Who should invest in specific training?*

Hessel Oosterbeek† Randolph Sloof Joep Sonnemans

March 29, 2005

Abstract

We study experimentally whether employers or workers should invest in specific training. Workers have an alternative trading opportunity that either takes the form of an outside option or of a threat point. Theory predicts that with outside options employers have (weakly) better investment incentives than workers and should therefore be the investing party. With threat points employers and workers are predicted to invest the same. Our results are by and large in line with these predictions. Due to offsetting inefficiencies in the bargaining stage, however, realized inefficiencies are remarkably similar across the different situations considered.

1 Introduction

This paper presents an experimental analysis of investments in specific training. More precisely, it addresses the question which party – either employer or worker – should invest in such training from an efficiency point of view. We will use the term specific training to refer to all types of training that create a rent that can be shared between the worker and the current firm. Traditionally, specific training has been interpreted as training that enhances skills that are only useful with the current employer, i.e. as training that is specific in a technological sense. As Stevens (1994) and Acemoglu and Pischke (1999a, 1999b) have made explicit, however, the creation of a rent, and hence the specificity of training, may also be due to market imperfections. Acemoglu and Pischke (1999a, 1999b) argue that market imperfections such

*We gratefully acknowledge the useful remarks of three anonymous referees which improved the presentation of the material considerably.

†Department of Economics, University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands. e-mail: h.oosterbeek@uva.nl.

as monopsony power, asymmetric information, unions and minimum wages may drive a wedge between the worker's productivity inside the firm and his outside options, and that this wedge may be increasing in the worker's skill level. Under such circumstances training that is general in a technological sense, is specific in an economic sense.[1] We use the term specific training to refer to this latter (broader) concept.[2]

When training is completely general in an economic sense, Becker's (1962) result that workers should bear the costs of such training still applies. In a competitive labor market, workers will always be able to reap the full marginal product of general training. Anticipating this, employers will not invest in such training. When training is not completely general (in an economic sense), however, there is no competitive market for trained workers. In that case the training firm and the worker are likely to share the additional surplus created by the investment in specific skills. The costs should therefore also be shared between the employer and the worker. Only when costs are shared proportional to the benefits, investments in specific training will be efficient (cf. Hashimoto 1981).

Explicit cost sharing assumes that investments are contractible. Typically, however, investments in training are not verifiable for a third party like a court.[3] In that case training investments are non-contractible and explicit formal cost sharing becomes cumbersome, if not impossible. Assuming that both parties have access to the same training technology, it is still possible for the employer and the worker to jointly invest in training in a non-cooperative way. Parties' *marginal* investment incentives are in that case, however, distorted. Each party bears the full marginal costs of its own investment, but cannot capture the full marginal returns of it. Theory predicts that in equilibrium the party that receives the largest share of the marginal return will do all the investment whereas the other party will invest nothing (cf. Proposition 4 in Acemoglu and Pischke 1999b). Only when the employer and the worker receive the same share of the marginal return they may both invest in equilibrium. But in that case their joint investment equals the amount each party would invest in isolation (i.e. if it were the only investing party).[4]

[1] Note that training that is specific in a technological sense, is necessarily specific in an economic sense.

[2] It should be noted that Becker (1962) already recognized this distinction between specificity in a technological and an economic sense (cf. Leuven 2005).

[3] See also Malcomson (1999, p. 2312) who notes that: "...investments may be too complex or too multidimensional for a court to verify whether they have been carried out as specified in a contract. Although it may, for example, be feasible to specify the number of hours of specific training unambiguously, specifying the quality of training during those hours is more problematic."

[4] We formally prove these predictions for the setting that we consider in

The above discussion implies that in the presence of a (potential) holdup problem, the party that is predicted to receive the largest share of the marginal returns should do all the investment in training from an efficiency point of view.[5] In this paper we investigate this prescription by means of a series of laboratory experiments. We consider a setting in which explicit cost sharing is not possible. Like in MacLeod and Malcomson (1993a), it is assumed that either the employer or the worker decides how much to invest in specific training and bears the full costs of it. After the investor has chosen the investment level, the employer and the worker bargain over the division of the surplus created by the investment. The bargaining stage is affected by the parties' alternative trading opportunities. Employers' alternative trading opportunities are normalized to zero. A worker's alternative opportunity consists of working for a different employer at a given outside wage.

The theoretical literature models alternative trading opportunities in bargaining situations in two ways. First, they has been modeled as *outside options*. The underlying assumption is that the worker can unilaterally quit the bargaining with the original employer, in order to take up an alternative employment opportunity. If the worker opts out, he cannot return to the bargaining with his original employer. If alternative opportunities take this form the so-called *outside option principle* applies. In equilibrium the worker gets the best of his outside option and the wage (bargaining share) he would obtain in the absence of outside options. The outside option thus acts as a lower bound on the equilibrium wage.

Second, the worker's alternative opportunity can be modeled as a *threat point*. The underlying assumption is then that the worker obtains his outside wage during the bargaining stage as long as agreement has not been reached. Threat point payoffs are thus collected while bargaining continues. In this case the equilibrium bargaining division equals the Nash bargaining solution, with the disagreement (status quo) payoffs equal to the values of the alternative opportunities.

Theoretical models of wage bargaining have incorporated workers' alter-

---

our experiment in Appendix B, which is available at the first author's website: http://www1.fee.uva.nl/scholar/oosterbeek/main.htm. The model analyzed in that Appendix is almost identical to the one analyzed in Acemoglu and Pischke (1999b, Section II.D). While the title of their paper refers to general training, they in fact analyze training that is specific in an economic sense.

[5] The theoretical literature proposes several contractual solutions to the holdup problem. But as Malcomson (1999, p. 2333) notes: "None of the contracts discussed here for inducing efficient specific investments by both parties thus seems unproblematic when applied to labor markets. This suggests a powerful case for, wherever possible, *all* the specific investments to be carried out by either the firm or the employee...". This provides an additional justification for our focus on settings in which only one of the parties invests.

native employment opportunities in either way. McDonald and Solow (1981) and Grout (1984), for instance, have modeled them as threat points, while e.g. MacLeod and Malcomson (1993b) and Kessler and Lülfesmann (2000) have incorporated them as outside options. Stole and Zwiebel (1996) consider both formulations. Which specification is the more relevant one in reality is not always easy to determine and depends on the situation considered. According to Malcomson (1997) the threat point specification is appropriate when the labor market operates frictionless (*no-friction case*), while the outside option formulation applies in case there are costs involved when switching trading partners (*turnover costs case*). Chiu and Yang (1999) formalize this distinction in terms of market frictions / rigidity. They consider an alternating offer bargaining framework in which the length of commitment (T) of the alternative trading opportunity is introduced as a parameter. Threat points and outside options then arise as the two limit cases where $T \longrightarrow 1$ and $T \longrightarrow \infty$ respectively. That is, if the mininum time period for which the worker has to stick to an outside employer before he can return to the bargaining with his original employer goes to one bargaining round, alternative trading opportunities act as threat points. In the opposite case where this minimum time period approaches infinity they act as outside options.[6]

The distinction can more generally be made on the basis of turnover costs, see Malcomson (1997, p. 1928). When, for instance, the worker's search and relocation costs associated with finding an equally good alternative job are high, the worker is only willing to incur these costs if the outside job is permanent. Taking up such a job effectively ends the current employment and thus acts as an outside option. Temporary nearby jobs that are easily found and readily available are better characterized as threat points though. As an alternative interpretation, within labor union-firm bargaining threat points are the payoffs the union and the firm would respectively receive in the event of a strike or a lockout. They thus represent insider factors in bargaining. The outside options reflect the parties' alternative market opportunities, i.e. outside market conditions. In a study of the UK manufacturing industry, Scaramozzino (1991) finds empirical evidence in line with this interpretation. The industry-wide (i.e. outside) wage level only affects wage levels within a firm when it is binding, whereas profits per employee (an insider factor) only do so when outside wages are non-binding.

In the experiment we consider both forms of alternative opportunities. We thus consider two different bargaining games: the outside option bar-

[6] Alternative trading possibilities are thus *less* powerful when they are *more* rigid. The intuition is that the benefits of choosing the alternative opportunity are twofold: the worker not only receives the outside wage, but also obtains the option of returning the original employer. The value of the latter option is decreasing in T.

gaining game and the threat point bargaining game. The experiment also looks at both the case where the employer invests and the one in which the worker invests. The main focus is on the comparison of investment levels of employers and workers for a given bargaining game.[7] In particular, given external market conditions (i.e. no-frictions or turnover costs), we investigate whether employers or workers should make the specific investment from an efficiency point of view.

Both bargaining games can be regarded as simplified or reduced forms of more complicated bargaining settings. The labor market imperfections which turn training that is general in a technological sense into training that is specific in an economic sense – e.g. monopsony power, asymmetric information, unions and minimum wages – are factors which may influence the bargaining stage. It is not obvious how the respective labor market imperfections exactly feed into the bargaining. A more structural model of the bargaining situation including these factors would shed light on this, and could learn us whether the threat point game or the outside option game is the more appropriate simplification. The analysis required to derive these results is beyond the scope of the present paper though. In Acemoglu and Pischke (1999b) only the threat point game is considered for the bargaining stage. They thus implicitly assume that this game is the appropriate reduced form.[8] By considering both types of bargaining games we allow for the possibility that the actual bargaining situation is better described by the outside option game.

The remainder of this paper is organized as follows. Section 2 describes the two-stage game that is used in our experiments. This section also summarizes the (subgame perfect) equilibrium predictions that are obtained for this model and spells out the hypotheses that are put to the test. Subsequently, Section 3 describes our experimental design. Experimental results are discussed in Section 4. The final section concludes.

[7] Two related papers focus on the effect of the value of alternative opportunities on the level of investment. Sloof et al. (2004) examine the situation where the non-investing party has alternative opportunities, Sonnemans et al. (2001) consider the situation in which the investor has alternative opportunities. Some of the data used in the current paper have also been analyzed in Sonnemans et al. (2001).

[8] It may appear inconsistent to use the broad definition of specific training, which also includes training that is specific due to labor market imperfections, in a bargaining model with threat points, which reflect a frictionless labor market. It should be noted, however, that the labor market imperfections that may cause training to be specific do not necessarily imply that workers cannot obtain their outside wage while bargaining. The two labels of the operation of the labor market just refer to different dimensions.

2 The model

2.1 Basic setup of the model

Consider a labor relationship between an employer and a worker. Their interaction can be described as a nested bargaining game with advance production. It has the following setup (cf. Malcomson 1997, 1999):

1. *Investment stage.* Either the employer (E) or the worker (W) makes a specific investment $I \in \{0, 1, 2, .., 100\}$. Investment costs equal $C(I) = I^2$ and are immediately borne by the investor.

2. *Bargaining stage.* The employer and the worker bargain over the division of the gross surplus $R(I)$ created by the investment. Bargaining either takes the form of the outside option (OO) or the threat point (TP) bargaining game. The employer and the worker are assumed to have equal bargaining power. The worker's outside wage equals $w \in \{w_l, w_m, w_h\} \equiv \{1800, 6800, 7800\}$.[9]

Gross surplus is linear in the investment and equals $R(I) = V + 100 \cdot I$, with $V > w$. The latter assumption, together with w being competitive, ensures that employment at the original employer is always efficient, irrespective of the actual level of investment chosen. The net social surplus created by the investment thus equals $V + 100 \cdot I - I^2$. It follows that the efficient level of investment equals $I^* = 50$.

The above description reveals that four different situations are considered, which differ by the identity of the investor and the type of bargaining game. Table 1 gives the abbreviations that we use to refer to each of these four situations. We focus in this paper on comparing the employer invests case with the worker invests case, while keeping the bargaining game fixed (either OO-game or TP-game). For this comparison the exact value of V is immaterial. Occasionally, however, we will compare the results of the two different bargaining games. With these latter comparisons in mind, we have chosen the value of V with special care.

[9] That the outside wage w is independent of the investment reflects the assumption that the investment is fully firm specific. In the formulation of Acemoglu and Pischke (1999b) the outside wage is allowed to increase with the level of the investment, i.e. $w = w(I)$. The key feature of their model, however, is that the difference between the worker's productivity in the current firm $R(I)$ and the outside wage $w(I)$ is *increasing* in the investment (i.e. $w'(I) < R'(I)$). This feature is also present in our setup. Specifying an outside wage that increases with the investment would thus only complicate the experiment for the subjects, without adding anything of substance.

Table 1: The four different situations considered

	Employer invests (E)	Worker invests (W)
Outside Option game (OO)	E-OO	W-OO
Threat Point game (TP)	E-TP	W-TP

Remark: In each cell appears the abbreviation used to refer to each of the four situations considered.

In the TP-game the worker receives the outside wage w in rounds of disagreement. The actual surplus up for renegotiation thus equals $R(I) - w$. In the OO-game the parties do not receive anything during rounds of disagreement. The actual bargaining surplus then equals $R(I)$. As a result, the *joint costs of delay* would be different in the two bargaining games if the total surplus $R(I)$ would be the same. Existing experimental studies indicate that the frequency of disagreement is inversely correlated with the joint costs of disagreement (e.g. Ashenfelter et al. 1992, Roth 1995). In our setup, delay of agreement and ex post bargaining inefficiency would therefore most likely be much smaller under the OO-game, because the marginal gain of agreement is (much) higher. To facilitate a comparison of the two different bargaining stages, especially with respect to delay of agreement and (in)efficiency, we therefore have chosen V differently in the two games: $V^{OO} = 10,000$ and $V^{TP} = 10,000 + w$. With this choice of V the joint costs of disagreement are independent of w and the same for the two bargaining games (for a given level of investment). It follows immediately from Subsections 2.2 and 2.3 that under the TP-game investment incentives are unaffected by the value of V.[10]

The different bargaining games yield different equilibrium divisions of the surplus. As these equilibrium divisions feed back into the investment stage, equilibrium investment levels will be different for the various situations considered. The next two subsections describe this in further detail.

[10] Although we feel that keeping the marginal gains from agreement constant accross bargaining games is the appropriate ceteris paribus condition (cf. Knez and Camerer 1995, p. 84), one could also argue that making the bargaining stages comparable requires the total surplus to be the same, i.e. $V^{OO} = V^{TP} = 10,000$. Erlei and Siemer (2004) exactly replicate our W-TP treatment, with the single exception that they choose $V^{TP} = 10,000$. (They do not consider the OO-game and the employer invests case.) Their findings replicate most of our results. Most notably, investment levels are by and large the same (although for $w = 1800$ they find a slightly lower investment rate). However, in their experiments the average length of the bargaining is longer, as one would expect given the lower joint costs of delaying agreement.

Table 2: Equilibrium bargaining divisions

	Employer's share	Worker's share
OO-game (DMO solution)	$R(I) - \max\{w, \frac{1}{2} \cdot R(I)\}$	$\max\{w, \frac{1}{2} \cdot R(I)\}$
TP-game (STD solution)	$\frac{1}{2} \cdot (R(I) - w)$	$\frac{1}{2} \cdot (R(I) + w)$

Remark: Equilibrium shares are independent of the identity of the investor and only depend on the type of bargaining game.

2.2 Equilibrium bargaining outcomes

In the OO-game the employer and the worker alternate in making offers about how to divide the joint surplus. If one party makes an offer, the other party can react in *three* ways: accept the offer, disagree and formulate a counteroffer in the next bargaining round, or quit the bargaining by opting out. If an offer is accepted, the employer and the worker receive payoffs according to this proposal. If there is disagreement both parties receive nothing during the round of disagreement. If a party opts out, the employer receives nothing and the worker receives his outside wage w. After opting out parties cannot return to the bargaining table.

In the TP-game parties also make alternating offers. The difference is that the responder can now react in only *two* ways: accept the offer or disagree and formulate a counteroffer. Opting out is not possible. If an offer is accepted the parties receive payoffs according to the proposal. In case of disagreement the payoffs during the round of disagreement are 0 for the employer and w for the worker.

In equilibrium agreement is reached immediately under both bargaining games. But the equilibrium division is different. Under the OO-game the so-called *outside option principle* applies. This principle entails that the gross surplus $R(I)$ is split evenly between the employer and the worker,[11] unless this yields the worker less than his outside wage. In that case he just obtains a share of the surplus equal to w, while the employer gets the remaining part $R(I) - w$. The outside wage w thus only acts as a constraint on the equilibrium division. In the words of Binmore et al. (1989), in the OO-game the equilibrium division equals the 'deal-me-out' solution (DMO).

The equilibrium prediction for the TP-game is that the surplus in excess of the outside wage w, i.e. $R(I) - w$, is split evenly. On top of that the worker obtains his outside wage w. Binmore et al. refer to this equilibrium

[11] The equal split follows from our assumption of equal bargaining power.

outcome as the 'split-the-difference' solution (STD).[12]

The equilibrium predictions under the two bargaining games are summarized in Table 2. In the TP-game neither party ever becomes residual claimant. Neither of them thus obtains the incentives to invest efficiently. In contrast, when $w > \frac{1}{2} \cdot R(I)$ under the OO-game the employer gets the whole surplus over and above the outside wage w. In that case the employer gets the full (marginal) return on investment and has the right incentives to invest. The worker never becomes residual claimant under the OO-game.

2.3 Equilibrium investment levels

Anticipating the equilibrium shares in Table 2, the investor invests the amount that maximizes his net payoffs. The equilibrium investment level depends on the identity of the investor and on the bargaining game that applies. Table 3 presents the predicted investment levels in the four situations.[13]

In the E-OO situation three relevant ranges for the outside wage w can be distinguished. First, w can be so low that it does not put a constraint on the equilibrium division. Our choice of $w_l = 1800$ represents this case. Second, w can be that high such that the outside wage constraint is strictly binding and fully determines the equilibrium division. This applies for our choice of $w_h = 7800$. The remaining third case refers to the in-between situation where the wage w is such that the equilibrium investment level makes the outside wage exactly binding, i.e. $\frac{1}{2} \cdot R(I^{eq}) = w$. $w_m = 6800$ belongs to this range and yields an equilibrium investment level of $I = 36$. Note that in this case no local interior solutions exist in the relevant sub-ranges $I \leq 36$ and $I > 36$.

The above in-between situation cannot occur when the worker makes the investment. This holds because when $\frac{1}{2} \cdot R(I) \leq w$, the worker simply obtains his outside wage w and he gains by not investing at all, saving him the investment costs. Therefore, when the worker invests there are just two relevant ranges for the outside wage w: w_l belongs to the first one and both w_m and w_h belong to the second one. Note that under the OO-game the equilibrium investment of the employer (worker) is increasing (weakly decreasing) in w. Moreover, the employer always invests weakly more than

[12] Owing to our specification of $R(I) = 10,000 + w + 100 \cdot I$ under the TP-game, we have that the employer always receives $5,000 + 50 \cdot I$ according to STD. The amount the employer receives in equilibrium is thus independent of w.

[13] The row corresponding to 'all' just reflects the average over the three different values of w. It can be shown that in an alternative model with exogenous uncertainty, in which the true value of w becomes publicly known only *after* the investment is made (and in which the three values of w have ex ante equal probabilities), the equilibrium investment levels are as follows: 36 under E-OO, $8\frac{1}{3}$ under W-OO and 25 under both TP situations.

Table 3: Equilibrium investment levels

	w	Employer invests	Worker invests
	1800	25	25
OO-game	6800	36	0
	7800	50	0
	all	37	$8\frac{1}{3}$
	1800	25	25
TP-game	6800	25	25
	7800	25	25
	all	25	25

Remark: Parameter w represents the worker's outside wage. The row corresponding to 'all' reflects the average investment level over the three values of w. The efficient investment level equals 50.

the worker does.

Under the TP-game the equilibrium level of investment is independent of the outside wage. Both the employer and the worker always underinvest. Due to the assumption of equal bargaining power they are predicted to invest the same amount in equilibrium.

It can be easily shown that, in terms of overall (in)efficiency, nothing is lost by having only one party making the investment.[14] Under two-sided investments, where both the employer and the worker can invest, the joint equilibrium investments never exceed the amounts invested by the employer in isolation (cf. Table 3).

2.4 Hypotheses

Equilibrium predictions based on subgame perfection are summarized in Tables 2 and 3 above. Our focus is on which party should make the investment from an efficiency point of view. We are therefore mainly interested in comparing, for a given bargaining game, the situation in which the employer invests with the one in which the worker does so. The corresponding hy-

[14] This is formally derived in Appendix B, which is available at the first author's website: http://www1.fee.uva.nl/scholar/oosterbeek/main.htm.

potheses are summarized below.

The predictions concerning investment levels are guided by those regarding bargaining outcomes. We therefore also want to test these latter predictions. In particular, we want to relate the different investment levels of the employer and the worker to the returns they (are predicted to) obtain on their investment. Under both bargaining games it is predicted that agreement is reached immediately. But the equilibrium division of the surplus agreed upon, and thus the return on investment, differs. In practice substantial delays in reaching agreement may be observed, as well as large deviations from the predicted divisions. This may have an (additional) adverse effect on efficiency. For a final judgement of who should make the investment, a comparison of overall efficiency levels is therefore needed. In sum, we obtain the following three hypotheses:

INV *Investment levels.* (a) Under the OO-game the employer (weakly) invests more than the worker for any value of the outside wage. (b) Under the TP-game investment levels are independent of the identity of the investor.

BAR *Bargaining outcomes.* (a) When the outside option of the worker is binding under the OO-game, the employer gets a larger return on his investment than the worker does. (b) In case the outside option of the worker is non-binding, the employer and the worker get an equal return on their investment. (c) The latter also applies under the TP-game.

EFF *Efficiency.* (a) Under the OO-game overall efficiency losses (from investments and bargaining) are smaller when the employer invests then when the worker invests. (b) Under the TP-game overall efficiency losses are independent of the identity of the investor.

3 Experimental design

The experiments cover the four situations of Table 1. We ran two sessions per investor-bargaining game combination, giving 8 sessions in total. Overall 160 subjects participated, with 20 participants per session. The subject pool was the undergraduate student population of the University of Amsterdam. Most of them were students in economics. They earned on average 27 euros in about two and a half hours.

In each session subjects played the two-stage game described in the previous section 18 times (periods) in a row. Half of the subjects performed the role of employer, the other half were assigned the role of worker. Subjects kept the same role during the whole session. To rule out reputational considerations subjects were in each period anonymously paired, using a rotation scheme for the first nine periods and a different one for the last nine periods. The experiment was computerized. The instructions and the experiment were phrased neutrally.

Like in Binmore et al. (1998) the three different values of the outside wage w were considered within each session.[15] In each session we used the same ordering of w's over the 18 periods. Each of the three values of w occured 3 times in the first block of 9 periods, and 3 times in the second block of 9 periods. Subjects were told how the ordering was generated (each of the three values of w had an equal chance of $\frac{1}{3}$ in each period), but were not told ex ante what this ordering was. At the beginning of each period they were informed about the relevant value of w.

We provided subjects with an initial endowment. Employers received 60,000 points at the beginning of the experiment, workers received 10,000 points. The conversion rate was 1 euro for 5500 points. Endowments were used to provide investors with funds to invest in the first periods. Asymmetric endowments were needed to equalize at least somewhat the unequal (equilibrium) payoffs employers and workers obtain in the game. Actual endowments were chosen such that over all four treatments employers and workers theoretically would earn about the same.

The bargaining stage was framed as a finite horizon multiple-pie alternating offer bargaining game in which one pie vanishes in each round of disagreement.[16] The employer always made the first offer, and thus could formulate the proposal in all odd rounds. Bargaining lasted for exactly 10 rounds. The gross surplus $R(I)$ and the outside wage w were spread evenly over these 10 rounds. Hence in each round a pie of $\frac{R(I)}{10}$ was to be divided between the employer and the worker. As soon as agreement was reached all remaining pies, including the one of the current round, were divided accord-

[15] A potential disadvantage of our design might be that variations in w theoretically do not affect investment levels under TP. But by confronting the same subjects with different values of w we may have created the impression that subjects are expected to change their behavior. It is, however, a priori far from clear in what direction investors should change their behavior and our experimental results confirm this.

[16] Our bargaining game thus differs from a Rubinstein (1982) like setup in which there is a *single* pie that shrinks over time at a constant rate (due to discounting when agreement is delayed). In practice, bargaining over wages in an employer-worker relationship typically concerns the division of a stream of future payoffs, instead of the division of a single once and for all payoff that is obtained immediately when agreement is reached. The multiple-pie framework nicely takes account of this aspect (cf. Manzini 1998). In Sloof (2000) it is shown that the subgame perfect equilibrium predictions of the OO-game and the TP-game employed here equal the equilibrium divisions spelled out in Table 2.

ing to the proposal agreed upon and the period ended. In the TP-game the worker received $\frac{w}{10}$ for every round that agreement was postponed, while the employer received nothing during rounds of disagreement. Postponement of agreement in round 10 ended the bargaining stage. In the OO-game both parties received nothing during a round of disagreement. Here opting out in round t resulted in a payoff of $(11-t)\cdot\frac{w}{10}$ for the worker and 0 for the employer, and ended the bargaining stage.

Finally we discuss the framing of the investment stage. At the beginning of each period subjects were informed about both the size of the initial (per bargaining round) pie $\frac{V}{10}$ and the value of the outside wage $\frac{w}{10}$. (Recall that V differs between the OO-game and the TP-game.) Then investors were asked how much they wanted to add to the initial pie size. Effectively, they chose the amount $10\cdot I$ $(=\frac{100\cdot I}{10})$ at costs I^2. The size of the actual (per bargaining round) pies was then set at the sum of the initial pie and the amount added. Subsequently, subjects bargained over the division of the ten actual pies as described above.

4 Results

The presentation of our findings follows the three hypotheses formulated in Subsection 2.4. In three subsections we deal respectively with investment levels, bargaining outcomes and efficiency. Observed investment levels did not vary significantly between the different sessions of the same investor-bargaining game situation.[17] We have therefore pooled the observations from sessions that consider the same treatment.

4.1 Investment levels

Our first result relates to hypothesis *INV* and compares employers' investments with workers' investments.

Result 1. *(a) Under the OO-game the employer invests significantly more than the worker when the outside wage is intermediate or high (w=6800 or w=7800), but not when the outside wage is low (w=1800). (b) Under the TP-game the worker invests significantly more than the employer when the outside wage is high (w=7800), but not when the outside wage is low or intermediate (w=1800 or w=6800).*

[17] Twelve Mann-Whitney ranksum tests are performed to compare mean individual investment levels conditional on the value of w. No significant differences between similar sessions are found at the 5% level.

Evidence supporting Result 1 is provided in Table 4, which reports average investment levels by treatment. Statistical tests are based on the average investment levels of individual investors (rather than on separate investment decisions).[18] In 4 out of 8 cases average investment levels of employers differ significantly from those of workers. Under the OO-game employers typically invest significantly more than workers do. Only when $w = 1800$, such that the outside wage never binds, the observed difference is not significant at the 5% level. For the TP-game an opposite conclusion holds. There employers and workers typically invest about the same: differences are insignificant for the low and intermediate outside wage. Only when $w = 7800$ workers invest significantly more than employers do. In all treatments average investment levels are below the efficient level.

To make sure that our conclusions are not biased due to ignoring learning effects, we also analyzed the data from the last nine and the final three periods separately. A full discussion can be found in Appendix A. The results reported there indicate that Result 1 is not contaminated by learning effects. An observation which is not incorporated in Result 1 is that there is quite some heterogeneity in individual investment behavior. The standard deviations reported in Table 4 already indicate this. More detailed information regarding individual mean investment levels is provided in Figure 1. This figure shows the frequency distributions of individual mean investment levels (based on six investment decisions) by treatment. The means are rounded to the nearest multiple of 10. For each treatment we have 20 individual investors. A fraction of 0.1 thus corresponds to two investors etc.. The figure confirms that there is quite some dispersion among individual investors. Take, for instance, the W-OO situation for which theory predicts that workers do not invest at all for high outside wages. In practice only three out of the 20 workers consistently choose such low investment levels (for both $w = 6800$ and $w = 7800$). At the same time there are some other investors who on average choose the efficient level of 50 in these treatments.

[Figure 1 about here]

The findings for the OO-game reported in Result 1 correspond with theoretical predictions. In line with hypothesis INV(a) holdup is less severe when the employer invests than when the worker invests. The results under the TP-game deviate from theoretical predictions only when the outside wage is

[18] Recall that per value of w we have six observations for each individual investor, and that for each of the four situations considered we have 20 investors.

Table 4: Mean investment levels by treatment

		Employer invests			Worker invests		
	w	mean	s.d.	pred.	mean	s.d.	pred.
OO-game	1800	38.7	(11.9)	[25]	28.0	(17.2)	[25]
	6800	^{a}37.9	(13.3)	[36]	^{a}21.5	(15.5)	[0]
	7800	^{b}40.0	(14.9)	[50]	^{b}21.9	(15.8)	[0]
	all	^{c}38.9	(9.0)	[37]	^{c}23.8	(14.5)	[8.3]
TP-game	1800	29.9	(14.5)	[25]	30.9	(12.5)	[25]
	6800	32.9	(12.6)	[25]	39.2	(15.5)	[25]
	7800	^{d}32.5	(13.5)	[25]	^{d}43.5	(15.1)	[25]
	all	31.8	(12.7)	[25]	37.9	(13.3)	[25]

Remark: Superscripts indicate significant differences at the 5% level (ranksum test; m=n=20 individual investors). Standard deviations appear within parentheses (s.d.), theoretically predicted investment levels are within square brackets [pred.]. Row 'all' agggregates over the three different values of w.

high ($w = 7800$). Here no significant differences were expected, but in practice workers invest significantly more than employers do. Thus in contrast with hypothesis INV(b), the extent of the underinvestment problem under the TP-game is not completely independent of the identity of the investor. Holdup is less severe when the worker invests. Finally, except for the single case where holdup is predicted not to occur (viz. $w = 7800$ in E-OO), average investment levels are always above the level predicted by subgame perfectness.

The finding that underinvestment is less severe than predicted (except when $w = 7800$ in E-OO) is in line with a number of earlier experimental studies. These papers indicate that a partial solution to holdup is provided by fairness and reciprocity considerations. Here reciprocity refers to the motivation to reward fair behavior and to punish unfair behavior, even though these rewards and punishments are costly to carry out. In a holdup context this entails that investment is typically seen as a kind act, which is therefore rewarded by the non-investor with a larger than predicted return.[19] The investment levels reported in Table 4 suggest that this informal mechanism is also at work here. In the next subsection we return to this issue when we look whether the actual investment levels can be explained from actual bargaining behavior.[20]

We end this subsection by briefly discussing two other interesting observations that can be obtained from Table 4. First we look at the comparative statics with respect to w, for each of the four different situations considered. Under the OO-game investment levels appear to be constant over the different values of the outside wage, irrespective of whether the employer or the worker invests.[21] This also holds for the E-TP situation. But for W-TP we

[19] See e.g. Berg et al. (1995), Ellingsen and Johannesson (2004a, 2004b), Gantner et al. (2001), Hackett (1993, 1994) Königstein (2000, 2001) and Oosterbeek et al. (2003).

[20] Recently some alternative behavioral theories focus on the importance of fairness and reciprocity, see Fehr and Schmidt (2002) for a comprehensive overview. These theories can only make precise predictions under strong additional assumptions (e.g. homogeneity and/or common knowledge of social preferences), but we can touch upon some indications. Fairness theories are outcome based and predict that participants will only agree on outcomes in which both parties earn about the same. This means that the outside option will never be exercised. If both parties earn the same, the efficient investment of 50 will be optimal, in all treatments. Reciprocity theories are intention based; (un) friendly actions are rewarded (punished). If investments that are higher (lower) than the equilibrium prediction are considered (un) friendly, investments will be a little more profitable for the investor. In that case investments will be higher than the equilibrium prediction in all treatments, but the relative investments levels will stay about the same.

[21] Although for the worker invests case the means and standard deviations might suggest otherwise, the Wilcoxon signed-rank test for matched pairs does not yield a significant difference between $w = 1800$ and, respectively, $w = 6800$ and $w = 7800$. The former

observe that workers' investment levels are increasing in the outside wage w.[22] The observed comparative statics for the E-OO situation suggest that contractual solutions to holdup that rely on the working of the outside option principle are unlikely to solve holdup in practice. The comparative statics in w for the W-OO situation strengthen this conclusion; they also reject the predicted relationship between investment incentives and the outside option principle.[23]

Second, we can compare average investment levels between the two bargaining games. When the employer invests investment levels are typically higher under the OO-game than under the TP-game. In case the worker makes the investment this is exactly the other way around. These results are in line with the theoretical predictions.

4.2 Bargaining outcomes

The investor's actual return on investment is determined by the offers finally accepted and the actual number of bargaining rounds. In this subsection we present results on these (and other) aspects. Our focus is on whether observed investment levels are optimal – from the selfish point of view of the investor – given actual bargaining behavior.

Before we turn to the investment incentives induced by actual bargaining behavior, we first take a descriptive look at the actual bargaining outcomes. These are also briefly compared with the outcomes found in previous bargaining experiments. Subsequently we turn to the evaluation of hypothesis BAR, i.e. we compare the actual returns on investment with the predicted ones. At the end of this subsection we then establish whether actual investments can be considered optimal given observed bargaining outcomes.

Descriptive analysis of bargaining outcomes When the bargaining stage starts investment costs are sunk. Subgame perfection then predicts that the bargaining outcome is independent of the identity of the investor. In case the OO-game applies it is predicted that the parties immediately agree on the DMO-outcome. Under the TP-game the predicted division equals the STD-solution (cf. Table 2). In practice, however, fairness considerations

p-value equals $p = 0.198$, the latter $p = 0.139$.

[22] Since these comparisons are not the main focus of this paper, we have suppressed symbols indicating results from these statistical tests in Table 4.

[23] The experiments reported in Sloof et al. (2004) provide an explanation for these findings. There it is found that a self-serving bias induces investors to overlook the difference between a binding and a non-binding outside option when making their investment decision.

are likely to play an important role. Camerer (2003, p. 196) summarizes the results of a large number of (alternating-offer) bargaining experiments and concludes that: "[G]enerally, players deviate from (perfect) equilibrium predictions in the direction of equal splits...". Besides DMO and STD we therefore consider the equal split division ES, in which the employer and the worker both receive half of the gross surplus $R(I)$, to be a relevant benchmark as well.

Table 5 provides an overview of the actual bargaining outcomes by comparing the (means of the) employers' first offers and the finally accepted offers with the three relevant benchmarks. The term offer always refers to the amount the employer gets. For the OO-game we make a distinction between the cases in which the outside wage is binding (i.e. $w \geq \frac{1}{2} \cdot R(I)$) and the ones where it is not. In the former the predicted DMO share for the employer equals $R(I) - w$, while in the latter cases DMO corresponds with the equal split division $\frac{1}{2} \cdot R(I)$.

Table 5 is divided into two panels: the left one considers the case where the employer made the investment and the right hand panel gives the case where the worker invested. The two panels have exactly the same setup. The columns labelled 'Offer' report the means of the first offers (in rows 'First') and the finally accepted offers (in rows 'Acc.'), respectively. The next three columns give the values of the three relevant benchmarks. In the OO binding situation it holds that STD<DMO<ES, while in the other two bargaining situations we have STD<DMO=ES. The fifth columns express the (means of) actual offers as share of the gross surplus whereas the two columns labeled (6) give the actual offers as share of the predicted offer. The last columns give the number of observations.

In the OO binding situation mean first offers are below the mean DMO-prediction. Employers are thus typically prepared to give workers somewhat more than their outside wage. This can be interpreted as some minimum amount needed to make the worker prefer the employer's first offer over his outside wage. The finally agreed offers are somewhat below the first offers.[24] However, these finally agreed offers remain much closer to the predicted DMO-outcome than to the STD-outcome.

In the OO non-binding situation the DMO-prediction coincides with the equal split ES. Mean first offers are much closer to DMO than to STD. This

[24] The finding that accepted offers are on average below first offers is not surprising. In our setup the employer always makes the first offer and offers are expressed as the amount the employer obtains. If the worker accepts the first offer, the finally agreed offer equals the first offer. If not, it is likely that the worker does so in anticipation of either a higher total or a higher relative payoff. In both cases the finally agreed offer must be lower than the first offer.

Table 5: Mean bargaining outcomes in relation to three relevant benchmarks

		Employer invests							Worker invests						
		Offer	STD	DMO	ES	Off./$R(I)$	Off./Pred.	n	Offer	STD	DMO	ES	Off./$R(I)$	Off./Pred.	n
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
OO binding	First	523	283	566	658	0.39	0.92	149	398	217	435	586	0.33	0.94	206
	Acc.	490	285	571	660	0.37	0.86	125	363	219	437	587	0.30	0.84	168
OO non-binding	First	813	518	720	720	0.56	1.13	211	643	517	663	663	0.49	0.98	154
	Acc.	711	525	719	719	0.49	0.99	196	575	518	663	663	0.44	0.88	147
TP-game	First	903	659	932	932	0.49	1.37	360	842	689	963	963	0.44	1.22	360
	Acc.	788	659	933	933	0.42	1.19	356	767	690	963	963	0.40	1.12	354

Remark: Columns denoted (1) report mean first offers (in rows 'First') and mean accepted offers (in rows 'Acc.'), respectively; 'offer' always refers to the amount the employer gets. Columns (2) through (4) report the average values of the three relevant theoretical benchmarks DMO (predicted outcome under OO game, cf. Table 2), STD (predicted outcome under TP-game, cf. Table 2) and ES (equal split, i.e. $R(I)/2$). Columns (5) give the mean actual offers as share of the gross surplus, columns (6) as share of the predicted offers. Columns labeled (7) report the number of observations n on which the averages are based. The rows labeled 'OO binding' belong to the OO-game with a binding outside wage (i.e. $w > R(I)/2$), the rows 'OO non-binding' to the OO-game when the wage is non-binding (i.e. $w < R(I)/2$).

also holds for finally agreed offers when the employer made the investment. In case the worker invested finally accepted offers are in between the DMO and STD outcomes. In the TP-game the difference between the ES-division and the STD-prediction is the largest. On average employers' first offers are well above the STD-division and more in line with the DMO-outcome of an equal split. This suggests that employers are not prepared to let workers fully exploit their advantageous bargaining position. Finally agreed offers are much more line with the STD-prediction, especially in the worker invests case.

Comparison with previous experimental studies The findings reported above are in line with earlier alternating offer bargaining experiments (without an investment stage). Binmore et al. (1989) consider an outside option bargaining game for three different values of the outside option. The surplus up for division was equal to £7, while the outside option of the worker (their player 2) equaled either £0, £2 or £4, respectively. The first two outside option values are non-binding whereas the outside option of £4 is binding. Binmore et al. (p. 756) conclude that in the treatments with outside option values equal to £2 and £4, "[D]eal-me-out predicted the outcomes overwhelmingly better than split-the-difference."[25] This conclusion is largely in line with our findings. Binmore et al. also find that the observed comparative statics are in line with the outside option principle: the share of the available pie finally obtained by the worker is equal for the £0 and £2 treatments where the outside option is non-binding (and typically equal to the predicted equal split), but significantly higher when the outside option is binding (i.e. equals £4).[26] The fractions reported in columns (5) of Table 5 replicate these results.

The experiments reported in Binmore et al. (1991) are based on a 2×2 design. For both the outside option and the threat point bargaining game they consider two outside wage values, viz. £1.8 and £3.2, with a pie of £5 up for division. For the outside option game DMO again predicts the actual outcomes better than STD. Moreover, comparative statics are again in line with the outside option principle; the worker obtains a significantly larger share when the outside option is binding (£3.2). In the threat point game with a high threat point value of £3.2 the outcomes are well in line with the predicted STD solution. For the low threat point value actual outcomes are

[25] Note that for an outside option value of £0 these two outcomes are identical.

[26] Binmore et al. (2002), Kahn and Murnighan (1993) and Knez and Camerer (1995) also obtain experimental evidence that the worker's final payoff does not rise one-to-one with his outside option.

in between the DMO and STD solution, just like we observe.[27] Comparing the two different bargaining games final outcomes differ significantly when the outside option is high (£3.2), and less so when it is low (£1.8).[28] Thus, subjects do in general recognize the (subtle) difference between threat points and outside options in alternating offer bargaining games.

Evaluation of hypothesis BAR After having described the actual bargaining outcomes, we next want to establish how the actual divisions agreed upon affect the investor's investment incentives. Unlike theory predicts, Table 5 indicates that the bargaining outcome depends on the identity of the investor. Especially under OO non-binding and under the TP-game the employer gets a larger share of the gross surplus when he made the investment. This suggests that the actual returns on investment differ from the predicted ones. Even then, the *relative* investment incentives of the worker and the employer may be unaffected though. Our next result relates to this.

Result 2. *(a) When the worker's outside option is binding under the OO-game, the finally accepted offers yield a larger private return on investment when the employer invests than when the worker invests. (b) In case the outside option is non-binding finally accepted offers give a private return on investment that is independent of the investor's identity. (c) The latter also applies under the TP-game.*

Result 2 follows from the regression results reported in Table 6. Here we have regressed the amount the investor receives according to the finally accepted offer on the base amount V, the level of investment I and the outside wage w. In order to determine whether the identity of the investor matters we have also included interaction terms with the dummy variable D_W. This dummy equals 1 when the worker makes the investment and 0 otherwise. Time trend t controls for potential learning effects. Observations in which the worker opted out under the OO-game (84 observations) and in which no agreement was reached in the TP-game (10 cases) are left out. For the OO-game the regressions have been estimated separately for the case where the

[27] Because we chose $V^{TP} = 10,000 + w$, our setup corresponds with the low threat point payoff in Binmore et al. (1991, Figures 3 and 4) where the threat point value is less than half of the overall surplus. Oosterbeek et al. (2003) also consider a threat point bargaining game treatment in which the pie up for division is exogenously fixed and find that in their setup STD predicts the actual bargaining outcomes better than the DMO or ES outcome does.

[28] For the £5 cake the latter result may be due to round number focal point effects. Binmore et al. (1991) therefore also consider a larger cake size ($11), and there sharper differences between the two bargaining games are obtained.

Table 6: Regressions explaining investors' finally accepted shares

	OO binding	OO non-binding	TP game
V	.850 [1]	.549 [.5]	.436 [.5]
	(.084)**	(.028)**	(.020)**
$V \cdot D_W$	−.770 [−1]	−.036 [0]	.127 [0]
	(.095)**	(.026)	(.026)**
I	.852 [1]	.719 [.5]	.653 [.5]
	(.070)**	(.074)**	(.044)**
$I \cdot D_W$	−.613 [−1]	−.016 [0]	.039 [0]
	(.079)**	(.083)	(.104)
w	−.825 [−1]	−.335 [0]	−.164 [−.5]
	(.111)**	(.046)**	(.038)**
$w \cdot D_W$	1.78 [2]	.439 [0]	.221 [1]
	(.129)**	(.055)**	(.042)
t	−1.47 [0]	−2.46 [0]	−.507 [0]
	(1.03)	(1.19)*	(.894)
n	293	343	710
adj. R^2	.910	.666	.988

Remark: In each regression the dependent variable is the amount the investor gets according to the finally accepted offer. Regressors are the base amount V, the level of investment I, the outside wage w, a time trend t and interaction terms based on dummy variable D_W, which equals 1 only when the worker makes the investment. For the OO-game there are separate regressions for the case where the outside wage is binding (i.e. $w > R(I)/2$) and the case where it is not. Numbers within square brackets refer to the predicted coefficients. Robust standard errors (in parentheses) take account of correlated disturbance terms of multiple observations per subject. Significant coefficients at the 1% level (5% level) are marked with a ** (*).

outside wage is binding ($w \geq \frac{1}{2} \cdot R(I)$) and the case where it is not. Because the regressions are based on multiple observations per subject, we calculated robust standard errors that take this into account.

The estimated coefficients for the level of investment are of particular interest. In the OO binding situation the employer is predicted to be residual claimant.[29] This is apparently not the case, because the estimated coefficient on I is significantly below one (.852) and also exceeds the coefficient on $I \cdot D_W$ in absolute value ($.852 > .613$). When the employer invests he thus gets about 85% of the marginal return on his investment, where a 100% return is predicted. The worker still gets about 25% of the marginal return on his own investment ($.852 - .613 = .239$). Here a zero return is predicted. But in line with theoretical predictions, finally accepted offers give the employer a substantially larger (marginal) private return on his investment than the worker gets if he is the investor; the coefficient on $I \cdot D_W$ is significantly negative in this case.

Under both the OO-game with a non-binding outside wage and the TP-game it is predicted that the finally agreed offer gives the investor half of the return on his investment. This is not exactly the case though. In the second and third columns of Table 5 the estimated coefficients for I exceed one half. The investor gets a return of about $65 - 70\%$ on investment. But as theory predicts, the employer and the worker get an equal return; the coefficients on $I \cdot D_W$ are not significantly different from zero in both cases. For the W-TP situation with $V = 10,000$ (instead of $10,000 + w$ as in our case), Erlei and Siemer (2004, Table 6) obtain almost exactly the same return on investment. Their estimated coefficient on I equals 0.648.

Other interesting observations follow from Table 6.[30] When theory pre-

[29] Formally the DMO and STD predictions within square brackets do not exactly apply for proposals made by the worker. Specifically, these two predictions have to be multiplied by $\frac{(9-t)}{(10-t)}$ to obtain the worker's equilibrium proposal in even round t (cf. Sloof 2000). Out of the 1346 interactions that finally ended in agreement, 400 (29.7%) were concluded upon in an even bargaining round.

[30] Instead of using the investors' finally accepted shares as independent variable we can alternatively regress the *employers'* finally accepted shares on the explanatory variables of Table 6. Because the bargaining outcome is predicted to be independent of the identity of the investor, the three interaction terms are then predicted to have no effect. These regressions (not reported here) reveal that this is almost always the case for $V \cdot D_W$ and $w \cdot D_W$ (the single exception occurs for $w \cdot D_W$ under OO non-binding), but not so for $I \cdot D_W$. The latter interaction term is significant (5% level) in OO non-binding and the TP game. We have chosen to report the regressions of Table 6 because they immediately reveal – through the (in)significance of the coefficient on $I \cdot D_W$ – whether the private investment returns differ significantly between employers and workers. Clearly the two types of regressions lead to similar conclusions.

dicts the outside wage w to be irrelevant for the final bargaining division (OO non-binding), it still has a significant influence. A higher outside wage then yields the worker a larger final share. When w is predicted to have a significant impact, its influence appears to be smaller than expected. For the OO binding situation the estimated coefficient on w equals $-.825$ when the employer invests while the prediction is -1. In case the worker invests the net coefficient on w equals .95 ($= 1.78 - .825$) and is not signficantly different from the predicted net coefficient of 1. Under the TP-game the difference is much larger. Here the outside wage is predicted to have a substantial impact, yet its actual impact is very small. By and large workers appear to be unable to exploit their bargaining advantage stemming from more favorable threat points. Finally, the coefficient on t reveals that some (weak) learning occurs only in the OO non-binding situation. Here the investors' finally accepted shares decrease during the course of the experiment. Even for this bargaining situation, however, the other coefficients are almost identical whether a time trend is included or not.

The higher than predicted return on investment under OO non-binding and the TP-game can be explained by positive reciprocity. Reciprocal behavior is a common finding in a large number of different experimental games (cf. Fehr and Gächter 2000). In our experiment, investing more than predicted can be interpreted as something that is kind and therefore warrants a larger than predicted return. If the investor anticipates such a reciprocal reaction, it is optimal for him/her to invest more than predicted. As noted before in Subsection 4.1, a number of other experiments also find that investors are partly compensated for the sunk investment costs borne.

Result 2 only considers the return on investment as it is present in the finally agreed offers. But for the actual investment incentives it is also important to know how long it takes to reach agreement. Our next result compares delay of agreement accross treatments.

Result 3. *(a) Under the OO-game agreement is reached sooner when the worker invests than when the employer invests. (b) Under the TP-game with a low outside wage (w=1800) agreement is also reached sooner when the worker invests. In case the outside wage is high (w=6800 or w=7800) there are no significant differences in delay.*

Result 3 follows from Table 7 which reports for given levels of w the percentage of cases in which agreement is reached immediately and the mean number of bargaining rounds needed to reach agreement. The columns labeled 'n' give the numbers of observations on which the latter averages are based. These numbers differ from the maximum of 120 (360 for all) due to opting out (in the OO-game) and no agreement (in the TP-game).

Theory predicts that agreement is always reached immediately (in the first round) and that opting out does not occur. Actual outcomes deviate from these predictions: the percentages of immediate agreement are well below 100% and opting out does occur. For each level of the outside wage w, ranksum tests on the number of bargaining rounds reveal that under the OO-game agreement is reached significantly sooner when the worker invests than when the employer invests. The same holds for the TP-game when $w = 1800$, but not when $w = 6800$ or $w = 7800$. In the latter cases the mean number of bargaining rounds do not vary with the identity of the investor.[31] An explanation why on average agreement is reached sooner when the worker has made the investment is that the employer then demands for less in her/his first offer, both as share of the overall surplus and as share of the predicted outcome (cf. Table 5).

Results 2 and 3 can be used to evaluate hypothesis *BAR*. The first part of this hypothesis *BAR*(a) receives mixed support. Result 2(a) provides evidence in favor of the employer getting a larger return on investment in the OO binding situation, Result 3(a) provides evidence against it. Here the employer's larger marginal share of the final agreement and the longer delay when he invests work in opposite directions. Hypothesis *BAR*(b) is rejected, because the speed with which agreement is reached gives the worker a larger return on investment when the outside option is non-binding.[32] The third part of hypothesis *BAR* receives qualitative support. Although there is some evidence that when the outside wage is low agreement is reached sooner when the worker invests (Result 3(b)), overall the return on the investment indeed seems fairly similar for the worker and the employer.

'Optimum' investment levels The mean investment levels reported in Table 4 differ substantially from the predicted levels in most treatments. But given the (at best) mixed support for hypothesis *BAR*, the observed investment levels may be still be optimal (from a selfish point of view) given actual bargaining behavior. The last result of this subsection relates to this.

Result 4. *(a) When the outside wage is low (w=1800) under the OO-*

[31] Appendix A reports the average number of bargaining rounds for the last nine and final three periods separately (cf. Table 11). Average delay is typically shorter in later periods. Apparently subjects learn to avoid costly delay. Result 3 is, however, not affected by this. It is supported when we consider only the periods 10 to 18 or when we just look at the final three periods.

[32] The opting out rates under E-OO (11%) and W-OO ($12\frac{1}{2}$%) are fairly similar and thus do not affect the relative return on the investment in the two cases.

Table 7: % immediate agreement and mean # of rounds before agreement

		Employer invests			Worker invests		
	w	% imm.	# rnds	n	% imm.	# rnds	n
OO-game	1800	39.2%	$^{a}2.39$	118	57.5%	$^{a}1.85$	115
	6800	51.7%	$^{b}1.60$	101	72.5%	$^{b}1.27$	105
	7800	64.2%	$^{c}1.36$	102	70.0%	$^{c}1.15$	95
	all	51.7%	$^{d}1.82$	321	66.7%	$^{d}1.44$	315
TP-game	1800	35.0%	$^{e}2.85$	118	46.7%	$^{e}2.23$	118
	6800	40.0%	2.44	118	35.0%	2.34	118
	7800	32.5%	2.49	120	37.5%	2.19	118
	all	35.8%	2.59	356	39.7%	2.25	354

Remark: The columns '% imm.' give the percentage of cases in which agreement is reached immediately. Columns labeled '# rnds' report the mean number of bargaining rounds needed to reach agreement and the columns labeled 'n' the number of observations on which these avarages are based. Row 'all' aggregates over the three different values of w. Superscripts indicate significant differences at the 5% level (ranksum test).

game the 'optimum' investment level of the employer is below the worker's 'optimum' investment level. (b) In case the outside wage is high (6800 or 7800) this is the other way around. (c) Under the TP-game 'optimum' investment levels are similar for both investors when the outside wage is low (w=1800), while in case w is high (6800 or 7800) the 'optimum' investment level is higher for the worker.

We estimated regression equations with the investors' net payoffs as dependent variable, and the level of investment and investment squared as independent variables. To control for potential learning effects we included a variable that measures the time that the investor was confronted with the particular outside wage (hence this variable ranges from 1 to 6).[33] The 'optimum' levels of investment can be directly obtained from the estimated coefficients. Table 8 reports these 'optimum' investment levels, along with their standard errors.[34] In two out of twelve treatments the estimated coefficient for I and I^2 were both negative and insignificant, yielding an optimal investment level of 0.[35] This actually corresponds with the theoretical predictions for these two cases.

When $w = 1800$ under the OO-game the selfish worker should invest more than the employer according to our estimates of 'optimum' investment levels. (But note that the standard error on the worker's estimated optimum is particularly large.) For $w = 6800$ and $w = 7800$ this is exactly the other way around. When the TP-game applies the estimated optima are always larger for the worker, although the differences are very minor when the outside wage is low ($w = 1800$). One potential caveat applies. In some of the cases the estimated standard errors are rather high. This suggests that the variance in observed bargaining behavior is so large that it is of little help in determining the appropriate investment level.

Comparing Results 1 and 4 we observe that in four out of six relevant cases (i.e. bargaining game-outside wage combinations) actual bargaining behavior can explain differences between employers' and workers' actual investment levels. The first exception occurs when $w = 1800$ under the OO-game. In this case the worker's outside option is always non-binding. Theory then predicts that employers and workers invest the same, while actual bargaining behavior

[33] Except for $w = 7800$ under W-OO and $w = 1800$ under W-TP these time trends were never significant (at the 5% level).

[34] We again calculated robust standard errors that take account of correlated disturbance terms of multiple observations per subject.

[35] When we use the negative coefficients to calculate the unconstrained optima, the following 'optimum' investment levels are obtained (both for W-TP) : -11.48 when $w = 6800$ and -2.40 when $w = 7800$. The reported robust standard errors for these two treatments (Table 8) in fact refer to these unconstrained optima.

Table 8: "Optimum" investment levels

	w	Employer invests		Worker invests	
OO-game	1800	28.47	(3.56)	34.43	(13.04)
	6800	32.17	(5.23)	0	(24.50)*
	7800	38.07	(4.28)	0	(5.46)*
TP-game	1800	23.90	(4.76)	24.28	(7.60)
	6800	18.04	(7.68)	30.10	(3.52)
	7800	30.25	(3.42)	35.71	(2.51)

Remark: In each cell appears the estimated optimal (from a selfish point of view) investment level given actual bargaining behavior. These estimates follow from regressing the investors' net payoffs on the investment level, investment squared and a time trend. Robust standard errors (in parentheses) correct for multiple observations per subject. n=120 in all cases. Superscript * indicates that the level is set at the lower border of 0, because the estimated unconstrained optimum is negative. Reported standard errors refer to these unconstrained optima.

seems to indicate that workers should invest more (Result 4(a)). This is not what we observe though. Employers invest more than workers do, albeit not significantly so (Result 1(a)). Given the calculated 'optimum' investment levels we conclude that employers substantially overinvest from a selfish point of view, while workers seem to underinvest in this case.

The second exception concerns the TP-game when $w = 6800$. From Result 4(c) we can conclude that workers should invest more in this case. We do indeed observe that workers on average invest more (cf. Table 4), but the difference lacks significance. Given that the differences in average observed investment levels are fairly large (and that the p-values are always below .14), we conclude for this case that actual bargaining behavior can provide a reasonable explanation for employers' and workers' investments.

In case the outside wage is high under the OO-game ($w = 6800$ and $w = 7800$), relative investment levels are in line with actual private returns. Yet our calculations suggest that workers now substantially overinvest from a selfish point of view, whereas employers do only slightly so. These deviations are not substantive enough to alter the relative investment levels from the predicted direction; employers invest significantly more, in line with their higher 'optimum' investment level.

Also under the TP-game the observed differences between investment levels can partly be explained on the basis of the actual private returns on the investment. There workers invest significantly more when the outside wage equals $w = 7800$ (Result 1(b)), in line with the ordering of the 'optimum' investment levels presented in Table 8. The observed differences are, however, somewhat larger than one would expect on the basis of actual private returns. Both the employer and the worker seem to overinvest from a selfish point of view. A possible explanation for the observed overinvestment is the presence of self-serving biases. Both employers and workers may have a self-serving assessment of what is a fair or reasonable return on investment. Some experimental evidence that such self-serving biases can indeed explain over-investment is provided by Sloof et al. (2004). Overall we conclude that actual private investment returns can explain observed investment patterns reasonably well.

4.3 Efficiency

Subgame perfection predicts that there will be no efficiency losses in the bargaining stage. Inefficiencies will be solely due to underinvestment. Our final result concerns actual inefficiency.

Result 5. *(a) Under the OO-game investment (bargaining) inefficien-*

cies are smaller (larger) when the employer invests than when the worker invests. Total inefficiencies are by and large independent of the identity of the investor. (b) Under the TP-game with a high outside wage (w=7800) investment inefficiencies are larger when the employer invests. As a result overall inefficiencies are somewhat smaller when the worker invests.

Table 9 presents the evidence for this result. Investment inefficiencies are calculated as the difference between the maximum net surplus of 2500 achieved at $I = 50$ and the actual net surplus of the investment made, which is equal to $100 \cdot I - I^2$. Under the OO-game investment inefficiencies are always significantly smaller when the employer invests. Under the TP-game with $w = 7800$ investment inefficiencies are larger when the employer invests. For lower values of w the investment inefficiencies do not vary significantly with the identity of the investor.

Bargaining inefficiencies reflect our earlier findings regarding delay of agreement and opting out, which are the two sources of this type of inefficiency. Under the OO-game bargaining inefficiencies are significantly smaller when the worker invests.[36] Under the TP-game average bargaining inefficiencies are also smaller when the worker invests, but the observed differences are not significant at the 5% level.

Investment and bargaining inefficiency together yield total inefficiency. Interestingly, not many significant differences are observed between different investor's identities. By and large we can conclude that under the OO-game overall inefficiencies tend to be independent of the investor, while in case the TP-game applies overall inefficiencies are somewhat smaller when the worker invests.

Result 5 provides qualitative support for hypothesis *EFF*. Under the OO-game we do observe that efficiency losses due to suboptimal investment are smaller when the employer invests than when the worker invests, in line with *EFF(a)*. But overall observed inefficiencies are not in line with theoretical predictions. For the TP-game we indeed observe few significant differences between the two investor types. Only when $w = 7800$ observed inefficiencies contrast with hypothesis *EFF(b)*. Inefficiencies are then smaller when the worker makes the investment.

A final interesting observation that follows from Table 9 is that inefficiencies owing to delayed agreement and opting out are substantial. In all

[36] The part of the bargaining inefficiencies that can be attributed to the worker opting out (included in the bargaining inefficiencies reported in Table 9) does not vary significantly with the identity of the investor. For the E-OO situation the average 'opting out' inefficiencies equal 129, 923, 787 and 613 for w_l, w_m, w_h and all, respectively. For W-OO they are equal to 267, 559, 882 and 570.

Table 9: Average efficiency losses

		Employer invests			Worker invests		
	w	invest.	bargain.	total	invest.	bargain.	total
	1800	^{a}429	^{e}2148	2577	^{a}828	^{e}1545	2373
OO	6800	^{b}401	^{f}1940	2341	^{b}1130	^{f}894	2024
	7800	^{c}387	1441	^{h}1828	^{c}1158	1091	^{h}2249
	all	^{d}406	^{g}1843	2249	^{d}1039	^{g}1177	2206
	1800	698	2599	3296	759	1826	2585
TP	6800	581	2128	2709	470	2065	2536
	7800	^{i}642	1983	^{j}2625	^{i}402	1904	^{j}2306
	all	640	2237	^{k}2877	544	1932	^{k}2476

Remark: Columns labeled 'invest.' report average investment inefficiencies, calculated as $2500 - 100 \cdot I - I^2$. Columns 'bargain.' reports average bargaining inefficiencies (due to delay and opting out) and 'total' the average overall inefficiencies. Superscripts indicate significant differences at the 5% level (ranksum test). n=120 for each value of w.

treatments bargaining inefficiencies typically outweigh investment inefficiencies. Bargaining inefficiencies are smallest in the W-OO situation where the investment inefficiencies are the largest. This points to a potential trade-off between investment and bargaining inefficiency. Theoretically W-OO should perform worst because investment inefficiencies are predicted to be the largest. In practice bargaining inefficiencies are particularly large in the other three situations considered, causing them to perform even worse than the W-OO case.

5 Conclusion

In this paper we address the question whether the employer or the worker should make an investment in training that is specific in an economic sense. We consider a setting in which only the worker has alternative trading opportunities in the labor market. Depending on whether the market for alternative jobs operates frictionless (no friction case) or not (turnover costs case), bargaining between the employer and the worker is either modelled as a threat point game or as an outside option bargaining game. Theoretically investment incentives are the same for the employer and the worker when the former case applies. If, however, the outside option bargaining game applies the employer is predicted to have the better investment incentives. He therefore should make the investment from an efficiency point of view.

By and large our results are in line with these predictions. For the turnover costs case we indeed observe that only when the outside wage of the worker is high, employers invests more than workers do. For the no-friction case we obtain the opposite. There workers invest more than employers when the outside wage is high. Actual bargaining outcomes can provide a reasonable explanation for the observed differences between employers' and workers' investment levels. Only when the outside wage is low under the outside option game the observed private returns on the investment made do not justify the observed differences between employers and workers. Overall the observed inefficiencies are remarkably similar across the different situations considered. If anything, they suggest that the employer should make the investment in the turnover costs case, while the worker should invest when the no-friction case applies.

In our experiment we consider one-sided investments only. Although standard theory predicts for our setup that allowing for two-sided investments would not improve investment incentives, it may still be the case that in practice efficiency is higher when the two parties can effectively share the costs of investment. It might therefore be interesting to experimentally investigate treatments with two-sided investments as well. This would also allow for a direct test of the theoretical prediction that (generically) only one party will actually invest (cf. Proposition 4 in Acemoglu and Pischke 1999b). In contrast to this prediction, one may reasonably expect that fairness considerations make the outcome in which the parties each invest half of the efficient level (and share the gross surplus equally) focal, especially when the outside wage is low.[37] Whether this is indeed the case is left for future research.

Appendix A: learning effects

In each session subjects played 18 times the two-stage game of Section 2. During the course of the experiment they may have changed their behavior, for instance because over time they learned how to play the game. To make sure that our conclusions are not biased due to ignoring such learning effects, we consider in this Appendix also the data from the last nine and the final three periods separately. The focus is on investment levels (cf. Result 1 and Table 4) and on delay of agreement (cf. Result 3 and Table 7). Recall that in the regressions reported in the main text that lead to Results 2 and 4 we already control for potential learning effects.

The design of the experiment was such that the first and last nine periods included the same frequency of low, intermediate and high levels of the worker's outside wage. Moreover, each value of w was represented exactly once in the final three periods. Tables 10 (investment levels) and 11 (bargaining length) below report the same statistics as in Table 4 and 7 of the main text, but now also for the last nine and final three periods separately. The top panels of Tables 10 and 11 correspond exactly with the tables presented in the main text. The middle panels only consider the data from the second half of the experiment, while the bottom panels only use the data from the final three periods.

[37]Fehr et al. (2004) consider a setting in which paired subjects make *sequential* investment decisions. Before investments are made, subjects first bargain about the allocation of ownership (which in turn specifies how the returns to investments are shared). Under joint ownership standard theory predicts no investments at all whereas under single ownership only the owner (which is the 2^{nd} investor in the experiment) is predicted to invest. In contrast to these predictions, Fehr et al. find that joint ownership is chosen most often, and also that under this ownership structure typically both subjects invest. The latter can be explained by reciprocity; the second investor on average strongly reciprocates higher investments made by the first investor (although standard theory predicts that these investments are independent). This in turn makes it attractive for the first investor to invest under joint ownership. Note that, in contrast to our experiment, in these experiments there is no ex post bargaining.

Table 10: Mean investment levels by treatment

		Employer invests			Worker invests		
	w	mean	s.d.	pred.	mean	s.d.	pred.
OO-game {1-18}	1800	38.7	(11.9)	[25]	28.0	(17.2)	[25]
	6800	^{a}37.9	(13.3)	[36]	^{a}21.5	(15.5)	[0]
	7800	^{b}40.0	(14.9)	[50]	^{b}21.9	(15.8)	[0]
	all	^{c}38.9	(9.0)	[37]	^{c}23.8	(14.5)	[8.3]
TP-game {1-18}	1800	29.9	(14.5)	[25]	30.9	(12.5)	[25]
	6800	32.9	(12.6)	[25]	39.2	(15.5)	[25]
	7800	^{d}32.5	(13.5)	[25]	^{d}43.5	(15.1)	[25]
	all	31.8	(12.7)	[25]	37.9	(13.3)	[25]
OO-game {10-18}	1800	33.8	(13.9)	[25]	25.7	(17.1)	[25]
	6800	^{e}36.5	(14.0)	[36]	^{e}18.9	(17.0)	[0]
	7800	^{f}40.3	(15.9)	[50]	^{f}19.7	(17.3)	[0]
	all	^{g}36.9	(9.1)	[37]	^{g}21.4	(15.9)	[8.3]
TP-game {10-18}	1800	28.7	(14.6)	[25]	29.4	(14.4)	[25]
	6800	32.3	(17.2)	[25]	40.7	(15.8)	[25]
	7800	31.0	(17.2)	[25]	41.0	(16.2)	[25]
	all	30.7	(15.4)	[25]	37.0	(14.4)	[25]
OO-game {16-18}	1800	34.9	(14.1)	[25]	24.6	(18.3)	[25]
	6800	^{h}37.4	(13.3)	[36]	^{h}19.2	(17.5)	[0]
	7800	^{i}42.4	(15.1)	[50]	^{i}20.1	(17.7)	[0]
	all	^{j}38.2	(9.1)	[37]	^{j}21.3	(16.7)	[8.3]
TP-game {16-18}	1800	29.7	(17.3)	[25]	30.3	(24.7)	[25]
	6800	32.9	(23.3)	[25]	40.0	(17.5)	[25]
	7800	29.7	(19.7)	[25]	41.5	(20.1)	[25]
	all	30.8	(17.9)	[25]	37.3	(18.2)	[25]

Remark: Superscripts indicate significant differences at the 5% level (ranksum test; m=n=20 individual investors). Standard deviations appear within parentheses (s.d.), theoretically predicted investment levels are within square brackets [pred.]. Rows 'all' agggregate over the three different values of w. Periods considered are within curly brackets.

Table 10 reports average investment levels by treatment. Statistical tests are again based on the average investment levels of individual investors. The results in the middle and bottom panel almost exactly reproduce the results of the top panel. The single difference is that no significant differences are found anymore under the TP-game when $w = 7800$. This holds despite the fact that the mean levels over all investors are fairly far apart. Comparing for this particular case (i.e. $w = 7800$ under the TP-game) average investment levels across the different panels of Table 10 by means of a Wilcoxon signrank test, we find no significant differences when the worker invests (the lowest p-value equals $p = .304$). For the case in which the employer invests, however, the top panel differs significantly from both the middle ($p = .047$) and the bottom panel ($p = .032$). Over time employers thus tend to invest less in this case, while workers do not change their investment behavior. Based on these learning effects, we still conclude that under the TP-game workers invest more than employers when the outside wage is high.[38]

As an additional test of learning effects we regressed, for each of the twelve treatments, the investment levels on a variable which measures the time that the investor was confronted with this particular value of the outside wage (besides a constant term). Only in two treatments this time trend had a statistically significant (negative) coefficient: the case where $w = 1800$ under E-OO and the case where $w = 6800$ under W-OO. As can be seen from Table 10, however, the differences in overall mean investment levels for both these treatments in the three panels are small (and they display a non-monotonic pattern). Taking all the above checks together, we conclude that Result 1 on investment behavior is not contaminated by learning effects.

Table 11 presents for each treatment the percentage of cases in which agreement is reached immediately and the average number of bargaining rounds before agreement is reached. Columns labeled 'n' give the number of observations on which the latter averages are based; observations in which one of the parties opted out (OO-game) or no agreement was reached are left out. Comparing the middle with the top panel it is observed that average delay is typically shorter in the second half than in the first half of the exper-

[38] An explanation for why we do not get significant results (at the 5% level) in the middle and bottom panel is that the averages used there are based on fewer observations (3 and 1, respectively) and thus are more noisy. The reported standard deviations reflect this; in the middle and the bottom panel standard deviations are higher than in the top panel. Assume that each investor has a personal inclination for a certain behavior. Actual behavior is determined by this inclination and some 'noise' or errors. The average behavior of the investor is an estimation of that individual inclination. More observations per individual means a better measurement of these inclinations, and so a better quality of the data used in the test. Although the test itself is independent of how many observations are included in the average, the power of the test increases if more data per individual is used.

Table 11: % immediate agreement and mean # of rounds before agreement

		Employer invests			Worker invests		
	w	% imm.	# rnds	n	% imm.	# rnds	n
OO-game {1-18}	1800	39.2%	[a]2.39	118	57.5%	[a]1.85	115
	6800	51.7%	[b]1.60	101	72.5%	[b]1.27	105
	7800	64.2%	[c]1.36	102	70.0%	[c]1.15	95
	all	51.7%	[d]1.82	321	66.7%	[d]1.44	315
TP-game {1-18}	1800	35.0%	[e]2.85	118	46.7%	[e]2.23	118
	6800	40.0%	2.44	118	35.0%	2.34	118
	7800	32.5%	2.49	120	37.5%	2.19	118
	all	35.8%	2.59	356	39.7%	2.25	354
OO-game {10-18}	1800	45.0%	2.30	60	58.3%	1.81	57
	6800	56.7%	[f]1.65	54	80.0%	[f]1.13	53
	7800	71.7%	[g]1.30	54	83.3%	[g]1.04	52
	all	57.8%	[h]1.77	168	73.9%	[h]1.34	62
TP-game {10-18}	1800	40.0%	[i]2.61	59	61.7%	[i]1.92	59
	6800	38.3%	2.33	60	31.7%	2.42	59
	7800	33.3%	2.43	60	43.3%	2.07	59
	all	37.2%	2.46	179	45.6%	2.14	177
OO-game {16-18}	1800	55.0%	2.25	20	65.0%	1.37	19
	6800	50.0%	[j]1.94	18	85.0%	[j]1.11	18
	7800	75.0%	1.18	17	85.0%	1.06	18
	all	60.0%	[k]1.82	55	78.3%	[k]1.18	55
TP-game {16-18}	1800	40.0%	2.95	20	60.0%	1.75	20
	6800	30.0%	2.45	20	20.0%	2.3	20
	7800	20.0%	2.45	20	55.0%	1.75	20
	all	30.0%	2.62	60	45.0%	1.93	60

Remark: The columns '% imm.' give the percentage of cases in which agreement is reached immediately. Columns labeled '# rnds' report the mean number of bargaining rounds needed to reach agreement and the columns labeled 'n' the number of observations on which these avarages are based. Rows 'all' aggregate over the three different values of w. Superscripts indicate significant differences at the 5% level (ranksum test). Periods considered are within curly brackets.

iment. This follows because in almost all treatments the average number of bargaining rounds before agreement is reached decreases when we take only the last nine periods into account (the two exceptions occur when $w = 6800$ under E-OO and W-TP). Apparently subjects learn to avoid costly delay when they play the game. Result 3 on delay of agreement is, however, not seriously affected by this. It is fully supported when we consider only the periods 10 to 18.[39] For the final three periods the same type of differences are found, although not all of them are significant. But also there we observe that overall agreement is reached sooner under the OO-game when the worker invests. In case of the TP-game there are no significant differences, although when $w = 1800$ the overall observed mean bargaining length before agreement is substantially larger when the employer invests than when the worker invests, in line with Result 3(b).

References

Acemoglu, D. and J.-F. Pischke (1999a). Beyond Becker: Training in imperfect labor markets. *Economic Journal 109*, F112–F142.

Acemoglu, D. and J.-F. Pischke (1999b). The structure of wages and investment in general training. *Journal of Political Economy 107*, 539–572.

Ashenfelter, O., O. Currie, H. Farber, and M. Spiegel (1992). An experimental comparison of dispute rates in alternative arbitration systems. *Econometrica 60*, 1407–1433.

Becker, G. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy 70*, S9–S49.

Berg, J., J. Dickhaut, and K. McCabe (1995). Trust, reciprocity, and social history. *Games and Economic Behavior 10*, 122–142.

Binmore, K., J. McCarthy, G. Ponti, L. Samuelson, and A. Shaked (2002). A backward induction experiment. *Journal of Economic Theory 104*, 48–88.

Binmore, K., P. Morgan, A. Shaked, and J. Sutton (1991). Do people exploit their bargaining power? An experimental study. *Games and Economic Behavior 3*, 295–322.

Binmore, K., C. Proulx, L. Samuelson, and J. Swierbinski (1998). Hard bargains and lost opportunities. *Economic Journal 108*, 1279–1298.

[39] The p-value of the ranksum test comparing E-OO with W-OO for $w = 1800$ equals $p = .060$, close to significance at the 5% level.

Binmore, K., A. Shaked, and J. Sutton (1989). An outside option experiment. *Quarterly Journal of Economics 104*, 753–770.

Camerer, C. (2003). *Behavioral game theory. Experiments in strategic interaction.* Princeton University Press.

Chiu, Y. and B. Yang (1999). The outside option, threat point, and Nash bargaining solution. *Economics Letters 62*, 181–188.

Ellingsen, T. and M. Johannesson (2004a). Is there a holdup problem? *Scandinavian Journal of Economics 106*, 475–494.

Ellingsen, T. and M. Johannesson (2004b). Promises, threats, and fairness. *Economic Journal 114*, 397–420.

Erlei, M. and J. Siemer (2004, April). Endogenous property rights in a hold-up experiment. Working Paper.

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

Fehr, E., S. Kremhelmer, and K. Schmidt (2004). Fairness and the optimal allocation of ownership rights. Working Paper.

Fehr, E. and K. M. Schmidt (2002). Theories of fairness and reciprocity - Evidence and economic applications. In M. Dewatripont, L. Hansen, and S. Turnovsky (Eds.), *Advances in Economics and Econometrics - 8th World Congress, Econometric Society Monographs.* Cambridge University Press.

Gantner, A., W. Güth, and M. Königstein (2001). Equitable choices in bargaining games with advance production. *Journal of Economic Behavior and Organization 46*, 209–225.

Grout, P. (1984). Investment and wages in the absence of binding contracts: A Nash bargaining approach. *Econometrica 52*, 449–460.

Hackett, S. (1993). Incomplete contracting: A laboratory experimental analysis. *Economic Inquiry 31*, 274–297.

Hackett, S. (1994). Is relational exchange possible in the absence of reputations and repeated contract? *Journal of Law, Economics and Organization 10*, 360–389.

Hashimoto, M. (1981). Firm-specific human capital as shared investment. *American Economic Review 71*, 475–482.

Kahn, L. and J. Murnighan (1993). A general experiment on bargaining in demand games with outside options. *American Economic Review 83*, 1260–1280.

Kessler, A. and C. Lülfesmann (2000). The theory of human capital revisited: on the interaction of general and specific investments. CEPR Discussion paper no. 2533.

Knez, M. and C. Camerer (1995). Outside options and social comparison in three-player ultimatum games experiments. *Games and Economic Behavior 10*, 65–94.

Königstein, M. (2000). *Equity, efficiency and evolutionary stability in bargaining games with joint production*. Berlin Heidelberg: Springer Verlag.

Königstein, M. (2001). Hold-up, inefficiency and fairness - An experimental study. Working paper, Humboldt-University at Berlin.

Leuven, E. (2005). The economics of training: A survey of the literature. forthcoming in: Journal of Economic Surveys.

MacLeod, W. and J. Malcomson (1993a). Investment, holdup and the form of market contracts. *American Economic Review 83*, 811–837.

MacLeod, W. and J. Malcomson (1993b). Specific investment and wage profiles in labour markets. *European Economic Review 37*, 343–354.

Malcomson, J. (1997). Contracts, holdup, and labor markets. *Journal of Economic Literature 35*, 1916–1957.

Malcomson, J. (1999). *Individual employment contracts*, Volume 3 of *Handbook of Labor Economics*, Chapter 35, pp. 2391–2372. North Holland.

Manzini, P. (1998). Game-theoretic models of wage bargaining. *Journal of Economic Surveys 12*, 1–41.

McDonald, I. and R. Solow (1981). Wage bargaining and employment. *American Economic Review 71*, 896–908.

Oosterbeek, H., J. Sonnemans, and S. van Velzen (2003). The need for marriage contracts: An experimental study. *Journal of Population Economics 16*, 431–453.

Roth, A. (1995). Bargaining experiments. In J. Kagel and A. Roth (Eds.), *Handbook of experimental economics*, Chapter 4, pp. 253–348. Princeton University Press, Princeton NJ.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica 50*, 97–109.

Scaramozzino, P. (1991). Bargaining with outside options: Wages and employment in UK manufacturing 1974-1982. *Economic Journal 101*, 331–342.

Sloof, R. (2000). Finite horizon bargaining with outside options and threat points. Forthcoming in: Theory and Decision.

Sloof, R., J. Sonnemans, and H. Oosterbeek (2004). Specific investments, holdup, and the outside option principle. *European Economic Review 48*, 1399–1410.

Sonnemans, J., H. Oosterbeek, and R. Sloof (2001). On the relation between asset ownership and specific investments. *Economic Journal 111*, 791–820.

Stevens, M. (1994). A theoretical model of on-the-job training with imperfect competition. *Oxford Economic Papers 46*, 537–562.

Stole, L. A. and J. Zwiebel (1996). Intra-firm bargaining under non-binding contracts. *Review of Economic Studies 63*, 375–410.

Figure 1. Frequency distributions of individual mean investment levels by treatment