Promotion rules and skill acquisition: An experimental study

Oosterbeek, H.; Sloof, R.; Sonnemans, Joep

Promotion rules and skill acquisition: An experimental study*

Hessel Oosterbeek [a]
Randolph Sloof [a]
Joep Sonnemans [b]

[a] University of Amsterdam, Department of Economics, Tinbergen Institute and NWO Priority Program ‘Scholar’, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands.
[b] University of Amsterdam, Department of Economics, Tinbergen Institute and CREED, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands.

First version: September 1999
This version: May 2001

Abstract.
Gibbons (1998) identifies a tradeoff between up-or-stay and up-or-out promotion rules. Up-or-stay never wastes skills of those not promoted but may provide insufficient incentives to invest in skills. Up-or-out on the other hand can always induce investment in skill acquisition but may waste the skills of those not promoted.

This paper reports about an experiment designed to study this tradeoff. Under the up-or-out rule parties behave (almost) just as theory predicts them to do. But under up-or-stay rules results differ markedly from theoretical predictions. Workers invest rather frequently although the subgame perfect prediction is that they should not do so. Deviations from theoretical predictions can be explained by reference to different reciprocity mechanisms.

* We gratefully acknowledge valuable comments from participants at seminars in Amsterdam and Groningen and the ESA-meeting in Lake Tahoe.

JEL Codes: J41, J24, J31, C91

"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)

1. Introduction

In a recent survey article Gibbons (1998) identifies a tradeoff 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." (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 ¼, with investment it raises to ¾. Before the second period starts the firm has to assign the worker to one of two job levels; an easy job or a difficult job. Alternatively the firm may decide to fire the worker. After the investment decision but before the firm makes her assignment decision, the worker's actual productivity level is revealed to both parties. When 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. When 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. In this situation efficiency requires that low productivity workers are assigned to the easy job and high productivity workers to the difficult job, and that the worker invests since the expected joint gain of (¾-¼)·(220-100) exceeds the cost of 25.

Assume an environment in which parties cannot contract upon the worker's investment decision. Instead, the firm can attach wage levels 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 1 gives the gross payoffs for the two parties depending on the worker's level of productivity and the firm's assignment decision.

Table 1. *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

When the worker is of high productivity, the firm will promote him to the difficult job. In that job the worker produces 220 and, earning a wage of 110, both parties have a payoff equal to 110 (ignoring the possible costs of investment). When the worker is of low productivity, the firm is best off when the worker is dismissed (since 0 exceeds -10 and -110). The worker who anticipates this assignment rule, will decide to invest. Investment increases the probability of being of high productivity from ¼ to ¾, and a high productivity worker earns 110 more than a low productivity worker. Hence the expected revenue of 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 2 gives the resulting gross payoffs.

Table 2. *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

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 in an expected gain of ½·(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.

Under the up-or-out rule efficient investment is achieved but the assignment decision is inefficient, while for the up-or-stay rule it is the other way around. These outcomes are determined by the choice of the wage levels. For the given values of the exogenous variables (the worker's four potential productivity levels, the value of the outside option, the cost of investment, and the probabilities to be of high productivity with and without investment), it is, however, not possible to specify fixed wages for the two jobs such that both the investment decision and the assignment rule are efficient (see Section 2). This means that in this example there is necessarily a tradeoff between efficient assignment and efficient investment. This tradeoff parallels the choice between wage contracts which *de facto* stipulate an up-or-out rule and an up-or-stay rule.

Whether the two promotion rules work as theory predicts and whether there indeed exists a tradeoff between the two types of efficiency is an empirical issue. By 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. Normally in this context the term third party refers to 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 is not assigned, or what he would have produced elsewhere.

With this kind 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 payoffs of different alternatives are set by the researcher.

This paper reports about an experiment designed to consider the tradeoff between investment inefficiency and job assignment inefficiency under the two wage contracts. Our main findings are that up-or-out contracts perform by and large as theory predicts. That is, there is almost no investment inefficiency and job assignment is typically inefficient. In contrast, up-or-stay contracts perform better than predicted. The investment inefficiency is much smaller than theory predicts, while assignment efficiency is only somewhat below the predicted level. This deviation from the subgame perfectness predictions can be explained by reciprocity mechanisms and is consistent with the notion that reciprocity can serve as an informal contract enforcement device (cf. Fehr et al 1997).

The remainder of this paper is organized as follows. The next section describes the model in more detail. This will show more generally under which conditions parties inescapably face a tradeoff between investment and assignment efficiency. This should convince the reader that this tradeoff not only occurs in a constructed example such as presented in this introduction. Section 3 describes the design of the experiment and the predictions. Section 4 presents and discusses the results. Section 5 concludes.

2. The model

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

- y_{dh}: the value of a high productivity worker in the difficult job;
- y_{eh}: the value of a high productivity worker in the easy job;
- y_{dl}: the value of a low productivity worker in the difficult job;
- y_{el}: the value of a low productivity worker in the easy job;
- p_t: the probability that a trained worker is of high productivity;
- p_n: the probability that an untrained worker is of high productivity;
- c: the cost of training;
- 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 $w_d \geq w_e$. Next the worker decides whether to invest and spend c, or not to invest and spend 0. This decision is observed by the firm.[1] 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_t, otherwise it equals p_n. Then the firm assigns the worker to one of the two inside positions or lays the worker off. Finally the worker decides whether to accept the firm's job offer (if any) or to quit.

Table 3. *Gross payoffs for firm and worker*

	low productivity		*high productivity*	
	firm	*worker*	*firm*	*worker*
Difficult	y_{dl} - w_d	w_d	y_{dh} - w_d	w_d
Easy	y_{el} - w_e	w_e	y_{eh} - w_e	w_e
Out	0	r	0	r

Table 3 generalizes Tables 1 and 2 from the introduction. Following Prendergast (1993) we assume that:

(1) $$y_{dl}<y_{el}<y_{eh}<y_{dh}.$$

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

(2) $$y_{el}>r.$$

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) $$(p_t-p_n)\cdot(y_{dh}-y_{el})>c.$$

[1] In the original model of Kahn and Huberman (1988) only the worker observes the training level chosen, and only the employer observes the worker's productivity. In our setup these two aspects are publicly observed. As Malcomson (1997, p. 1946, footnote 22) correctly notes, this difference does not matter theoretically. Predictions based on subgame perfectness are the same whether these aspects are privately or publicly observed. All that matters is that neither the investment level nor the worker's productivity is verifiable in court, such that enforceable contracts on these aspects cannot be written.

The first-best is achieved when parties' private incentives are aligned with the efficient outcomes. The firm's assignment rule will be efficient if: $y_{el}-w_e>0$ and $y_{dh}-w_d>y_{eh}-w_e$. While the worker invests and stays with the firm if $(p_t-p_n)\cdot(w_d-w_e)>c$ and $w_e>r$.

It is now straightforward to derive that there does not exist a fixed wage contract which achieves the first-best solution, if:

$$(p_t-p_n)\cdot(y_{dh}-y_{eh}) < c. \tag{4}$$

The assumption that training is efficient only implies (3), and since $y_{eh}>y_{el}$, (4) may hold. In that case, there is either investment inefficiency or assignment inefficiency. And if the wages are set at a suboptimal level, 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 necessarily a tradeoff between investment efficiency and assignment efficiency. That is, we assume that inequality (4) holds. (We also continue to assume (1), (2) and (3).)

Under these assumptions a fixed wage contract is an up-or-out contract when $y_{el}<w_e\leq w_d<y_{dh}$. The first inequality ensures that the firm will fire a low productivity worker because his wage exceeds his highest possible value. The last inequality ensures that the high productivity worker will be kept, and since $y_{eh}<y_{dh}$ this worker will be promoted to the difficult job. Moreover, the worker will make the investment when w_d is chosen such that $(p_t-p_n)\cdot(w_d-r)>c)$, which is always possible because of (2) and (3). We thus end up with a situation in which the investment decision is efficient, but the assignment decision is not.[2]

If instead the two wage levels are set such that $r<w_e<y_{el}$ and $(w_d-w_e)<(y_{dh}-y_{eh})$ then we effectively have an up-or-stay contract. The first expression guarantees that the firm will offer a low productivity worker the easy job and that the worker will accept this offer. The second expression guarantees that a high productivity worker will be offered the difficult job which the worker accepts. In this case, the worker has, however, insufficient incentives to invest. If $(w_d-w_e)<(y_{dh}-y_{eh})$ is combined with inequality (4), we have that $(p_t-p_n)\cdot(w_d-w_e)<c$, implying that the worker's costs of investing exceed his expected benefits from it.

While the values of w_d and w_e covered by these two contracts do not cover all possible combinations of w_d and w_e, it can be shown that other values do not improve in terms of efficiency upon the two contracts considered.[3]

3. Experimental design and predictions

In the experiment we use the parameter values which also appeared in the numerical examples in the introduction. That is, for the exogenous variables we chose the following values: $y_{dl}=0$, $y_{el}=100$, $y_{eh}=175$, $y_{dh}=220$, $p_t=¾$, $p_n=¼$, $c=25$ and $r=0$. These values satisfy the four conditions specified in the previous section.

Rather than letting the firm choose freely the values of w_d and w_e, we restrict this choice to three fixed wage packages. The first fixed package represents an up-or-out contract. The second fixed package depicts an up-or-stay contract, while the third package portrays a "stay-or-stay" contract.

[2] Inefficient assignment under the up-or-out contract is not an issue when investment leads to high productivity with certainty ($p_t=1$). This is the reason why we followed Kahn and Huberman (1988) and introduce some uncertainty by having $p_t<1$. While this uncertainty is not included in Prendergast's original setup, introducing it into his setup does not change the spirit of his model nor the results.

[3] Consider first contracts which specify that $w_e<r$. A low productivity worker will not accept the easy job and will thus quit. Consequently, these contracts can not 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 promise of promotion contract.

In all three cases the wage level in the difficult job is set equal to 110. This implies that when a high productivity worker is assigned to this job, the worker and the firm split the gross surplus equally. Under the up-or-out rule, the wage level in the easy job needs to exceed y_{el} (=100), and we chose w_e=w_d=110. As we showed in the introduction the backwards induction prediction is then that high productivity workers are assigned to the difficult job, low productivity workers are dismissed and workers will invest.

For the up-or-stay contract w_e needs to be in between r (=0) and y_{el} (=100), and w_e>w_d-y_{dh}+y_{eh}=65. We chose to set w_e for this case equal to 70. As we showed in the introduction the backwards induction prediction is then that high productivity workers are assigned to the difficult job, low productivity workers to the easy job and workers will not invest.

In addition to these two contracts we also investigate the performance of a third one, namely a contract where w_e falls short of the required level of 65. For this case we have chosen a wage level of 50. 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. This is an interesting variation as it constitutes the case where the firm's promise to promote the worker is not credible. Since now (y_{dh}-w_d)<(y_{eh}-w_e), the firm will not promote a worker who is of high productivity. Hence, whatever the worker's productivity level, the firm will offer him the easy job. Therefore this contract is dubbed stay-or-stay contract. The worker who anticipates that he will never be promoted abstains from investing 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 $(p_t-p_n)\cdot(w_d-w_e)>c$.

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

Predictions of the outcomes for each of the three contracts based on subgame perfectness are summarized in Table 5. In all contracts the prediction is that workers accept a job if offered (acceptance or rejection of an "out" offer is irrelevant).

Table 5. *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	up	up	stay	up
assignment low productivity workers	out	stay	stay	stay

When the firm has to choose between contracts, she prefers up-or-out to stay-or-stay, which she in turn 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 (82½) than with the up-or-stay (50) or stay-or-stay (68¾) contracts. Consequently, if firms would all make optimal choices we would never observe how subjects play the other two contracts. Therefore our experimental design also covers treatments in which the firm is not allowed to 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 subgame perfectness and match the predictions in the relevant theoretical literature. A large body of earlier experimental studies strongly suggests, however, that subjects are also motivated by considerations as altruism, equity, fairness etcetera. In the context of contract enforcement especially the experimental results of Fehr et al. (1997) are relevant. Their results suggest that potentially *reciprocity* might matter, and may lead to different predictions. Reciprocity entails that one is willing to forgo some money in order to punish behavior that is considered as unfair and to reward behavior 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 behavior that is perceived as fair (unfair)." Due to this willingness to pay to enforce fair behavior, 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 in the experiments studied in this paper. Consider the up-or-out rule. The firm is predicted to dismiss a low productivity worker, independent of the worker's investment decision. But assume now that the firm considers an investment as a fair action which deserves a reward. This reward can be given by assigning a low productivity worker who invested to the easy job rather than dismissing him. The firm then spends 10 to give the worker a reward of 110 (see Table 1). In this example *positive reciprocity* supports the first-best outcome. As another example, consider the up-or-stay rule. The firm is predicted to assign a low productivity worker to the easy job, again independent of the worker's investment decision. But assume now that the firm considers no-investment as an unfair action which deserves some punishment. The firm can punish a low productivity worker who did not invest by dismissing him instead of offering the easy job. This costs the firm 30 and gives the worker a loss of 70. A worker who anticipates this, will make the investment. Now the anticipation of *negative reciprocity* supports the first-best outcome. Reciprocity is thus very interesting because it directly relates to the efficiency of particular types of contract in practice.

Unfortunately, while the general idea is rather straightforward, there seems to be no single, clear cut and commonly accepted definition of reciprocity. The main problem with formulating predictions based on reciprocity is that in many circumstances it is unclear when an action will be considered as fair or unfair. For instance, in the first example just considered, should it be considered fair under an up-or-out contract when the worker invests although this action is in his own best interest? Or alternatively, is it unfair if the worker does not invest or is that just stupid because not investing does not only reduce the firm's expected payoff, but also that of the worker? This implies that many different reciprocity predictions are possible based on different assumptions of what is fair and what is unfair. We do not review all possibilities here in detail. Rather, we return to this issue when we discuss the results from the experiment. There we consider whether the more substantial differences between game theoretical predictions and actual behavior can be explained by reciprocity mechanisms.

Treatments and sessions

160 subjects participated in the experiment. The subject pool was the undergraduate student population of the University of Amsterdam. Most of them were students in economics (66%). They earned on average 44 Dutch guilders (approximately US$ 22) in one and a half hour. 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. The experiments took place in a computer laboratory in which subjects are separated through cubicles. Subjects could only communicate by means of a computer network and they did not know with whom they were connected.

The design of the experiment covers five treatments and eight sessions. Table 6 maps sessions to treatments. Three treatments consider the three contracts in isolation. In these treatments the stage in which the firm determines the wage levels is skipped. The game thus starts with the investment decision of the worker. In the remaining two treatments the employer chooses between up-or-out and up-or-stay, and between up-or-out and stay-or-stay, respectively.

Table 6: *Overview of the experimental sessions*

session number	*treatments*
1	10 rounds up-or-out 10 rounds up-or-stay 10 rounds up-or-out versus up-or-stay
2	10 rounds up-or-stay 10 rounds up-or-out 10 rounds up-or-out versus up-or-stay
3	10 rounds up-or-out 10 rounds stay-or-stay 10 rounds up-or-out versus stay-or-stay
4	10 rounds stay-or-stay 10 rounds up-or-out 10 rounds up-or-out versus stay-or-stay
5 + 6	30 rounds up-or-out versus up-or-stay
7 + 8	30 rounds up-or-out versus stay-or-stay

In session 1, subjects played 10 times the up-or-out contract in isolation, followed by 10 times the up-or-stay contract in isolation and finally they played 10 times the game including the choice between these two contracts. Session 2, is identical except that the order of the isolated up-or-out and up-or-stay contracts is reversed. Sessions 3 and 4 are exact copies of sessions 1 and 2, but now the up-or-stay contract is replaced by the stay-or-stay contract. Thus, in the first four sessions the same subjects are involved in three different treatments. The motivation for this is twofold. First, this 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. Second, this setup gives the subjects sufficient experience with the various contracts available before they actually have to choose between them.

In sessions 1-4 we used a rotating scheme with two groups of ten subjects. During 10 rounds each subject from one group met each subject from the other group once. This scheme preserves the one-shot nature of the game (cf. Kamecke 1997), that is: although subjects play the same game repeatedly, the play a series of one-shot games with different opponents. Subjects thus do not play a repeated game, thereby ruling out reputational considerations. Between treatments the rotating scheme was changed in order to establish that subjects were not matched in the same order to other subjects. Subjects were informed about these features of the rotating scheme.

The firm's choice of contract in the last 10 rounds of the sessions 1 to 4 is likely to depend on the experiences during the first 20 rounds. In order to control for learning effects sessions 5 to 8 were held. These sessions consider the single treatment with an endogenous choice between contracts. Sessions 5 and 6 consider the choice between the up-or-out and up-or-stay contracts, while sessions 7 and 8 deal with the choice between the up-or-out and stay-or-stay contracts. We ran two times two identical sessions to obtain sufficient observations. With 30 rounds, 10 workers and 10 employers it is impossible to accomplish that subjects are not matched to each other more than once. The rotation scheme used in these sessions guarantees that the same matchings were kept at a minimum. Again, subjects were informed about these characteristics of the rotation scheme.

The experiment is computerized. Subjects start with on-screen instructions. All subjects have to answer some questions correctly before the experiment starts. For example, they have to calculate the earnings of subjects for some hypothetical situations. Subjects also receive a summary of the instructions on paper. A translated copy of this hand-out is attached as

Appendix A1. The instructions and the experiment are phrased as neutral as possible; words like opponent, game, and player are avoided. At the start of the first game all subjects receive a message which informs them about their role (worker or employer). After the subjects have played 30 games, they fill out a short questionnaire. At the end of the experiment the earned experimental points are exchanged for money at a rate of 1 point=1 eurocent. Subjects are paid individually and discretely.

4. Results

4.1. Summary of the findings

The main results of the experiment are summarized in Table 7, which gives for each contract the mean investment rate, the share of high productivity workers that is efficiently assigned (to the difficult job) and the share of low productivity workers that is efficiently assigned (to the easy job). In parentheses are the backwards induction predictions. Results are pooled over sessions and over exogenous and endogenous treatments. More disaggregated results are used in the sequel of this section when we explain the patterns. See also Appendices A2 and A3 which present detailed investment and assignment results, respectively.

Table 7: *Summary of results*

	up-or-out	*up-or-stay*	*stay-or-stay*
Share of efficient investment decisions	0.85 (1)	0.46 (0)	0.43 (0)
Share of efficiently assigned high productivity workers	1.00 (1)	0.96 (1)	0.69 (0)
Share of efficiently assigned low productivity workers	0.19 (0)	0.92 (1)	0.97 (1)

The results for the up-or-out contract are closest to the game-theoretical 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 are instead dismissed. For the up-or-stay contract assignment coincides almost perfectly with the predicted assignment rule, but the mean investment rate is 0.46 while no investments are predicted. Finally, for the stay-or-stay contract deviations from game-theoretical predictions are the largest. The actual investment rate is 0.43 while theory predicts a zero investment rate. Moreover, 69% of the high productivity workers are promoted to the difficult job while it is predicted that such workers would be assigned to the easy job.

In terms of efficiency, these results imply for the up-or-out contract that it 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 almost 20% of the low productivity workers are not laid off as predicted, but instead assigned to the easy job. Taken together, the realized surplus is on average around 135 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 assignment does not exactly mimic the predicted efficient assignment rule. The associated losses are, however, small in comparison with the efficiency gain resulting from the 0.46 investment rate. In sum, the realized social surplus on average amounts to 140, while 130 was predicted under this contract.

Finally, for the stay-or-stay contract there are efficiency gains on two accounts. One is due to the 0.43 investment rate, the other to the 69% high productivity workers who are assigned to the difficult job. On the other side of the balance sheet is a small loss due to the 3% low productivity workers not assigned to the easy job. Moreover there is some additional loss, which cannot be read from Table 7, 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 126 while only 118¾ is predicted.

In the next three subsections we investigate in more detail subjects' behavior 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 the up-or-out contract gives almost no scope for reciprocity considerations.

4.2. Behavior 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. Although the exact values vary somewhat between sessions and treatments, formal tests reveal that the mean individual investment rates do not vary systematically between sessions or treatments (cf. Appendix A2).

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

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

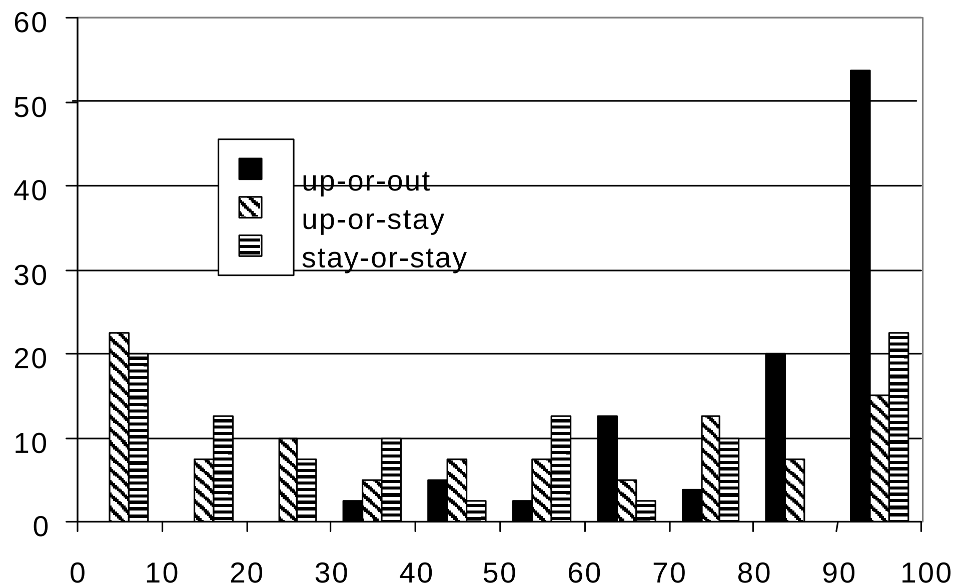

Assignment of workers to the different job options appears to be almost identical for the endogenous and exogenous treatments. See Appendix A3 for the full table. Table 8 reveals that the majority of assignments fits the subgame perfect predictions. When the worker is of high productivity, he is almost always assigned the difficult job, irrespective of whether he invested in skills acquisition or not. If the worker did not invest and turned out to be of low productivity the worker is typically dismissed, as predicted. The single contingency where the theory predicts less than (almost) perfect concerns the case where the worker is of low productivity despite his investment. Although the majority of observations (73%) is still in line with the theoretical prediction of dismissal, in 27% of the cases the worker is offered the easy job. This assignment behavior can be explained by positive reciprocity: the firm is 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 73% of the cases in fact suggests that this motive is not very important.

Table 8. *Worker assignment under up-or-out*

	No investment		*Investment*	
	low	high	low	high
Difficult	1% (1)	**98%** (47)	0% (1)	**100%** (751)
Easy	2% (2)	0% (0)	27% (65)	0% (2)
Out	**97%** (112)	2% (1)	**73%** (178)	0% (0)

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

4.3. Behavior under the up-or-stay contract

Averaged over sessions and treatments the mean individual investment rate under the up-or-stay contract equals 0.46. The actual mean investment rates fluctuate slightly between sessions and treatments, but none of the differences is significant (see Appendix A2). Although mean individual investment rates are very stable over the different treatments and sessions, individual investment rates are dispersed, see Figure 1. There is only a small peak in the frequency distribution near the game theoretical prediction of zero.

Under this contract subgame perfection predicts that high productivity workers will be offered promotion, while low productivity workers should be offered the easy job. No worker will be dismissed and, likewise, quits do not occur. Table 9 presents the realized assignment patterns. (Quits are unimportant here, see Appendix A3.)

Table 9. *Worker assignment under up-or-stay*

	No investment		*Investment*	
	low	high	low	high
Difficult	0% (0)	**91%** (53)	0% (0)	**98%** (163)
Easy	**91%** (193)	7% (4)	**98%** (65)	2% (4)
Out	9% (20)	2% (1)	2% (1)	0% (0)

Note: Endogenous and exogenous treatments are pooled. Number of cases is printed between brackets. Bold faced numbers represent assignments that are in line with theoretical predictions.

In two of the four relevant contingencies subgame perfection predicts less than perfect. 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 they are in 9% of the cases dismissed instead of being offered the easy job. In 9% of the cases workers who are of high productivity although they did not invest are offered the easy job or are even dismissed.

While only a small fraction of the workers who did not invest are actually punished, this fraction is large enough to motivate workers' investment behavior. Given the actual job offer patterns the expected payoff for the worker from investment equals about 74 points, while the expected payoff from not-investing also equals about 74. The ‘fictitious play’ strategy against the firms’ aggregate observed job offer behavior 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 the employer when he would not invest.

4.4. Behavior under the stay-or-stay contract

Averaged over sessions and treatments the mean individual investment rate under the stay-or-stay contract equals 0.43. 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 A2). We again find a large dispersion in investment rates across subjects (see Figure 1). Here the frequency distribution is bimodal; it has a peak at the game theoretical prediction of zero, and a somewhat smaller peak at the completely opposite prediction of one.

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.

Subgame perfection predicts that workers are always offered the easy job, independent of their investment and productivity. No worker will be dismissed and quits do no occur. Table 10 presents the actual assignment rule. Since we are now explicitly interested in workers quit behavior when they are offered an easy job, we added and extra row to the table.

Table 10. *Worker assignment under stay-or-stay*

	No investment		*Investment*	
	Low	high	low	high
Difficult	0% (0)	55% (54)	0% (0)	74% (196)
Easy and accept	**96% (271)**	**25% (25)**	**100% (92)**	**13% (33)**
Easy and quit	1% (4)	20% (20)	0% (0)	13% (35)
Out	2% (6)	0% (0)	0% (0)	0% (0)

Note: Endogenous and exogenous treatments are pooled. Number of cases is printed between brackets. Bold faced numbers represent assignments that are in line with theoretical predictions.

Actual assignment of low productivity workers is almost perfectly in line with predicted assignment. This is not true for high productivity workers. A majority of these workers is assigned to the difficult job rather than (as predicted) to the easy job. This assignment pattern is 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 offered. This costs the worker 50 points, and causes a damage of 125 points to the firm. As can be seen in the table this happens fairly frequently; 55 out of 113 high productivity workers who are offered the easy job, quit. Firms that anticipate this reciprocal response may want to avoid this by granting the worker promotion. 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 only applies to workers who did make the investment. Consequently our interpretation of the results in Table 10 is that the anticipated negative reciprocity causes the 55% promotion rate amongst the high productivity workers who did not invest, whereas the second form of positive reciprocity adds another 19% to the promotion rate of high productivity workers who did make the investment.

Here too, we can calculate the expected benefits for the worker of investment and no-investment under the assumption of 'fictitious play'. Averaged over endogenous and exogenous treatments the expected payoff of investment equals 58.39, and the expected payoff from not-investing equals 57.37. These expected payoffs of investment and no-investment are almost identical, 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 versus 50 points). In sum, also under the stay-or-stay contract investment is not that irrational after all, because the worker may anticipate the positive reciprocal response of the employer when he does invest.

4.5. The choice of contract

At the end of this section we turn to the analysis of the firm's choice of contract. Contract choices are only made in the endogenous treatments.

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 62%. When the choice is between up-or-out and stay-or-stay, this propensity equals 43%. In both cases subgame perfectness predicts 100%. The low mean rates are mainly caused by a large dispersion in the individual choice rates. For both the 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 large fractions of the subjects choose often for the up-or-out contract or often for the other contract. Only very few subjects choose about equally often the two contracts available to them.

Which factors determine the choice of contract? Firms will base their choices on expected actual payoffs under the different contracts. Actual payoffs may deviate from the theoretically predicted payoffs. Table 11 below provides information concerning the average net payoffs of the firm (and the worker).

Table 11. *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	82½	50	68¾	57½	80	50
Observed outcome	74.92	63.47	65.78	60.82	71.89	44.75

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. Formal statistical tests reveal that both the worker's and the firm's payoffs are not significantly different between the up-or-out and stay-or-stay contracts. Hence, firms should be indifferent between these contracts. The finding that in 57% of the cases the firm chooses the stay-or-stay contract 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. Therefore, the finding that in 38% 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 employers 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 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 and 2 subjects first played 10 rounds of the exogenous up-or-out treatment and 10 rounds of the exogenous up-or-stay treatment. Then they played 10 rounds of the endogenous game with a choice between these two contracts (cf. Table 6). 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.[4] This partly explains the anomaly that the up-or-stay contract is often chosen.

[4] The regression equation gives the following result (using OLS):
Frequency of choice up-or-stay contract = 0.488 (5.8) + 0.0015 (3.8) Earnings differential between up-or-stay and up-or-out; t-values in brackets, adjusted R squared equals 0.42.

5. Conclusion

This paper reports about an experiment designed to study the tradeoff between investment efficiency and job assignment efficiency under different promotion rules. The key finding is that in practice the rules perform differently from what theory (subgame perfectness) 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 different 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 to reciprocity. In an up-or-out contract, theory predicts that the worker starts with investing. If the worker 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 small number of the cases a low productivity worker is offered the easy job, even when the worker did invest. This in spite of the fact that offering the worker to stay rather than to dismiss him is a fairly cheap action of to 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 do provide scope for reciprocity in a way that is efficiency enhancing. The theoretical prediction is that the worker will not invest. The two contracts provide very 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 corresponds with his productivity. 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 the worker may want to invest if he anticipates that the firm will reward that by offering a more attractive job. This type of job assignment is further supported if the firm realizes that not offering a more attractive job can be retaliated by the worker in the form of quitting. Both reciprocity mechanisms appear to work, thereby increasing investment levels and also improving assignment.

From the experimental evidence we conclude that labor contracts that provide scope for reciprocity may induce efficiency levels above the levels 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 labor contracts. The difference between the contracts studied in that paper and the ones considered here is that Fehr et al. have designed labor contracts with the explicit aim to give scope for reciprocity, while we have analyzed labor contracts that emerge directly from the relevant theoretical literature.

References

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

Fehr, E., Gächter, S. and G. Kirchsteiger, 1997, Reciprocity as a contract enforcement device: experimental evidence, *Econometrica* 65, 833-860.

Gibbons, R., 1998, Incentives in organizations, *Journal of Economic Perspectives* 12, 115-132.

Kahn, C. and G. Huberman, 1988, Two-sided uncertainty and "up-or-out" contracts, *Journal of Labor Economics* 6, 423-444.

Kamecke, U., 1997, Rotations: Matching schemes that efficiently preserve the best reply structure of a one shot game, *International Journal of Game Theory* 26, 409-417.

Milgrom, P. and J. Roberts, 1992, *Economics, Organization and Management*, New Jersey: Prentice Hall.

Prendergast, C., 1993, The role of promotion in inducing specific human capital acquisition, *Quarterly Journal of Economics* 108, 523-534.

Appendix

A1: Summary of instructions

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

Besides the on line instructions a summary sheet of these instructions was handed out to the subjects. Below a direct translation of one these summary sheets is given to provide some information on exactly how the experiment was framed to the subjects. It concerns the summary sheet for the sessions 5 and 6 where only the endogenous treatments were considered.[5] In the other sessions we first explained the simpler (exogenous treatment) games. Instructions for the endogenous treatment were only given after the two exogenous treatments were 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 to the same subject in two consecutive rounds. Whenever you will 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 remain to be coupled to the same other subject. The four subsequent stages take the following form:

Stage 1: Subject A within a pair chooses between the UPPER and the LOWER table. The choice of A 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 further elaborate on this.

Stage 2: Subject B within a pair chooses between DISK 1 and DISK 2. When subject B chooses disk 1, this will cost him/her 25 points. In case subject B chooses disk 2, this does not cost anything. The choice of B for a particular disk is communicated to A.

Subsequently, the disk is turned round. When the disk has come to a stop, it will point at a particular color: blue or yellow. The color indicated by the disk is communicated to both subjects within a pair. With disk 1 the probability that the disk will point at yellow is 75%, whereas the probability of obtaining blue is 25%. With disk 2 these probabilities are exactly the other way around: 75% for obtaining blue and 25% for obtaining yellow. The indicated color co-determines the number of points the subjects receive at the end of the round.

The two disks 1 and 2 are reproduced below. The costs belonging to them that subject B has to bear are reflected within parentheses.

[5] In the experiment player A represents the employer and player B represents the worker. The upper and lower 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.

Disk 1
(25 points)

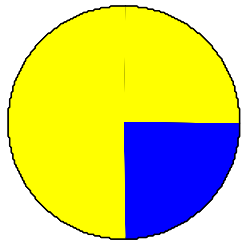

Disk 2
(0 points)

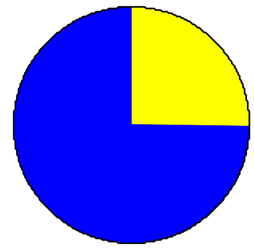

Stage 3: subject A makes a PROPOSAL to subject B. A may choose between three proposals X, Y and Z. The tables printed on the additional sheet handed out reflect the number of points both subjects earn according to the proposal. The number of points received depends on: the table (upper, lower) chosen in stage 1, the proposal (X, Y, Z) and the color indicated by the disk (blue, yellow).

Please note that with respect to subject B the tables reflect the GROSS EARNINGS. In case subject B has chosen disk 1at 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 he/she ACCEPTS or REJECTS the proposal made by subject A in stage 3. In case the proposal is accepted the earnings equal the corresponding number of points reflected in the table in question. In case 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 Euro) equals fl. 2.20. At the end of the experiment you will be paid in guilders, based on the total number of points you earned.

A2: Mean individual investment rates by session, treatment and contract

For each session Table A2 reports the mean individual investment rate under the three contracts, both for the exogenous and endogenous treatments separately and for these treatments together.

Table A2. *Mean of individual investment rates per session*

Session:		*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*
Up-or-Out	Exogenous	0.83	0.88	0.73	0.81				
	Endogenous	0.81	0.80	0.95	0.90	0.76	0.96	0.90	0.85
	Combined	0.83	0.85	0.82	0.84	0.76	0.96	0.90	0.85
Up-or-Stay	Exogenous	0.44	0.52						
	Endogenous	0.50	0.50			0.47	0.41		
	Combined	0.45	0.51			0.47	0.41		
Stay-or-Stay	Exogenous			0.35	0.52				
	Endogenous			0.27	0.52			0.31	0.68
	Combined			0.34	0.52			0.31	0.68

Using Mann-Whitney (rank-sum) tests we find no differences between the exogenous sessions of the up-or-out contract and between the endogenous sessions of the up-or-out contract. The same holds for the exogenous and endogenous sessions of the up-or-stay contract.

Under the stay-or-stay contract, the statistical tests performed reveal that individual investment rates sometimes vary significantly across sessions. For instance, in session 8 investment rates are highest. In session 7 that has an identical setup they are much lower. The observed overall dispersion, however, is *not* systematically related to the different treatments (exogenous or endogenous). The following statistical tests support these general conclusions. The exogenous treatments do not reveal a significant difference. Between the endogenous treatments significant differences do exist over the sessions. Of the 6 pairwise comparisons of individual investment rates two yield significant differences. The individual investment rates in session 8 are significant higher than the investment rates observed in sessions 3 and 7, respectively (p values equal 0.010 and 0.038). Although the significant lower investment rate in session 3 could be attributed to the different treatments involved, the difference observed for session 7 cannot. The setup of session 7 is namely identical to session 8. This seems to indicate that the complete session 8 is somewhat of an outlier. Indeed, when session 8 is ignored, no significant differences are found at all between the endogenous treatments.

Investment behavior also appears to be almost the same when comparing the exogenous and endogenous treatments within one session. Only in one out of 8 cases, there is a significant difference between exogenous and endogenous treatments. This concerns the up-or-out contract in session 3. When the up-or-out contract is endogenously chosen the propensity to invest appears to be significantly higher (Wilcoxon sign-rank test for matched pairs, p value=0.022). (The same conclusion is obtained when the first five periods of the exogenous treatment are discarded, so it seems that this difference is not caused by learning effects.) Based on these tests, we conclude that investment decisions are typically not affected by the endogeneity of the contract.

A3: Worker assignment

Table A3. *Assignment conditional on contract, investment decision and productivity*

		Exogenous				*Endogenous*			
		no investment		*investment*		*no investment*		*investment*	
Offer	*Reaction*	low	high	low	high	low	high	low	high
Up-or-Out									
Difficult	*Yes*	-	21	1	245	1	26	-	505
	No	-	-	-	1	-	-	-	-
Easy	*Yes*	-	-	18	-	2	-	47	2
	No	-	-	-	-	-	-	-	-
Out	*Yes*	26	-	19	-	30	-	32	-
	No	27	1	41	-	29	-	86	-
Up-or-Stay									
Difficult	*Yes*	-	25	-	70	-	28	-	92
	No	-	-	-	1	-	-	-	-
Easy	*Yes*	64	1	22	-	129	3	42	2
	No	-	-	1	1	-	-	-	1
Out	*Yes*	4	-	-	-	2	1	-	-
	No	10	-	1	-	4	-	-	-
Stay-or-Stay									
Difficult	*Yes*	-	10	-	33	-	44	-	163
	No	-	-	-	-	-	-	-	-
Easy	*Yes*	78	10	23	17	193	15	69	16
	No	1	9	-	14	3	11	-	21
Out	*Yes*	-	-	-	-	-	-	-	-
	No	5	-	-	-	1	-	-	-