pairs. We employed a significance level of 5% throughout.

We first discuss the results of row-wise comparisons in Table A1. These comparisons include all periods; learning effects are discussed below. Of the 45 bilateral comparisons that can be made for the exogenous sessions of the up-or-out contract (first row), we do not find any significant differences. Of the 91 pair wise comparisons for the endogenous sessions (second row), 14 point to a significant difference, 9 of which are due to session 8 in which the mean investment rate coincides with the theoretical prediction of 1. For up-or-stay we do not find significant differences between the

Table A1
Mean of Individual Investment Rates per Session

Session		Public information												Private	
		1	2	3	4	5	6	7	8	9	10	11	12	13	14
Up-or-out	Exo.	0.83	0.77	0.88	0.72	0.73	0.85	0.81	0.89					0.82	0.86
	Endo.	0.81	0.81	0.80	0.83	0.95	0.98	0.90	1.00	0.76	0.96	0.90	0.85	0.85	0.91
Up-or-stay	Exo.	0.44	0.54	0.52	0.53										
	Endo.	0.50	0.39	0.50	0.75					0.47	0.41				
Stay-or-stay	Exo.					0.35	0.31	0.52	0.53					0.34	0.35
	Endo.					0.27	0.15	0.52	0.44			0.31	0.68	0.46	0.36

TABLE A2
LEARNING EFFECTS IN INVESTMENT

		1st half	2nd half	*p*-value
Up-or-out	Exogenous	0.80	0.84	0.0634
	Endogenous	0.88	0.87	0.9076
Up-or-stay	Exogenous	0.49	0.53	0.0800
	Endogenous	0.53	0.48	0.5924
Stay-or-stay	Exogenous	0.44	0.36	0.1762
	Endogenous	0.39	0.42	0.4624
Up-or-out private info	Exogenous	0.76	0.88	0.2430
	Endogenous	0.90	0.86	0.1437
Stay-or-stay private info	Exogenous	0.40	0.28	0.1229
	Endogenous	0.35	0.45	0.1654

different exogenous sessions (third row; six comparisons). Between the different endogenous sessions (fourth row; 15 comparisons) only the difference between sessions 4 and 9 is significant. The 15 comparisons that can be made in the fifth row do not reveal any significant differences. Finally, for the endogenous sessions of the stay-or-stay contract (last row), four out of 28 differences are significant: between sessions 5 and 12, 6 and 11, 6 and 12, and 11 and 12. Individual investment rates appear abnormally high in session 12. Recall that the setups of sessions 11 and 12 are identical (cf. Table 7). It thus seems that session 12 is somewhat of an outlier.

The results from the row-wise comparisons justify the following two conclusions. First, no important subject pool and order effects can be detected in the data. (Recall that sessions 1–4, and 5–8 differ in the order of treatments; see Table 7.) Second, the introduction of private information about the investment made does not affect investment rates under up-or-out and stay-or-stay.

Next, we consider comparisons within columns for the same type of contract. In this way we can test whether endogeneity of the contract matters for investment behaviour. These comparisons are made on a within-subjects basis, so here signed rank tests are employed. Only 3 out of the 20 pair-wise comparisons yield a significant difference: in session 5, 6 and 8 individual investment rates are significantly smaller when the up-or-out contract is exogenously given than when it is endogenously chosen. Overall, however, we conclude that investment behaviour is very much the same in the exogenous and the endogenous treatments. For the up-or-out and stay-or-stay contracts this was also tested using session level means. This results in the same conclusion.

We finally look at learning effects. Table A2 breaks the investment rates down by first and second halves of the sessions. For the exogenous treatments, the investment level of the first (second) half is based on the first (second) five periods in which subjects were confronted with the particular treatment. For the endogenous treatments, the investment of the first (second) half is based in the first (second) five/fifteen periods in which subjects played the endogenous treatment (five periods for sessions 1–8 and 13 and 14, fifteen periods for sessions 5–8). The results in the table show that there are no significant differences between first and second halves. We therefore conclude that learning effects are unimportant with respect to investment behaviour.

APPENDIX C: WORKER ASSIGNMENT BY CONTRACT, INVESTMENT DECISION AND PRODUCTIVITY

Tables A3(a) and (b) list all the assignment made in the experiment (for public and private information treatments seperately). For each session, Table A4(a) reports the

Table A3(a)
Worker Assignment under Public Information about Investment

		Exogenous				Endogenous			
		No investment		Investment		No investment		Investment	
Offer	Reaction	Low	High	Low	High	Low	High	Low	High
Up-or-out									
Difficult	Yes	1	37	1	489	1	30	–	666
	No	–	–	–	1	–	–	–	–
Easy	Yes	2	–	34	–	2	–	66	4
	No	–	–	1	–	–	–	–	–
Out	Yes	59	–	43	–	42	–	39	–
	No	51	2	78	–	35	–	107	–
Up-or-stay									
Difficult	Yes	–	44	–	151	–	33	–	128
	No	–	–	–	1	–	–	–	–
Easy	Yes	125	5	44	2	153	6	58	2
	No	1	–	1	1	–	1	–	1
Out	Yes	6	–	–	–	3	1	–	–
	No	16	–	3	–	4	–	–	–
Stay-or-stay									
Difficult	Yes	–	26	–	73	–	49	–	177
	No	–	–	–	–	–	–	–	–
Easy	Yes	159	23	46	29	237	17	75	22
	No	1	10	–	23	3	14	–	23
Out	Yes	–	–	–	–	–	–	–	–
	No	10	–	–	–	1	–	–	–

TABLE A3(b)
WORKER ASSIGNMENT UNDER PRIVATE INFORMATION ABOUT INVESTMENT

		Exogenous				Endogenous			
		No investment		Investment		No investment		Investment	
Offer	Reaction	Low	High	Low	High	Low	High	Low	High
Up-or-out									
Difficult	Yes	–	6	–	126	–	4	–	55
	No	–	–	–	–	–	–	–	–
Easy	Yes	2	–	7	–	–	–	4	–
	No	–	–	–	–	–	–	–	–
Out	Yes	11	–	18	–	3	–	6	–
	No	13	–	17	–	5	–	11	–
Stay-or-stay									
Difficult	Yes	–	25	–	30	–	13	–	35
	No	–	–	–	–	–	–	–	–
Easy	Yes	98	6	18	9	52	–	10	–
	No	–	1	–	12	–	1	–	1
Out	Yes	–	–	–	–	–	–	–	–
	No	1	–	–	–				

share of high productive workers that is assigned to the difficult job. This is done separately for exogenous and endogenous treatments.

We first discuss the results of row comparisons in this table. Again, these comparisons include all periods; learning effects are reported in Table A5. Of the 45 bilateral comparisons that can be made for the exogenous sessions of the up-or-out contract (first row), we do not find significant differences. The same holds for the 91 pair-wise comparisons of the endogenous sessions (second row). For up-or-stay, out of six pair-wise comparisons only the difference between sessions 2 and 4 is significant (third row). Between the different endogenous sessions, four out of fifteen comparisons result in a significant difference; in all cases session 2 is involved (against sessions 1, 3, 4 and 9). For the exogenous sessions of the stay-or-stay contract (fifth row) four out of fifteen differences are significant. These are between sessions 5 versus 6, 8 and 14, and between sessions 7 and 14. Finally, for the endogenous sessions of the stay-or-stay contract (last row) nine out of twenty-eight differences are significant (8 *v*. 12; 11 *v*. 12, 13 *v*. 5, 6, 7, 8 and 11; and 14 *v*. 8 and 11). Taken together, these results (i) suggest that there are no systematic order effects, and (ii) give some indication of a difference in assignment outcomes between private and public information settings.

Next, we compare assignment patterns between exogenous and endogenous contracts. For the up-or-out contract there appears to be no significant difference in the share of high productivity workers that is assigned to the difficult job. For the up-or-stay contract, the difference between exogenous and endogenous contracts is significant only in session 2. For the stay-or-stay contract the difference in assignment is significant in sessions 5, 7 and 13. When the firm chooses the stay-or-stay contract, it is more likely that high-productivity workers are assigned efficiently than when the stay-or-stay contract is exogenously given. This conclusion also holds when tested using session-level means.

For each session, Table A4(b) reports the share of low-productivity workers that is assigned to the easy job. Again this is done separately for exogenous and endogenous treatments.

We start with the results of row comparisons in Table A4(b). For the exogenous sessions of the up-or-out contract (first row), significant differences are found between sessions 5 and 7, and 5 and 8. For the endogenous sessions of the up-or-out contract (second row), 23 out of 91 bilateral comparisons yield significant differences. In session 5 efficient assignment of low-productivity workers does not occur and is significantly lower compared with sessions 1, 6, 7, 8, 10, 11 and 14. Apart from that, efficient assignment is significantly higher in session 7 than in sessions 2, 3, 9, 12 and 13, and in session 8 than in sessions 1, 3, 9, 10, 11 and 12. Significant differences are also found between session 6 and sessions 3, 9 and 12, and session 10 versus sessions 2 and 12. Under the exogenous up-or-stay contract (third row), only one out of six differences is significant, viz. between sessions 1 and 3. When this contract is endogenously chosen no significant differences are found. Finally, in the fifth row, the efficient assignment rates in sessions 5 *v*. 6, 5 *v*. 7 and 7 *v*. 8 are different. These results indicate that assignment of low-productivity workers in sessions 6 and 7 is different from assignment of low-productivity workers in other sessions. This does not seem to be attributable to order effects, since sessions 5 and 8 and sessions 13 and 14 do not show such differences. Furthermore, the results do not point to differences in assignment of low-productivity workers to the easy job between the public and private information conditions.

Next, we compare assignment patterns of low-productivity workers between exogenous and endogenous contracts. For the up-or-out contract, we find significant differences in sessions 6 and 8. There also appears to be a difference between the endogenous and exogenous treatments under the up-or-stay contract in session 1. For stay-or-stay, no significant differences are found.

In Table A5 we check formally for learning effects in the assignment patterns. This table reproduces the fractions of efficient assignment of low and high-productivity workers from Table 8, but now broken down by exogenous and endogenous treatments and by the first and second halves of the sessions. In two out of twenty comparisons, we find a significant difference in assignment outcome between first and second halves. In

TABLE A4(a)
EFFICIENT ASSIGNMENT OF HIGH PRODUCTIVE WORKERS PER SESSION

		Public information												Private	
Session		1	2	3	4	5	6	7	8	9	10	11	12	13	14
Up-or-out	Exo.	0.99	0.97	1.00	1.00	1.00	1.00	0.99	1.00					1.00	1.00
	Endo.	1.00	0.96	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.00
Up-or-stay	Exo.	0.98	0.88	0.96	1.00										
	Endo.	1.00	0.60	1.00	1.00					0.94	0.88				
Stay-or-stay	Exo.					0.38	0.65	0.53	0.59					0.58	0.74
	Endo.					0.82	0.67	0.76	0.58			0.57	0.87	1.00	0.93

TABLE A4(b)
EFFICIENT ASSIGNMENT OF LOW PRODUCTIVE WORKERS PER SESSION

		Public information												Private	
Session		1	2	3	4	5	6	7	8	9	10	11	12	13	14
Up-or-out	Exo.	0.16	0.13	0.12	0.08	0.05	0.06	0.23	0.26					0.20	0.06
	Endo.	0.26	0.09	0.11	0.15	0.00	0.40	0.47	0.70	0.15	0.32	0.22	0.06	0.10	0.22
Up-or-stay	Exo.	0.77	0.88	0.92	0.89										
	Endo.	1.00	0.94	1.00	1.00					0.94	1.00				
Stay-or-stay	Exo.					1.00	0.92	0.88	1.00					0.98	1.00
	Endo.					1.00	1.00	1.00	1.00			0.99	0.97	1.00	1.00

TABLE A5
LEARNING EFFECTS IN ASSIGNMENT

		High productivity to difficult job			Low productivity to easy job		
		1st half	2nd half	*p*-value	1st half	2nd half	*p*-value
Up-or-out	Exogenous	0.99	0.99	0.9879	0.12	0.15	0.4178
	Endogenous	0.99	1.00	0.2484	0.22	0.24	0.6754
Up-or-stay	Exogenous	0.97	0.94	0.3368	0.86	0.86	0.9260
	Endogenous	0.93	0.94	0.7033	0.99	0.93	0.0181
Stay-or-stay	Exogenous	0.52	0.56	0.6257	0.97	0.93	0.1559
	Endogenous	0.71	0.79	0.1458	1.00	0.97	0.0393
Up-or-out private info	Exogenous	1.00	1.00	.	0.17	0.09	0.3310
	Endogenous	1.00	1.00	.	0.08	0.19	0.3987
Stay-or-stay private info	Exogenous	0.62	0.71	0.3980	1.00	0.98	0.3214
	Endogenous	0.96	0.96	0.9544	1.00	1.00	.

the up-or-stay and stay-or-stay treatments with endogenous contract choice, higher fractions of low-productivity workers are assigned to the easy job in the first half than in the second half. While significant, the sizes of these differences are modest.

ACKNOWLEDGMENTS

We gratefully acknowledge the valuable comments from two anonymous referees, which have improved the presentation of the material considerably.

NOTES

1. In an alternative interpretation, the trade-off between *ex ante* investment incentives and *ex post* trade efficiency that governs the choice between the up-or-out rule and the up-or-stay rule reflects the benefits and costs of prohibitively high re-contracting costs (cf. Schwartz and Watson 2004).
2. Inefficient assignment under the up-or-out contract is not an issue when investment leads to high productivity with certainty ($p_I = 1$). This is the reason why we followed Kahn and Huberman (1988) and introduced some uncertainty by having $p_I < 1$. While this uncertainty is not included in the original setup of Prendergast (1993), introducing it into his setup does not change the spirit of his model or the results.
3. Consider first contracts that specify that $w_E < r$. A low-productivity worker will not accept the easy job and so will quit. Consequently, these contracts cannot do any better than the up-or-out contract. Similarly, contracts specifying that $w_D < w_E$ do not give the worker the incentive to invest, and therefore such contracts cannot improve upon the up-or-stay contract.
4. Under up-or-out, the worker invests and thus has a $\frac{3}{4}$ probability of high productivity and a $\frac{1}{4}$ probability of low productivity. The firm's payoff from having a high-productivity worker under the up-or-out rule equals 110, while nothing is earned on a low-productivity worker. The firm's expected net payoff is thus $(\frac{3}{4} \times 110) + (\frac{1}{4} \times 0) = 82\,\frac{1}{2}$. Similarly, under up-or-stay—where the worker does not invest and therefore has a probability of only $\frac{1}{4}$ of being of high-productivity—the firm's expected payoff is $(\frac{1}{4} \times 110) + (\frac{1}{4} \times 30) = 50$. Finally, under stay-or-stay the firm can expect to earn $(\frac{1}{4} \times 125) + (\frac{3}{4} \times 50) = 68\frac{3}{4}$.
5. Let π be the fraction of reciprocal firms, i.e. the firms that will reward low-productivity workers for having made the investment. The remaining firms $(1 - \pi)$ are not willing to do so. Clearly, the worker then still prefers to invest. Expected joint surplus now equals $(\frac{3}{4} \times 220) + \frac{1}{4}\{[(1{-}\pi) \times 0] + [\pi \times 100]\} - 25 = 140 + 25\pi$. Efficiency gains thus are 25π and they increase proportionally with π.
6. To illustrate this, let π be the fraction of reciprocal firms that will punish low-productivity workers for not investing; the remaining firms $(1 - \pi)$ are not willing to do so. The worker does not know which type the firm is. His expected payoff of investing then equals $(\frac{3}{4} \times 110) + (\frac{1}{4} \times 70) - 25 = 75$, while his expected payoff of not investing equals $(\frac{1}{4} \times 110) + \frac{3}{4}\ \{[(1 - \pi) \times 70] + (\pi \times 0)]\} = 80 - 52\frac{1}{2}\pi$. Now $75 > 80 - 52\frac{1}{2}\pi$ whenever $\pi > 10/105 \approx 0.095$. The efficiency gain is then 35 points.
7. Let π now denote the fraction of reciprocal firms that will reward high-productivity workers for investing. Again, the worker does not know the type of the firm. His expected payoff of investing then equals $\pi\ [(\frac{3}{4} \times 110) + (\frac{1}{4} \times 50)] + [50(1 - \pi)] - 25 = 45\pi + 25$, while not investing yields him 50. Now $45\pi + 25 > 50$ whenever $\pi > 25/45 \approx 0.555$. The efficiency gain is then $(33\,\frac{3}{4})\pi + 12\frac{1}{2}$ points.
8. Let λ be the fraction of reciprocal workers who punish non-promotion of a high-productivity worker by quitting. A worker's type is private information. Offering promotion to a high-productivity worker yields the firm 110, non-promotion gives $125(1 - \lambda) + (\lambda \times 0)$. When $\lambda > 3/25 = 0.12$, promotion yields the firm more. And when a high-productivity worker is promoted, the worker in turn has an incentive to invest because of condition (9).
9. For $\frac{3}{4} < \alpha < 2 + \frac{5}{3}(\beta)$, the low-productivity worker is dismissed rather than offered the easy job only when he did not invest. Note that this condition cannot hold at the same time with the one for up-or-out given in the main text, because the latter necessarily requires $\alpha < 3/19$ (as $\beta < 1$ under quasi-maximin).
10. A similar remark applies for deriving predictions under the assumption that subjects make decision errors. The quantal response equilibrium (QRE) concept as developed by McKelvey and Palfrey (1998) provides a statistical generalization of Nash equilibrium which assumes noisy optimizing behaviour. Although QRE is a better solution concept than standard Nash equilibrium, without significant *a priori* restrictions on the distribution of

payoff perturbations QRE does not have much predictive content; as Haile *et al.* (2004) point out, any distribution of observed outcomes is consistent with a QRE.

11. Ten out of the overall 14 sessions were held before the official introduction of the euro, and four sessions afterwards. Because of the relatively large time span in-between the two sets of sessions, we paid a show-up fee of €5 in the second set of sessions (viz. sessions 2, 4, 6 and 8) to keep rewards in real terms the same. Average earnings in these sessions were around €25.
12. A potential advantage of role-dependent instructions would be that then it is easier for subjects to understand their own roles. In our experiment though understanding does not seem to be a problem, because no significant learning effects can be detected in the data (cf. Appendices B and C).
13. Hence during the second and third treatment of a session subjects were matched to someone they had played against in a previous treatment. Reputation effects across treatments therefore cannot be completely excluded. However, because the order was changed, and because the decision space gives little scope to signal one's identity, we are confident that reputation effects play no role.
14. The exception is an abnormally high investment rate under the stay-or-stay contract in session 12.
15. The appendices also report investment rates and assignment outcomes separately for the first and second parts of the treatments. The results indicate that subjects behave almost similarly in these two parts. Table A5 reveals that some assignment outcomes are significantly different between the two halves. The sizes of these differences are, however, modest.
16. Only in the endogenous treatment of session 8 is the assignment of low-productivity workers different. In that case 70% of the low-productivity workers is assigned to the easy job rather than dismissed (see Table A4(b)).
17. This percentage differs greatly between sessions and endogeneity of contracts. The lowest share is realized in the exogenous treatment of session 5 and equals 0.38. Even here, this share is well above the prediction of zero. In all other cases, the share exceeds 0.50 (see Table A4(a) for details).
18. With sessions as units of observation the number of observations is obviously smaller, which makes it more difficult to attain statistical significance. In particular, with six matched pairs the lowest possible p-value a two-sided Wilcoxon signrank test can attain equals 0.03125. Note that we observe differences between up-or-out and the other two rules at this level of significance. Because standard theory predicts that the investment rate is higher under up-or-out, we could employ a one-sided test here as well. Then the corresponding p-value is $p = 0.0156$.
19. This also holds if we consider only endogenous contracts.
20. With a fraction of $\pi = 0.27$ of such firms, the efficiency gain is $0.27 \times 25 = 6.75$ points (cf. n. 5).
21. The quit rates after being offered the easy job reported in Table 11 do indeed lend some support to this. A high-productivity worker rejects the easy job offer with a probability of 30% after no investment and a probability of 44% after investment.
22. These numbers reflect the potential expected payoffs, assuming that workers accept the job offered by the employer. Based on actual acceptance behaviour, the expected payoff of investment under the exogenous (endogenous) treatment equals 44.0 (57.3), whereas not investing yields 52.2 (56.5).
23. The choice between up-or-out and stay-or-stay bears some resemblance to the choice between a complete incentive contract and an incomplete bonus contract studied by Fehr *et al.* (2004) in a principal–agent experiment (see also Fehr and Schmidt 2000). In our setup, under up-or-out, the promise to promote high-productivity workers is made credible through the wages attached to the easy and difficult jobs. Under stay-or-stay, a promise to promote a high-productivity worker is not credible. Interestingly, Fehr *et al.* (2004) find that their incomplete bonus contract is more efficient than the more complete incentive contract because it gives more freedom to the parties to reciprocate. We also find that the stay-or-stay contract gives more scope for efficiency-enhancing reciprocity (but in our experiment it is not more efficient than the up-or-out rule).
24. Kahn and Huberman (1988) also assume that only the employer observes the worker's productivity. Again, this does not matter for the standard theoretical predictions, but it precludes a negative reciprocal reaction of the worker to the offer of an easy job.
25. In an alternating offer bargaining setup with advance investments, Hackett (1994) finds experimental evidence that private information about the actual investment made may interfere with a (positive reciprocity) mechanism of rewarding the investor for making a high investment.
26. In the up-or-out treatment we find no significant differences in investment behaviour with the public information situation (cf. Appendix B). The assignment decisions under up-

or-out illustrate the diminished possibility of positive reciprocity: overall only 13% of the low-productivity workers in the private information treatment are offered to stay in the easy job (compare this with the 27% in Table 9).

27. In the endogenous treatments with private information, the stay-or-stay rule is chosen in 56% of the cases. Comparing the private information treatments with the public information treatments, we find that mean individual choice rates do not differ significantly (rank sum test, $p = 0.7886$).
28. In the experiment player A represents the employer and player B the worker. The upper and lower parts of the table represent the two different contracts. The choice of disk represents the investment decision and the colour of the disk the worker's productivity. Proposals X, Y and Z correspond to Out, Easy and Difficult, respectively.

REFERENCES

BAKER, G., JENSEN, M. and MURPHY, K. J. (1988). Compensation and incentives: practice versus theory. *Journal of Finance*, **43**, 593–616.

CAMERER, C. F. (2003). *Behavorial Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.

CHARNESS, G. and RABIN, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, **117**, 817–69.

COX, J. C., FRIEDMAN, D. and GJERSTAD, S. (2006). A tractable model of reciprocity and fairness. *Games and Economic Behaviour*, forthcoming.

DUFWENBERG, M. and KIRCHSTEIGER, G. (2004). A theory of sequential reciprocity. *Games and Economic Behaviour*, **47**, 268–98.

ELLINGSEN, T. and JOHANNESSON, M. (2004). Promises, threats and fairness. *Economic Journal*, **114**, 397–420.

FALK, A. and FISCHBACHER, U. (2006). A theory of reciprocity. *Games and Economic Behaviour*, **54**, 293–315.

FEHR, E. and SCHMIDT, K. (1999). A theory of fairness, competition and cooperation. *Quarterly Journal of Economics*, **114**, 817–68.

——— and SCHMIDT, K. (2000). Fairness, incentives and contractual choices. *European Economic Review*, **44**, 1057–68.

——— and SCHMIDT, K. (2002). Theories of fairness and reciprocity: evidence and economic applications. In M. Dewatripont, L. Hansen and S. Turnovsky (eds.), *Advances in Economics and Econometrics*. Cambridge: Cambridge University Press, pp. 208–57.

——— GÄCHTER, S. and KIRCHSTEIGER, G. (1997). Reciprocity as a contract enforcement device: experimental evidence. *Econometrica*, **65**, 833–60.

——— KLEIN, A. and SCHMIDT, K. (2004). *Contracts, Fairness, and Incentives*. CESifo Working Paper no. 1215, June.

GÄCHTER, S. and FEHR, E. (1998). Reciprocity and contract enforcement. Forthcoming in C. Plott and V. Smith (eds.), *The Handbook of Experimental Economic Results*. Amsterdam, North-Holland; see http://www.elsevier.com/wps/find/bookbibliographicinfo.cws_home/601504/description#bibliographicinfo.

GIBBONS, R. (1998). Incentives in organizations. *Journal of Economic Perspectives.*, **12**, 115–32.

HACKETT, S. C. (1994). Is relational exchange possible in the absence of reputations and repeated contract? *Journal of Law, Economics and Organization*, **10**, 360–89.

HAILE, P. A., HORTAÇSU, A. and KOSENOK, G. (2004). *On the Empirical Content of Quantal Response Equilibrium*. Cowles Foundation Discussion Paper, Yale University, December.

KAHN, C. and HUBERMAN, G. (1988). Two-sided uncertainty and 'up-or-out' contracts. *Journal of Labor Economics*, **6**, 423–44.

KAMECKE, U. (1997). Rotations: matching schemes that efficiently preserve the best reply structure of a one shot game. *International Journal of Game Theory*, **26**, 409–17.

MALCOMSON, J. M. (1997). Contracts, hold-up, and labor markets. *Journal of Economic LiteratureI*, **35**, 1916–57.

MCKELVEY, R. D. and PALFREY, T. R. (1998). Quantal response equilibrium in extensive form games. *Experimental Economics*, **1**, 9–41.

MILGROM, P. and ROBERTS, J. (1992). *Economics, Organization and Management*. Englewood Cliffs, NJ: Prentice-Hall.

PRENDERGAST, C. (1993). The role of promotion in inducing specific human capital acquisition. *Quarterly Journal of Economics*, **108**, 523–34.

RABIN, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review*, **83**, 1281–1302.

SCHWARTZ, A. and WATSON, J. (2004). The law and economics of costly recontracting. *Journal of Law, Economics and Organization*, **20**, 2–31.