The Economic Journal, **111** (*October*), 791–820. © Royal Economic Society 2001. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 350 Main Street, Malden, MA 02148, USA.

ON THE RELATION BETWEEN ASSET OWNERSHIP AND SPECIFIC INVESTMENTS*

Joep Sonnemans, Hessel Oosterbeek and Randolph Sloof

Experimental results are presented for a simplified version of Hart's (1995) theory of the firm. Theory predicts that investment levels remain constant when investors' no-trade pay-offs increase, if these pay-offs are threat points. While they may decrease when no-trade pay-offs are outside options. Our results support these predictions in a relative sense. Average investment levels exceed the predicted level. Actual investment behaviour is consistent with the outcomes of the bargaining stage. The play of the game is supported by a reciprocity mechanism in which non-investors consider higher investment levels as fair behaviour which deserves a reward. Investors anticipate this.

In a series of important contributions, Grossman and Hart (1986), Hart and Moore (1990) and Hart (1995) develop the property rights theory of the firm. They make the case that the ownership structure may have important implications for incentives to invest. Within their framework, the optimal ownership structure is determined as the one which results in the smallest deviations from the socially efficient investment levels.

The general set-up of the Grossman–Hart–Moore framework is as follows. Two managers own two assets. Before they trade, both managers can choose investment levels, which affect the size of the surplus if they trade and (possibly also) their respective pay-offs if no trade occurs. After the investments are sunk, the parties negotiate about the division of the surplus. When the parties agree to trade, the ownership structure is unimportant because then both managers have access to both assets. If no trade occurs, however, the ownership structure is important as the possession of assets affects the parties' no-trade pay-offs. According to Hart (1995, p. 49), having more assets encourages investment and therefore the party whose investment is more important should own more assets.

In some recent papers, this result has been criticised (de Meza and Lockwood, 1998; Chiu, 1998; Bolton and Xu, 1999). These papers point to the fact that the driving force behind Hart's conclusion is that no-trade pay-offs have the form of *threat points*. They show that when no-trade pay-offs have the form of *outside options* instead of threat points, asset ownership may actually discourage investment. Conceptually, threat points and outside options differ in the following way. Threat points are no-trade pay-offs which parties receive in case of disagreement while bargaining continues. Outside options are no-trade pay-offs which are only available when one of the parties terminates the bargaining.

With threat points as no-trade pay-offs, the subgame perfect equilibrium of

* We gratefully acknowledge useful comments from three anonymous referees and the editor David de Meza.

the bargaining stage is that both players receive their no-trade pay-offs plus 50% of the remaining surplus.[1] This is the so-called split-the-difference (STD) solution (Binmore *et al.*, 1989). With outside options as no-trade pay-offs, the equilibrium prediction of the bargaining stage is that both players receive 50% of the total surplus,[2] unless this gives one of the parties less than her/his outside option. In that case, the player with the binding outside option receives this outside option, leaving the remaining surplus to the other player. This is the so-called deal-me-out (DMO) solution (Binmore *et al.*, 1989).

Fig. 1 exhibits the crucial difference between the bargaining solutions under these two types of no-trade pay-offs, and the relation with investment incentives. To keep things simple, we assume that only one of the players (M_1) can make an investment and that this investment is completely relation-specific.[3] Moreover, we normalise the no-trade pay-off of the other player (M_2) to zero. The left-hand panel gives the case where the no-trade pay-offs have the form of threat points. Without investment the surplus up for division equals S. When M_1's threat point equals r, the STD-solution gives M_1 a pay-off of $r + \frac{1}{2}(S - r)$. Now assume that M_1 makes a relation-specific investment which enlarges the surplus to S'. Again with r as M_1's threat point, M_1 will now receive $r + \frac{1}{2}(S' - r)$. Hence M_1's gross gain from the investment equals $\frac{1}{2}(S' - S)$. In other words, M_1 receives only half of the return to her investment. Now consider what happens if M_1 owns more assets. In Fig. 1, this comes down to giving M_1 a larger no-trade pay-off, say r^*. Repeating the previous analysis with r^* instead of r reveals that M_1's gain from her investment still equals $\frac{1}{2}(S' - S)$.

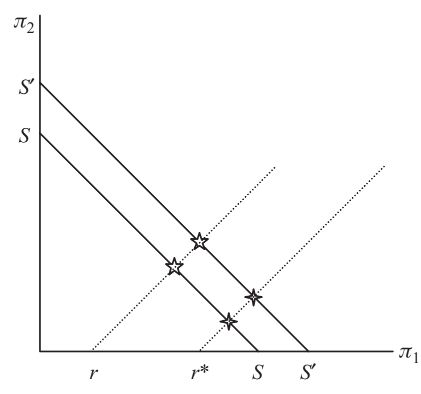

Fig. 1a. *Bargaining Outcomes with Threat Points*

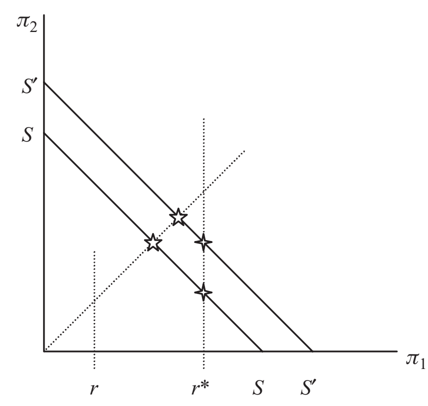

Fig. 1b. *Bargaining Outcomes with Outside Options*

[1] 50% follows if we make the standard assumption that both players have equal bargaining power. With unequal bargaining power, the remaining surplus is split in proportion to the relative bargaining powers.

[2] Again assuming equal bargaining power.

[3] In the papers referred to above the setup is more general. Both parties invest, and their investments are typically not completely relation-specific; ie they also affect the investors' no-trade payoffs. These differences do not, however, affect the crucial aspects of the analysis. In Section 2 we elaborate on this.

Hence, with no-trade pay-offs having the form of threat points, we find that owning more assets has no effects on the incentives to make the specific investment.

This is different in a situation where the no-trade pay-offs take the form of outside options; see the right-hand panel. At level r, the outside option is non-binding. Hence, the DMO-solution gives player M_1 $\frac{1}{2}S$ when there is no investment and $\frac{1}{2}S'$ when the investment has been made. Thus again M_1 gains $\frac{1}{2}(S' - S)$ from the investment. But if M_1 now owns more assets such that her no-trade pay-off is binding as is the case with r^*, both without and with investment, she receives r^*. Her gain from the investment then equals 0. Hence, with no-trade pay-offs having the form of outside options, we find that owning more assets may weaken M_1's incentives to make the relation-specific investment.

This paper reports an experiment designed to test whether the different investment incentives under the two different bargaining regimes show up in practice.[4] Experiments are very suitable to test the predictions of well articulated formal theories in a controlled environment that allows the observations to be unambiguously interpreted in relation to the theory (Roth, 1995*a*). If the experimental design complies with all the conditions set by the theory, the results of the experiment should confirm the predictions of the theory, if the theory is sound. Outside the laboratory, such tests are more difficult or even impossible to obtain because of unavailability of data and difficulties to control other, possible intervening, factors (the *ceteris paribus* condition). Subjects in our experiment play an alternating offer bargaining game that is preceded by an initial investment stage in which one of the two parties makes a relation-specific investment. We distinguish 2×3 treatments. First, there is a division into outside option and threat point bargaining games. Second, within each of the two bargaining games, the no-trade pay-off of the investor can take three levels: low, intermediate and high. Higher no-trade pay-offs reflect ownership of more assets.

Our experiment basically adds an investment stage to an alternating offer bargaining game with different forms of no-trade pay-offs. Predictions regarding investment behaviour are based on the premise that the bargaining stage results in the divisions reflected in Fig. 1. Of course, when subjects in a bargaining experiment without our initial investment stage reach bargaining outcomes that deviate substantially from these predicted outcomes, one might question the usefulness of adding an investment stage. Experimental studies by Binmore *et al.* (1989, 1991), however, provide ample support that actual bargaining outcomes are in line with the predicted outcomes. Subjects in their experiments do, in general, recognise the (subtle) difference between threat points and outside options in alternating offer bargaining games. Moreover,

[4] We analyse a game where the player who has positive no-trade pay-offs is also the player who invests. In a companion paper (Sloof *et al.*, 2000), we consider the case where only the non-investing player has positive no-trade pay-offs. There the focus is on the operation of the outside option principle as a solution to the holdup problem. Appendix A provides a brief summary of the results of that paper.

their results lend some support to the theoretical predictions about bargaining behaviour under the two bargaining regimes. One might therefore reasonably expect that also the theoretical relationship between investment behaviour and the form of the no-trade pay-offs will appear in the laboratory. From this perspective, our focus on a set-up that extends the one of Binmore *et al.* with an initial investment stage seems justified.

A number of experimental papers already study the two-stage nested bargaining game in which parties bargain over the division of a surplus created by advance investments (Gantner *et al.*, 1997; Hackett, 1993, 1994; Königstein, 1997; Oosterbeek *et al.*, 1999; Ellingsen and Johannesson, 2000). In all these studies the focus is mainly on how (relative) investment levels affect subjects' subsequent bargaining behaviour and the ultimate division of the surplus. Generally, it is concluded that 'endogeneity matters'. The sunk investments made at the earlier stage do have a significant influence on bargaining behaviour at the later stage, and thus the final division of the surplus. It appears that, typically, concepts of equity, fairness and reciprocity play a role when subjects divide a surplus. In particular, subjects take also the (relative) investment levels into account besides their gross pay-offs when bargaining over a particular division. This contradicts with subgame perfection, which predicts that sunk investments do not affect subsequent bargaining behaviour.[5] As noted, all these studies focus mainly on the 'one way' influence of how different (sunk) investment levels affect subsequent bargaining behaviour. The primary focus of this paper is the influence in the 'other' direction: the influence of different bargaining situations on initial investment levels.

The remainder of this paper is organised as follows. Section 1 summarises Hart's model and the variation proposed by de Meza and Lockwood, and establishes that the simple set-up described above captures the essential difference between the two approaches. Section 2 presents the details of the experimental design. Section 3 presents and discusses the experimental results. Section 4 summarises and concludes.

1. Theory

This section briefly presents a simplified version of Hart's model (1995) and the variation proposed by de Meza and Lockwood (1998), and then discusses the key difference between the two.

Consider two managers (M_1 and M_2) operating two assets (a_1 and a_2). M_2 in combination with a_2 supplies a single unit of input (a widget) to M_1. M_1 uses this input together with a_1 to produce a final output. Before trade takes place, M_1 can make a specific investment which makes the assets more

[5] See, however, Ellingsen and Robles (2000) and Tröger (2000) for recent theoretical contributions which use an evolutionary approach to explain why sunk cost may affect bargaining outcomes.

productive.[6] M_1's investment affects M_1's revenue of selling the product. Only when the investment costs are sunk, is uncertainty about the exact type of widget needed resolved. This makes it impossible for the parties to contract a price for the widget *before* the investment is sunk. Three ('leading') ownership structures are distinguished: non-integration where M_1 owns a_1 and M_2 owns a_2; type 1 integration where M_1 owns both assets, and type 2 integration where M_2 owns both assets. Formally, the ownership structure is represented by A which is the set of assets to which M_1 has access; $A = \{a_1\}$ under non-integration, $A = \{a_1, a_2\}$ under type 1 integration and $A = \varnothing$ under type 2 integration.

M_1's investment is denoted by i, which represents both the level and cost of the investment. If M_1 and M_2 agree to trade, M_1's revenue is denoted $R(i)$ (with $R'(i) > 0$ and $R''(i) < 0$), which is the price for which M_1 can sell her output if she could use M_2's input. The *ex post* pay-off of M_1 equals $R(i) - p$, where p is the price of the widget agreed between M_1 and M_2. M_1's *ex ante* pay-off equals $R(i) - p - i$. If M_1 and M_2 do not trade, M_1's revenue is denoted by $r(i; A)$ (with $r'(i; A) \geqslant 0$ and $r''(i; A) < 0$). The *ex post* pay-off of M_1 when there is no trade with M_2 then equals $r(i; A) - p^s$, where p^s is the spot market price of a widget (which is then not especially accommodated to M_1's needs).

For M_2, production cost in the case when trade between M_1 and M_2 takes place equals C. M_2's payoff then equals $p - C$ (*ex post* and *ex ante*). If M_1 and M_2 do not trade, M_2's cost equals $c(B)$, where $B = \{a_1, a_2\}\backslash A$. The *ex post* pay-off of M_2 then equals $p^s - c(B)$.

The *ex post* surplus if there is trade between M_1 and M_2 equals $R(i) - C$, while the *ex post* surplus without trade between M_1 and M_2 is equal to $r(i; A) - c(B)$. To express the notion that investment i is relation-specific, it is assumed that the *ex post* surplus in case of trade between M_1 and M_2 exceeds the *ex post* surplus when there is no trade between M_1 and M_2. That is: $R(i) - C > r(i; A) - c(B)$, for all i, A and B. It is further assumed that specificity also holds in a marginal sense, meaning that M_1's marginal returns on her investment are larger the more assets she has access to.[7] Formally:

$$R'(i) > r'(i; a_1, a_2) \geqslant r'(i; a_1) \geqslant r'(i; \varnothing) \quad \forall\, 0 < i < \infty \tag{1}$$

Given these assumptions the first-best level of investment (i^*) is easily derived as: $R'(i^*) = 1$. The first-best level will, in general, however, not be achieved. This follows because M_1 chooses her investment level in light of her own pay-offs, which depend on the outcome of the subsequent bargaining stage in which the price of the widget, p, is determined.

[6] In Hart's presentation, both managers make an investment decision. For our purposes, however, nothing essential is lost when we restrict attention to one investor. We are interested in the effect of owning more assets on incentives to invest and not in the optimum ownership structure *per se.* Following Hart's terminology, we study the case where M_2's investment is unproductive. Alternatively, our set-up can be interpreted as M_1's investment decision problem given the result of M_2's investment decision.

[7] The fact that $R'(i) > r'(i; a_1, a_2)$ reflects that under trade M_1 has also access to M_2's human capital, while there is no access to that asset if there is no trade and M_1 owns both assets.

In Hart's model, the pay-offs in case of trade result from dividing the *ex post* gains from trade equally (split-the-difference). That is, both M_1 and M_2 receive 50% of $[(R-C)-(r-c)]$. Each party receives her/his no-trade pay-offs and the remaining surplus is divided equally. Net *ex ante* payoffs for M_1 are then equal to

$$\pi = r - p^s + \tfrac{1}{2}[(R-C)-(r-c)] - i = \tfrac{1}{2}R + \tfrac{1}{2}r - \tfrac{1}{2}C + \tfrac{1}{2}c - p^s - i$$

Given this bargaining solution, M_1 will choose the investment level i for which

$$\tfrac{1}{2}R'(i) + \tfrac{1}{2}r'(i;\, A) = 1 \tag{2}$$

Given inequality (1) and the fact that $R''<0$ and $r''<0$, this results in $i^* > i_1 \geqslant i_0 \geqslant i_2$ (where inequalities are strict when (1) holds with strict inequalities). Here i_0 denotes M_1's investment level under non-integration, i_1 is M_1's investment level under type 1 integration and i_2 is M_1's investment level under type 2 integration. This establishes that owning more assets always results in an investment level which is closer to the first-best level.

de Meza and Lockwood (1998) draw attention to the fact that the exact form of no-trade pay-offs is crucial for Hart's result. They consider the situation where no-trade pay-offs have the form of outside options rather than threat points. In an alternating offer bargaining game with outside options, subgame perfection predicts that each party receives the best of her/his outside option and the share s/he would obtain in the absence of outside options. This is the deal-me-out solution. Again, it is possible to determine the optimal investment level for M_1 under different ownership structures. Assuming that the non-investor's outside option does not bind, the investor's pay-off in this bargaining game equals

$$\pi = \max\{\tfrac{1}{2}(R-C),\ r - p^s\} - i$$

Given this pay-off, M_1's optimum investment level is determined either (if the investor's outside option is non-binding) by

$$\tfrac{1}{2}R'(i) = 1 \tag{3a}$$

or (if the investor's outside option is binding) by

$$r'(i;\, A) = 1 \tag{3b}$$

The relation between asset ownership and investment levels now depends on how asset ownership affects whether (3*a*) or (3*b*) applies. Following de Meza and Lockwood, we assume that the total return to investment in the outside option is (weakly) increasing in the number of assets owned: $r(i;\, a_1,\, a_2) \geqslant r(i;\, a_1) \geqslant r(i;\, \varnothing)$. It can now happen that a shift from type 2 integration to no integration, or from no integration to type 1 integration, turns M_1's outside option from non-binding into binding. M_1's optimum investment level then switches from the solution to (3*a*) to the solution to (3*b*). If $R' > 2r'$ this results in a drop of M_1's investment level. Thus, owning more assets may result in a lower level of investment. This is in sharp contrast with the result for the case where no-trade pay-offs have the form of threat points.

Above, we claimed that the essential mechanism causing the differences between the models of Hart (1995) and de Meza and Lockwood (1998) also shows up when the investment is completely specific. Complete specificity implies that the investment has no effect on the no-trade pay-offs, hence: $r'(i;\ A) = 0$ for all A. Imposing this condition, (2) reduces to (3a), implying that, in the case of threat points, the investment level is independent of the ownership structure. Complete specificity also affects (3b) in the sense that $r'(i;\ A)$ cannot be equal to unity (the first-order condition cannot be met with equality). Given that investment involves costs and investment has to be nonnegative, the optimum investment level in case of a binding outside option equals zero (second-order condition).

2. Experimental Design and Hypotheses

Our design covers 2×3 treatments, corresponding to two bargaining games and three no-trade pay-off levels. We refer to the two bargaining games as the *OO-game* and the *TP-game*. In each single session, only one bargaining treatment was considered. We ran two sessions per bargaining treatment, such that we had four sessions in total. Overall 80 subjects participated in the experiment, 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 60 Dutch guilders (approximately US\$ $28\frac{1}{2}$) in about two and a half hours.

In the next subsection, we discuss the basic set-up of each experimental session. Subsequently, we describe how the bargaining and the investment stage were framed and presented to the subjects. Finally, the hypotheses that follow from our parameter choices are summarised in Subsection 2.3.

2.1. *Basic Set-up of a Session*

In each session, pairs of subjects repeatedly played a nested bargaining game. Specifically, each session contained 18 periods, and each single period consisted of a single play of a two-stage game. In the first stage, a subject with the role of player M_1 chooses an investment level; in the second stage, this subject together with a subject who has been assigned the role of player M_2 bargain over the division of the surplus. All subjects thus played 18 games. Half of the 20 subjects that participated in a session were assigned the role of M_1, the remaining 10 were assigned the role of M_2. Each participant kept the same role during the whole session. (These roles were communicated only after the complete instructions were read and understood.) In each single period, all M_1s were paired anonymously to a different M_2. We used a rotating scheme to ensure that the same subjects were not matched more than once during the first nine periods, and a different rotating scheme to ensure the same for the last nine periods. Within the two intervals of nine periods, the rotating schemes also ensured that a subject x was never matched with a subject y who met a subject previously matched with subject x. The subjects were explicitly informed about this matching procedure. We did this to rule out any reputational considerations.

The experiment was computerised.[8] Subjects started with on-screen instructions. All subjects had to answer some questions correctly before the experiment started. For example, they had to calculate the earnings of subjects for some hypothetical situations. Subjects also received a summary of the instructions on paper.[9] The instructions and the experiment were phrased as neutrally as possible; words like opponent, game, player, manager, buyer or seller were avoided. At the start of the first period, all subjects received a message informing them about their role. After the subjects had played 18 games, they filled out a short questionnaire. At the end of the experiment, the earned experimental points were exchanged for money. Subjects were paid individually and discretely.

To operationalise the experiment, concrete choices had to be made regarding:

- the feasible range of investment levels
- the form of revenue function $R(i)$
- the cost of investment, and
- the values of the no-trade pay-offs for different amounts of assets owned by M_1.

First, because our main focus is on investment behaviour, we allowed for enough variation in possible investment levels. Following Hackett (1993, 1994), investments could take any integer value between 0 and 100. Second, to keep the pay-off structure as simple as possible for the subjects, we chose $R(i)$ to be linear and the costs of investment to be quadratic in *i*. The linear revenue function has the general form $R(i) = V + v\text{i}$, with V the base amount and v the constant increment resulting from investment. In the OO-game, we set $V = 10{,}000$ and $v = 100$, along with the following values for the no-trade pay-offs: $r(\emptyset) = 1{,}800$, $r(a_1) = 6{,}800$ and $r(a_1, a_2) = 7{,}800$. A level of 1,800 is always non-binding, hence M_1 will collect half of the surplus $(5{,}000 + 50i)$. Given the quadratic cost function, i^2, the optimum investment level for M_1 then equals 25, which is exactly half of the socially optimum level of 50. No-trade pay-off levels of 6,800 and 7,800 are, given $R(i)$, both binding.[10] Half of the surplus $R(i) = 10{,}000 + 100i$ exceeds 6,800 (7,800) if i exceeds 36 (56), but then the net pay-off for M_1 would be negative. Hence, in these cases, M_1 will receive the outside options and investment will be set

[8] The experiments took place in the laboratory of the CREED research centre in Amsterdam, in which subjects are separated through cubicles. Subjects do not know with whom they are paired.

[9] A translation of this summary in English can be found on http://www.fee.uva.nl/creed/people/joeps.htm.

[10] At $r = 5{,}625$ the outside option turns from non-binding into binding. There the investor is indifferent between investing 25 and receiving half of the gross surplus, and investing zero and receiving the outside option.

equal to zero.[11] In the TP-game, M_1 always receives half of the surplus created by the investment. Hence, in that game, the optimum investment level equals 25 irrespective of the level of the no-trade pay-off. Furthermore, we normalised such that $C = c = p^s = 0$.

In the TP-game, M_1 receives the no-trade pay-off in case of delay of agreement. As a result, the joint cost of delay would be different in the two bargaining games if the total surplus would be the same. To enhance comparability between the OO-game and the TP-game, we have added M_1's no-trade pay-off to the base amount of $R(i)$. Thus, in the TP-game, we have $V = 10{,}000 + r$. This facilitates the comparison of delay of agreement under the two types of bargaining games (Knez and Camerer 1995, p. 84). Note that under the TP-game, investment incentives are not affected by the value of the base amount.

The conversion rate used in the experiments was 1 guilder for 2,500 experimental points. At the time that the experiments were run (June 1999), one guilder was about 48 US cents, such that 1 dollar corresponded to about 5,200 points.

As in Binmore *et al.* (1998), the three different values of the no-trade pay-offs were considered *within* one session. In each session, we used the same ordering of r's over the 18 periods. Subjects were told how the ordering was generated (each of the three values of r had an equal chance of $\frac{1}{3}$ in each period), but were not told what this ordering was. At the beginning of each period, they were simply informed about the value of r that applied in that period. In each period, all ten pairs were confronted with the same value of r. The fixed sequence of rs used was not ordered in an ascending or descending order as in Binmore *et al.* (1998). Rather, to allow for additional test possibilities of comparative statics from period to period, we used a sequence in which frequent upward and downward changes in r occurred (cf. Appendix B). To investigate whether subjects play the game differently in the first half of the experiment from the second half, for instance because they learn during the experiment how to play the game, the 18 periods are divided into two blocks of 9 in which the three different values of r occur with the same frequency. The two different rotating schemes for the two blocks of 9 periods facilitate such a test. Moreover, the last three periods include exactly one of each level of no-trade pay-offs, so that these final periods can also be considered separately.

The last important element common to all sessions is that we provided the subjects with an initial endowment. Subjects with role M_2 received 60,000 points (\$11$\frac{1}{2}$) at the beginning of the experiment, M_1s received 10,000 points (about \$1.90). We used initial endowments for two reasons. First, we wanted M_1s to have already some amount to spend when they had to take their first

[11] The intermediate level of 6,800 and high level of 7,800 are relatively close to each other. This is caused by the fact that the intermediate level cannot be too close to the threshold level of 5625 and that the high level cannot be too high because otherwise subjects with the role of M_2 would earn close to nothing in this treatment. Although the values 6,800 and 7,800 are so close together that one might expect subjects not to behave differently for these two values (thereby confirming theory), results below prove otherwise.

investment decision. Otherwise, to avoid an immediate debt, they may have felt somewhat reluctant to invest. Second, asymmetric initial endowments were needed to equalise at least somewhat the unequal (equilibrium) pay-offs M_1s and M_2s obtain in the game. As initial endowments may lead to undesirable wealth effects, we used the same endowments in both bargaining treatments to facilitate the comparison between them. Wealth effects then may occur, but are expected to be the same in both bargaining treatments. The actual endowments were chosen such that, over both treatments, M_1s and M_2s theoretically would earn about the same.

Note that the provision of asymmetric initial endowments may create an environment where the theory is more likely to 'work'. Typically, game theoretical predictions fare badly when they prescribe an unequal division of the surplus as the outcome of the bargaining, because in bilateral bargaining experiments subjects usually agree to divide the surplus fairly equally (Roth, 1995*b*). The game prescribes an unequal division (when the outside option binds), and is thus likely not to be supported in the lab. Asymmetric endowments may counteract the tendency to distribute the period pay-offs more equally, and may therefore give the theory a better chance.

2.2. *Framing of the Bargaining and the Investment Stage*

We make use of the well-known Rubinstein alternating offer bargaining game in framing the bargaining stage. Instead of working with a single pie that shrinks over time by some discount factor, we use a multiple pie framework in which one pie vanishes in each round of disagreement. The costs of delay of agreement are thus modelled in terms of foregone trade opportunities. (Bolton and Whinston, 1993). See also Chiu (1998), MacLeod and Malcomson (1993) and Sloof *et al.* (2000).)

The general characteristics of our bargaining stage that apply to both the OO-game and the TP-game can be described as follows. First, the bargaining stage consists of exactly ten rounds. Second, in each of these ten rounds there is a pie to be divided between M_1 and M_2. The size of the round-pies equals $R(i)/10$. That is, the gross surplus $R(i)$ obtained when the players immediately reach agreement is spread evenly over the ten bargaining rounds. Third, M_1 and M_2 alternate in making offers of how to split the ten round-pies. M_2 always makes the proposal in the first round. Hence, in all odd rounds, he formulates the proposal, while, in all even rounds, M_1 has the opportunity to make a proposal. Fourth, as soon as agreement is reached the pie of the current round, as well as the pies of *all remaining* rounds, are divided according to the proposal agreed upon. For example, if the players agree on an equal split in round 3, all the eight pies from round 3 and onwards till round 10 are divided equally. In case the players have not reached agreement before or in the final period (and they also did not choose to break off the negotiations in the OO-game), the bargaining ends and both agents receive nothing.

The OO-game and the TP-game differ in the possible responses to a proposal and in the pay-offs associated with these responses. In the OO-game, if one party makes an offer, the other party can react in *three* different ways: accept the offer, disagree and formulate a counteroffer in the next bargaining round, or unilaterally quit the bargaining by opting out. If a proposal is accepted, the players receive pay-offs according to this proposal. If there is disagreement and a counteroffer is formulated, both parties receive nothing during the round of disagreement. If the responder opts out, M_1 receives her no-trade pay-off. With the total surplus divided in ten equally sized round-pies, these no-trade pay-offs are also divided into ten pieces. Opting out in round t thus results in a pay-off for M_1 equal to $(11 - t)\,r/10$. In that case, M_2 receives nothing. Parties then cannot return to the bargaining table.

In the TP-game, the responder can only react in *two* different ways to an offer: accept or disagree and formulate a counteroffer. Hence, opting out unilaterally is not possible here. If a proposal is accepted, the players receive pay-offs according to this proposal. If there is disagreement and a counteroffer is formulated, the pay-offs during the round of disagreement are $r/10$ for M_1 and 0 for M_2.

Finally, we discuss the framing of the investment stage. At the beginning of each period, subjects were informed about both the size of the base round-pie and the value of the no-trade pay-off. In the OO-game, the size of the base round-pie equalled 1,000 experimental points, in the TP-game it was equal to $(1,000 + r/10)$. (Recall that, within one period, all ten pairs were confronted with the same value of r, and thus with the same base round-pie.) Subsequently, M_1 was asked how much she wanted to add to the base round-pie. Thus, instead of choosing the investment level i directly, M_1 in effect chose the amount $10i$. For each possible choice of $10i$ between 0 and 1,000, the costs of this particular choice (ie i^2) could be read directly from a table that was handed out to *all* subjects (thus also to M_2s). The size of the actual round pies was then set at the sum of the base round pie and the amount $10i$ chosen in the first stage. The game then continued to the second stage in which the two subjects bargained over the division of the ten actual round pies, as described above.

2.3. *Hypotheses*

Equilibrium predictions based on subgame perfection are summarised in Table 1. The bargaining outcome is stated as the share M_1 (as investor) obtains in equilibrium.

The predictions in Table 1 lead to hypotheses regarding investment behaviour and bargaining behaviour. With respect to investment behaviour, we formulate hypotheses based on comparative statics predictions and hypotheses based on point predictions.

Investment behaviour

- Under the OO-game, investment levels decrease when the no-trade pay-off goes from low to intermediate or high.

Table 1
Equilibrium Predictions

		OO-game			TP-game
		$r = 1{,}800$	$r = 6{,}800$	$r = 7{,}800$	All r
Stage 1	Investment	25	0	0	25
Stage 2	Outcome	*DMO*: $\max\{r, \frac{1}{2}(10{,}000 + 100i)\}$			*STD*: $r + \frac{1}{2}(10{,}000 + 100i)$
	Agreement	Immediate			Immediate

Remark: the socially efficient investment level equals 50

- Under the TP game, investment levels are independent of the value of the no-trade pay-off.
- Under the OO-game, investment equals 25 when $r = 1{,}800$, and 0 when $r = 6{,}800$ or $r = 7{,}800$.
- Under the TP-game, investment equals 25 for all r.

Bargaining behaviour

- Under the OO-game, M_1 receives 50% of the surplus created by the investment when the no-trade pay-off is non-binding.
- Under the OO-game, M_1 receives 0% of the surplus created by the investment when the no-trade pay-off is binding.
- Under the TP-game, M_1 always receives 50% of the surplus created by the investment.
- There is no delay of agreement.

3. Results

Our findings are presented in seven results. The presentation is divided into two parts which deal respectively with investment behaviour and bargaining behaviour. Actual investment behaviour did not vary significantly between the different sessions of the same treatment.[12] Therefore, we have pooled the observations from the sessions that considered the same bargaining game.

3.1. *Investment Behaviour*

Our first result concerns the comparative statics of investment behaviour.

[12] Six Mann–Whitney rank-sum tests are performed to compare mean individual investment levels conditional on the value of the no-trade pay-off. No significant differences between similar sessions are found at the 5% level.

RESULT 1 *In the OO-game, average investment levels are constant over different values of the no-trade pay-off. In the TP-game, average investment levels increase when the no-trade pay-off increases.*

Evidence for this result is obtained from the top panel of Table 2, which reports mean investment levels by bargaining game and no-trade pay-off. Statistical tests are based on the average investment levels of individuals (rather than on separate investment decisions). Within a column (of a panel), we compare average investment levels from the same investors for different levels of the no-trade pay-offs. For the OO-game, Wilcoxon sign-rank tests for matched pairs do not reject the hypothesis of equality of the investment levels for different values of the no-trade pay-offs. For the TP-game, on the other hand, sign-rank tests do reject the hypothesis of equal investment levels for different values of the no-trade pay-offs. Thus, under the TP-game, there is a significant pattern for the investment levels to rise with the value of the no-trade pay-offs. Notice that, although the no-trade pay-off levels of 6,800 and 7,800 may seem fairly similar, in the TP-game, subjects behave as if these values are different.[13]

These comparative statics results contrast with the theoretical predictions (the figures in square brackets). In the OO-game, investment levels remain the same when the no-trade pay-off goes up, while it is predicted that they should

Table 2
Mean Investment Levels

Periods		OO-game	TP-game
All (1-18)	1,800	28.0 [25]	^{cd}30.9 [25]
	6,800	21.5_a [0]	$^{ce}39.2_a$ [25]
	7,800	21.9_b [0]	$^{de}43.5_b$ [25]
Second half (10–18)	1,800	^{cd}25.7 [25]	^{ef}29.4 [25]
	6,800	$^{c}18.9_a$ [0]	$^{e}40.7_a$ [25]
	7,800	$^{d}19.7_b$ [0]	$^{f}40.0_b$ [25]
Final three (16–18)	1,800	24.6 [25]	^{cd}30.3 [25]
	6,800	19.2_a [0]	$^{c}40.1_a$ [25]
	7,800	20.1_b [0]	$^{d}41.5_b$ [25]

Note: superscripts c to f (within a column) indicate significant differences at the 5% level (signrank-test) (the individual letters denote which pair of numbers have been tested against the hypothesis); subscripts (within a row) indicate significant differences at the 5% level (ranksum-test) (again with the pair of valves in the same row with the same subscript letter being tested against the hypothesis); theoretical predictions are given in square brackets.

[13] A Mann–Whitney test is the nonparametric equivalent to a *t*-test, and a Wilcoxon test is the equivalent to a paired *t*-test. A *t*-test assumes that the underlying densities are normally distributed, Wilcoxon and Mann-Whitney tests make no assumptions about the exact distributions. Even when the tested variables are from normal distributions, the power-efficiency of the nonparametric tests is only slightly below that of the *t*-tests. But if (one of) the distributions differ only to a small extent from a normal distribution a nonparametric test is superior; see, for example, Winkler and Hays (1975) and Siegel and Castellan (1988). In the experimental literature, nonparametric tests are most commonly used; see Roth *et al* (1991, p. 1085) for a brief exposition about such tests in the context of experimental economics.

decrease when the no-trade pay-off changes from low to intermediate or high. In the TP-game, investment levels increase when the no-trade pay-off goes up, where it is predicted that they remain constant.

To make sure that our conclusions are not biased due to ignoring learning effects, the middle and bottom panels of Table 2 report the exact same statistics as the top panel, but now only for data from respectively the last nine and final three periods. As mentioned already in section 2, the design of the experiment was such that the first and last nine periods included the same frequencies of low, intermediate and high levels of M_1's no-trade pay-off.[14] Moreover the last three periods contained one period for each of the three levels of no-trade pay-offs each. The results in the middle and bottom panels almost exactly reproduce the results of the top panel.[15] As a further test, for each of the six treatments we regressed the investment levels on a variable which measures the time that the investor was confronted with this treatment (hence this variable ranges from 1 to 6). The time trend had a statistically significant (negative) coefficient only for the treatment of the OO-game with the no-trade pay-off equal to 6,800. As can be seen from the results in Table 2, however, the differences in mean investment levels for this treatment in the three panels are small. Given the results from these checks, it is safe to conclude that Result 1 is not contaminated by learning effects.

Further evidence that investors do not respond to changes in the no-trade pay-off in accordance with the subgame perfect predictions can be obtained from the period-to-period changes of the no-trade pay-offs (cf. Appendix B). Each investor makes 18 investment decisions. Hence, for each individual investor, we observe 17 (potential) changes in the investment level chosen. Depending on the no-trade pay-offs in two adjacent periods, we then predict that the investment level will increase, decrease or remain the same. These predictions can be confronted with the actually observed changes. The results in Table B1 (in Appendix B) reveal that, in most cases, a majority of the investors do not adjust their investment level in the direction predicted. Notice from this table also that there is no improvement during the periods of the experiment in the fraction of investors that make the predicted adjustment. When we compare, for example, the rows for rounds 2–3, 8–9, and 14–15 (all relating to a change from a high to a low level of no-trade pay-off), there are 9 (4+5) investors in round 3 adjusting their investment in the right direction, 16 in round 9, and again 9 in round 15. This is similar for other comparisons of the same change of no-trade pay-off across different periods.

Our second result concerns absolute investment levels.

[14] That is: in the first and second block of nine periods, subjects were confronted with three low, three intermediate and three high levels of the no-trade pay-off.

[15] In many cases, the investment levels for the same treatment differ significantly across the panels in Table 2 (by a Wilcoxon signrank-test). Notice, however, that there is no clear pattern of increasing or decreasing investment levels during the periods of the experiment. Between the top and middle panel the investment levels in the OO game go down, but between the middle and bottom panel these go up again for the higher levels of no-trade pay-offs.

RESULT 2 *In all treatments, average investment levels are below the socially efficient level of 50. Average investment levels are, however, always above the level predicted by subgame perfection.*

Result 2 follows immediately from comparing the realised mean investment levels from Table 2 with the theoretically predicted levels. More detailed information regarding investment levels is given in Fig. 2. This figure show plots of the distribution of investment levels by treatment and M_1's no-trade pay-off: 2*a* and 2*b* relate to separate investment decisions as level of observation while 2*c* and 2*d* relate to average investment levels per individual investor. The line in the middle of the boxes gives the mean, while the upper and lower sides of the boxes are located one standard deviation from the mean. The lines connected to the boxes extend to the largest and smallest data point observed.

The general picture that emerges from these figures is the same. In all cases (i.e. bargaining game/no-trade pay-off combinations), a vast majority of the investment levels falls short of the socially optimum level of 50. Indeed, considered at the level of investment decisions, 76% of the 720 observed decisions resulted in underinvestment; 16% of the investments were exactly at the efficient level; and only a small minority of observed investment decisions (8%) is above that socially efficient level. For these latter cases, the mean

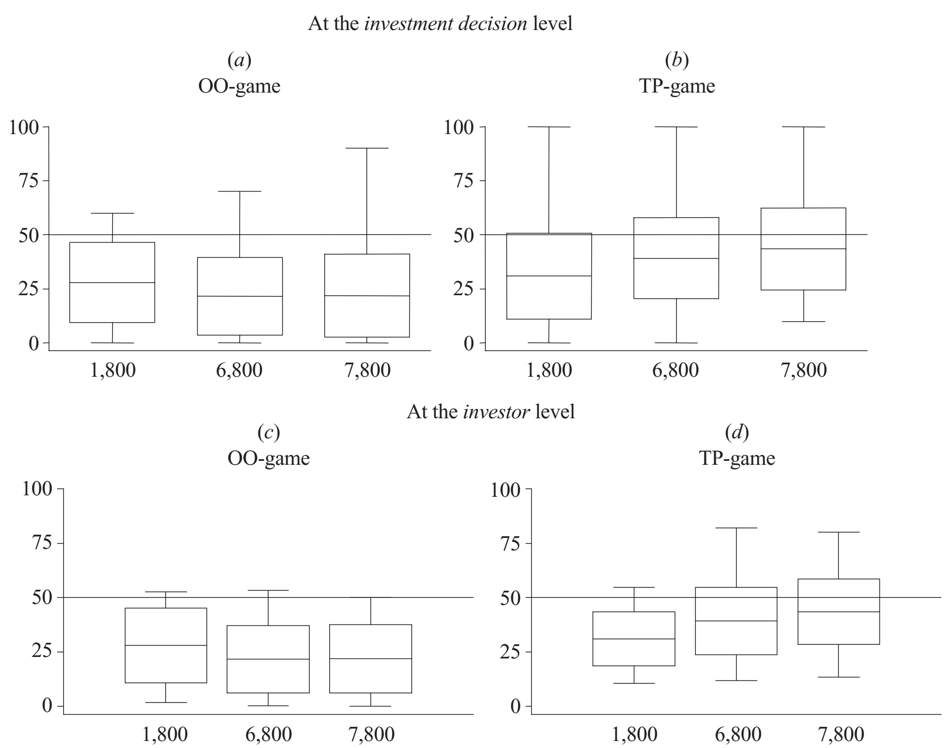

Fig. 2. *Box Plots of Distribution of Investment Levels by Bargaining Game and No-trade Pay-offs*

investment rate is 73. Although over-investment thus may occur, it does not seem to pose a serious problem in practice. Of the 720 investment decisions only 52 (7%) coincide with the theoretically predicted level. M_1s typically over-invest from an individually rational point of view. The similarity between the results in the three panels of Table 2 and the earlier mentioned insignificance of 5 out of 6 coefficients in regressions of investment on a time trend, indicate that Result 2 is not diluted by learning effects.

Our final result regarding investment behaviour compares investment levels under the two bargaining regimes.

RESULT 3 *Average investment levels are higher under the TP-game than under the OO-game.*

Support for Result 3 is again provided in Table 2. For given levels of the no-trade pay-offs, the average investment levels of individuals are higher under the TP-game than under the OO-game. This difference is not significant when the no-trade pay-off equals 1,800, but is significant for no-trade pay-offs equal to 6,800 and 7,800. This is true for the results in all three panels, indicating that learning effects do not change this finding. Result 3 is in line with theoretical predictions. Thus, while the relation between investment levels and no-trade pay-offs for a given bargaining regime contradicts theoretical predictions (cf. Result 1), the relation between investment levels and bargaining regime for a given level of the no-trade pay-off is in accordance with theoretical predictions.[16] This means that a key point of de Meza and Lockwood (1998) is corroborated by our results: when no-trade pay-offs have the form of outside options rather than threat points, there is a depressing effect on incentives to invest, and this depressing effect becomes larger when no-trade pay-offs become larger. However, Results 1 and 2 indicate that subgame perfection does not give a full explanation for observed investment behaviour. In subsection 3.2, we turn to the analysis of actual bargaining behaviour to see whether the deviations between predicted and realised investment levels can be explained by features of observed bargaining behaviour.

3.2. *Bargaining Behaviour*

The return M_1 actually obtains on her investment is determined by the offers finally accepted and the number of bargaining rounds needed to reach agreement. We start with presenting results on M_2's first offers and on finally accepted offers. The term offer always refers to the pay-offs for the investor (M_1). Then we turn to the number of rounds required to reach agreement.

An important reason why actual offers may differ from game-theoretical predictions is given by considerations of fairness. One of the main regularities obtained from a vast number of bargaining experiments is that first offers and

[16] This is supported by the period-to-period changes analysed in Appendix B. In eight instances, theory predicts a difference between the two bargaining regimes in the period-to-period changes. In seven cases, we do indeed find significant differences in the predicted directions.

final agreements typically deviate from equilibrium predictions in the direction of an equal split (Bolton, 1991; Ochs and Roth, 1989; Roth, 1995*b*). This may point at a concern for fairness of the proposer, but it may also simply reflect proposers' strategic considerations in view of the anticipated reaction of responders. Indeed, when the reaction of the responder is guided by the notion of being treated fairly, considerations of fairness may interact with the strategic features of the game (Prasnikar and Roth, 1992; Fehr *et al.* 1997). To complicate matters further, what people perceive as being fair may strongly be influenced by the strategic situation they find themselves in. Binmore *et al.* (1991), for instance, report that what subjects describe as being a fair bargaining outcome strongly depends on whether subjects played a bargaining game like our OO-game, or one like the TP-game. Based on a large number of experimental studies, Roth (1995*b*, p. 328) more generally concludes that 'notions of fairness are labile and appear to respond to strategic considerations'. Given these problems, it may in general be rather difficult to disentangle strategic (game-theoretic) and fairness considerations. That will not be the purpose here. We just confine ourselves to reporting descriptive statistics of the offers observed that are of interest from a fairness point of view.

The exposition in this section distinguishes between three bargaining situations: the threat point game (TP), the outside option game with binding no-trade pay-off (OO binding) and the outside option game with non-binding no-trade pay-off (OO non-binding). Whether the no-trade pay-off is binding in the outside option game or not, depends on whether $r \geqslant \frac{1}{2}(10{,}000 + 100i)$ or $r < \frac{1}{2}(10{,}000 + 100i)$.[17] Our first result in this subsection relates to the amounts of M_2's first offers and of finally agreed offers, and how these relate to the relevant theoretical benchmarks.

RESULT 4 *In all three bargaining situations, the average values of first offers and finally agreed offers are in between the DMO- and STD-solutions. In the OO-games, mean first offers are closer to the DMO-solution than to the STD-solution, while, in the TP-game, the reverse holds. Mean values of finally agreed offers are closer to the STD-solution than first offers are.*

Evidence for Result 4 is given in Tables 3a and 3b which presents the frequency distributions of the first offers (3*a*) and the finally agreed offers (3*b*). The first column reports the means of the first offers and the finally agreed offers. The next three columns give the values of three relevant benchmarks: the equal split (ES) where both players receive 50% of the gross surplus, the deal-me-out solution, and the split-the-difference solution. The outcomes predicted by subgame perfection appear in italics in the table. The fifth columns express the actual offer as a share of the gross surplus (the pie). The next six columns

[17] For the predicted levels of investment, the OO-binding situation applies for both $r = 6{,}800$ and $r = 7800$ and the OO non-binding situation for $r = 1{,}800$. As the actually observed investment levels differ substantially from theoretical predictions, this correspondence does not apply for our experimental data. But, even when investment levels differ from their equilibrium predictions, subgame perfection predicts that, in the bargaining stage, the first offer under the OO-game (TP-game) is given by DMO (STD).

Table 3*a*

Mean of First Offers and Predictions, Together with Frequency Distribution

Situation	Offer	ES	DMO	STD	Offer/ Pie		ES		DMO		STD		*n*
							↓		↓		↓		
OO binding	774	586	*737*	955	0.66	9	4	8	*4*	180	0	1	206
OO non-binding	684	663	*663*	809	0.52	32		*23*		65	3	31	154
TP-game	1,083	963	963	*1,236*	0.56	12		4		282	*11*	51	360

Remark: ES stands for equal split. Theoretical predictions are in italics. In the frequency distributions (next to last column) numbers straight below the predictions represent the number of observations that exactly equal this prediction. The numbers between them represent the number of observations that fall in between these predictions.

Table 3*b*

Mean of Finally Accepted Offers and Predictions, Together with Frequency Distribution

Situation	Offer	ES	DMO	STD	Offer/ Pie		ES		DMO		STD		*n*
							↓		↓		↓		
OO binding	810	587	*737*	955	0.69	0	0	3	*2*	162	0	1	168
OO non-binding	751	663	*663*	808	0.57	12		*11*		72	2	50	147
TP-game	1,159	963	963	*1,236*	0.60	2		2		228	*13*	109	354

See notes for Table 3a.

give the frequency distributions of the offers in terms of the three benchmarks ES, DMO and STD. In the OO binding situation, ES < DMO < STD holds, while, in the other two situations, ES = DMO < STD. For example, in the first row of Table 3*a* we read that in the OO binding situation, 9 first offers are below the ES outcome, 4 first offers are exactly equal to the ES outcome, 8 first offers are in between the ES and DMO outcomes, and so on.[18]

In the OO binding situation, the predicted DMO-solution gives M_1 her no-trade pay-off leaving the remainder to M_2. First offers below the DMO-solution are in the direction of an equal split. Apparently, this is not a common result, only 10% of the first offers are below the DMO-prediction. Moreover, no more than 2% of the first offers are exactly equal to the DMO-prediction. A vast majority of the first offers are above the DMO-prediction but below the STD-solution. Thus, players with the role of M_2 are prepared to give M_1 somewhat in excess of her no-trade pay-off. This can be either interpreted as a minimum amount needed to make M_1 prefer M_2's first offer over opting out, or as M_2 offering M_1 some positive return on her investment. The mean value of the first offer relative to the ES, DMO and STD benchmarks is in line with this.

[18] Also for the bargaining stage we ran analyses separately for all periods together, only for the last 9 periods, and only for the last three periods. If we compare the shares of offers to the total pies (the Offer/Pie columns in Tables 3*a* and 3*b*), there is no indication that subjects make different first offers or reach different agreements when they have gained experience in the experiment.

The mean offer is close to the mean DMO-solution, indicating that most of the 180 offers between STD and DMO are nearer to the predicted DMO-solution than to the STD-solution.

In the OO non-binding situation, subgame perfection predicts that the gross surplus is split equally (ES and DMO coincide here). In 21% of the cases, M_2 starts with offering less than this predicted solution. In another 15% of the cases, the first offer exactly equals the equal split while, in all other cases (65%), M_2 offers more than half of the gross surplus to M_1. The natural interpretation of this result is that M_2 is prepared to offer M_1 a return on her investment larger than 50% (which is the return in the case of an equal split).

In the TP-game, subgame perfection predicts that M_1 receives her no-trade pay-off and that the remaining surplus is divided equally. In a vast majority of the cases (83%), M_2 offers M_1 a smaller amount of the surplus than should be offered according to the STD-solution. This is a deviation in the direction of an equal split (and DMO), suggesting that M_2s are not prepared to let M_1 fully exploit her advantageous bargaining position. For the 282 cases with first offers in between the DMO and STD-solution, the first offer is substantially closer to the equal split (and DMO) solution than to the STD-solution. In fact, while under the TP-game, STD is the predicted outcome, M_2's first offers are closer to the DMO prediction. This observation corresponds with the findings of Binmore *et al.* (1991, Figs 3 and 4) for the TP-game. Of 360 first offers, 51 (14%) are, however, above the STD-solution. An explanation for these high first offers is that here M_2 is prepared to give M_1 a higher return to her investment than the 50% predicted by subgame perfection.

The finally agreed offers in the OO binding situation are slightly above the first offers. More precisely, of the 168 cases in which agreement is reached, in 150 cases, the finally agreed offer equals the first offer because of immediate agreement while, in the remaining 18 cases, the first offer is below the finally agreed offer. In 36 instances, M_1 found M_2's offer so unattractive that she opted out. These results reinforce the finding we already achieved with regard to first offers in the OO-game with binding no-trade pay-off: offers somewhat above the DMO-solution are the typical bargaining behaviour in this game.

The same pattern emerges even more strongly for the OO non-binding situation. While 36% of the first offers are below or equal to the DMO-solution, this is only true for 16% of the finally agreed offers. We also observe that the mean value of the finally agreed offers is substantially above the mean value of the first offers (751 versus 684).

For the TP-game, we also observe a substantial upward shift when going from the distribution of first offers to the distribution of finally agreed offers. Nevertheless, for over 65% of the cases, the finally agreed offer is below the theoretically predicted STD-solution, and are thus biased in the direction of an equal split. However, while the mean of the first offers in the TP-game is closer to the mean of the DMO-solution than to the mean of the STD-solution, we now find the opposite.

The finding that average finally agreed offers are higher than average first offers is not surprising. Player M_2 always makes the first offer and offers are expressed as the amount player M_1 obtains. If M_1 accepts the first offer, finally agreed offer and first offer are equal. If M_1 rejects M_2's first offer, it is likely that she does so in anticipation of a higher total pay-off or of a higher relative pay-off. In both cases, the finally agreed offer needs to exceed the first offer: in the first case, the increase of the share has to offset the losses due to delay; in the second case, this is not necessary.

In the discussion above, we related the differences between actual offers and theoretically predicted offers to possible rewards on M_1's investment and to not letting M_1 fully exploit her bargaining power. We tested these explanations more formally by regressing first offers and finally agreed offers on the level of investment and the value of M_1's no-trade pay-off. Results relates to this.

RESULT 5

(i) *In all bargaining situations, first offers and finally agreed offers give M_1 a larger return on her investment than predicted.*
(ii) *In the OO non-binding situation, there is an unpredicted positive effect of the no-trade pay-off on first and finally agreed offers.*
(iii) *In the TP-game, the impact of M_1's no-trade pay-off on first and finally agreed offers is smaller than predicted.*

Evidence for Result 5 is given in Tables 4a and 4b which contain for each of the three bargaining situations results from regressing M_2's first offer (4*a*) and

Table 4*a*

Regression Results Explaining M_2's first offers

Situation	DMO	STD	Realised	n	Adj R^2
OO binding	$\boldsymbol{r}$	$\frac{1}{2}(V+vi+r)$	$0.073V + 0.264vi + 0.811r + 5.87t$ (0.103)# (0.045) (0.143) (1.32)	206	0.35
OO non-binding	$\boldsymbol{\frac{1}{2}(V+vi)}$	$\frac{1}{2}(V+vi+r)$	$0.422V + 0.569vi + 0.209r + 1.60t$ (0.034) (0.067) (0.061) (2.19)#	154	0.47
TP-game	$\frac{1}{2}(V+vi)$	$\boldsymbol{\frac{1}{2}(V+vi+r)}$	$0.508V + 0.614vi + 0.062r + 3.26t$ (0.017) (0.027) (0.032)# (0.99)	360	0.83

Table 4*b*

Regression Results Explaining the finally agreed offer M_1 receives

Situation	DMO	STD	Realised	n	Adj R sq
OO binding	$\boldsymbol{r}$	$\frac{1}{2}(V+vi+r)$	$0.076V + 0.248vi + 0.927r + 0.78t$ (0.058)# (0.026) (0.081) (0.78)#	168	0.65
OO non-binding	$\boldsymbol{\frac{1}{2}(V+vi)}$	$\frac{1}{2}(V+vi+r)$	$0.493V + 0.710vi + 0.111r - 0.69t$ (0.020) (0.039) (0.036) (1.29)#	147	0.77
TP-game	$\frac{1}{2}(V+vi)$	$\boldsymbol{\frac{1}{2}(V+vi+r)}$	$0.543V + 0.696vi + 0.076r + 1.47t$ (0.018) (0.028) (0.033) (1.02)#	354	0.85

Remark: Theoretical predictions are bold faced. All coefficients are significant at the 5%-level, except when marked with a # which indicates no significant difference from zero at conventional levels.

the finally agreed offers (4*b*) on the level of investment and the value of the no-trade pay-off.[19,20] The first two columns give the expressions for the DMO and the STD solutions; the theoretically predicted expressions are bold faced.[21] The estimated coefficients are reported in the third column, the fourth column, the number of observations, and the last column the adjusted R^2. In addition to the base amount, the size of the investment and M_1's no-trade pay-off, the equations also include the period *t*, in which the bargaining took place as regressor. This was done to correct the estimates for possible learning effects. For the OO binding situation and for the TP-game, we find that during the course of the experiment M_2's increase their first offers. Even in these cases, however, the coefficients for V, vi and r are almost identical whether or not a period trend is included.

The first rows relate to the results for the OO binding situation. Subgame perfection predicts here that M_1 receives her no-trade pay-off and that M_2 becomes the residual claimant and thus receives the full returns on M_1's investment. The estimation results show that, on average, M_2's first offers and finally agreed offers deviate from this prediction. While the no-trade pay-off comes in with a coefficient somewhat smaller than one, M_1's investment has a positively significant coefficient. First offers and final agreements give M_1 a return on her investment of 25% whereas the prediction is a zero return. The finding that investors are compensated for sunk investments is not uncommon in the experimental economics literature; see, for instance, Ellingsen and Johannesson (2000), Hackett (1993, 1994) and Oosterbeek *et al.* (1999).

In the OO non-binding situation (second rows), first offers and finally agreed shares are significantly positively affected by the no-trade pay-off while theory predicts these to be independent of *r*. In her first offer, M_2 leaves slightly more than one half of the return to the investment to M_1. In the finally agreed amounts, M_1's share of the return of the investment rises to 71%. Thus, in this regime, two mechanisms operate in the same direction of making M_1 better off than theory predicts. M_1 receives a return on her investment exceeding one half, and M_1 receives some compensation on her non-binding outside option.

In the TP-game, the regressions for first offers and finally agreements show the same pattern: M_1 gains a return on her investment of more than 50%, but too low a weight is put on M_1's no-trade pay-off. This implies that M_1 is unable to exploit her bargaining advantage stemming from more favourable threat points. In the opposite direction operates that M_1 receives a higher return on

[19] Similar conclusions as in Result 5 are obtained when we estimate regression equations for each of the three values of the no-trade pay-off separately.

[20] Observations in which M_1 opted out in the OO-game (45 cases) and in which no agreement was reached in the TP-game (6 cases) were deleted from the regressions of finally agreed offers.

[21] Formally, the DMO and STD predictions in the case of finally agreed offers do not exactly apply for proposals made by M_1. Specifically, these two predictions have to be multiplied by $(9 - t)/(10 - t)$ to obtain M_1's equilibrium proposal in even round t (Sloof, 2000).

her investment than predicted by subgame perfection; more than 60% instead of the predicted return of 50%.[22]

Results 4 and 5 establish that first offers and finally agreed offers differ from the amounts predicted by subgame perfection. As noted in the introduction of this subsection, this is not an uncommon finding in bargaining experiments. Earlier bargaining experiments typically find deviations in the direction of an equal split. The preceding investment stage in our experiment is an additional source for deviations from subgame perfection. Giving the investor a larger return on her investment than predicted can be interpreted as reciprocal behaviour of the other player. Reciprocity refers to the motivation that a player is willing to forgo some amount of money to reward (punish) behaviour that is perceived as fair (unfair). Reciprocal behaviour is a common finding in many different experimental games (Fehr *et al.*, 1997; Fehr and Gächter, 2000). Making a larger investment than predicted can be interpreted by the non-investor as something that is fair and therefore warrants a reward in the form of a larger return than predicted. If such reciprocal behaviour is anticipated by the investor, it is optimal for the investor to invest more than predicted.

Results 4 and 5 thus provide a rationale for M_1's observed investment behaviour. First offers and finally agreed offers give M_1 a larger return on her investment than predicted. Hence, if M_1 accepts M_2's first offer, she already receives a higher return than predicted. If M_1 rejects M_2's first offer, it is likely that she does so in anticipation of higher total payoffs (which requires that the loss due to delay is offset by a larger share).[23]

Our next result relates to losses due to delay.

RESULT 6 *Agreement is not always immediate. Agreement is on average reached sooner in the OO-game than in the TP-game. In the OO-game, the average length of the bargaining game is shorter for higher values of the no-trade pay-offs. In the TP-game the average length of the bargaining game is equal for different no-trade pay-offs. When opting out is possible (OO-game), M_1 is more likely to unilaterally quit the bargaining for higher values of the no-trade pay-offs.*

Result 6 follows immediately from Table 5 which reports the mean number of bargaining rounds before agreement is reached, or before M_1 opts out. It also

[22] Comparing the regression results across the three situations considered in Table 4*a*, indicates that the three coefficients for V differ significantly (at the 5%-level) from each other. The coefficient for vi is significantly lower in the OO-game with binding outside option than in the other two regimes. The three coefficients for the no-trade pay-off all differ significantly from each other (again at the 5%-level). In Table 4*b*, the coefficients for the no-trade pay-off differ between the three regimes. The coefficient for vi is significantly lower in the first row than in the other two (which are not different from each other). The coefficients for V are all significantly different from each other.

[23] That M_1s are likely to reject first offers to receive a larger total pay-off rather than to obtain a larger share of the finally divided surplus, is indicated by the very small percentages of disadvantageous counter-offers. In the OO binding situation, 4 out of 20 counter-offers are disadvantageous; in OO non-binding, this holds in 2 out of 62 cases, and, in the TP game, in only 14 out of 217 cases. The fractions of disadvantageous counter-offers are substantially smaller than the figures reported by Roth (1995*b*).

Table 5

Mean Number of Bargaining Rounds Conditional on Outcome

	OO-game				TP-game	
	Accept		Opt out		Accept	
No trade pay-off	periods 1–9	periods 10–18	periods 1–9	periods 10–18	periods 1–9	periods 10–18
$r = 1,800$	$^{jk}1.90\ (58)_{ab}$	$1.81\ (57)_{cd}$	$^{jl}7.00\ (2)_{ef}$	$3.67\ (3)_{gh}$	$^{kl}2.68\ (60)^{*}$	$2.06\ (60)_{i}^{*}$
$r = 6,800$	$^{m}1.40\ (52)_{a}^{*}$	$^{o}1.13\ (53)_{c}^{*}$	$^{n}1.00\ (8)_{e}$	$^{p}1.29\ (7)_{g}$	$^{mn}2.40\ (60)$	$^{op}2.57\ (60)_{i}$
$r = 7,800$	$^{q}1.28\ (43)_{b}^{*}$	$^{s}1.04\ (52)_{d}^{*}$	$^{r}1.24\ (17)_{f}$	$^{t}1.00\ (8)_{h}$	$^{qr}2.45\ (60)$	$^{st}2.22\ (60)$

Remark: number of cases within parentheses. An asterisk (*) indicates a significant difference between mean number of bargaining rounds for each of the two sets of periods of first and last 9 periods according to a Mann–Whitney test at the 5% level. Subscripts (superscripts) indicate within-column (-row) significant differences between mean number of bargaining rounds according to a Mann–Whitney test at the 5% level. (The individual letters denote which pair of values are tested against the hypothesis as in Table 2).

presents the number of cases that lead to the particular outcome.[24] In the OO-game subjects can opt out. We observe 45 cases ($12\frac{1}{2}\%$) in which M_1 opts out (M_2 never chooses to opt out in the experiment).

The theoretical prediction is that agreement is reached in the first round and that opting out does not occur. These predictions are not affected by the no-trade pay-off being binding or not. We therefore do not distinguish here between non-binding and binding cases but rather distinguish between different levels of the no-trade pay-off. Obviously, actual outcomes deviate from the theoretical predictions. Besides the occurrence of opting out, we also see that, in the cases where opting out does not occur, the average number of rounds needed to reach agreement exceeds one. We also observe that in the OO-game the required number of bargaining rounds to reach agreement decreases when the value of the no-trade pay-off goes up.

In the TP-game, the average number of bargaining rounds is between 2 and 3, which is higher than in the OO-game. Also, the number of bargaining rounds in the TP-game is independent of the level of M_1's no-trade pay-off. These results point to an important difference between the OO-game and the TP-game: delay of agreement is shorter in OO-games than in TP-games. The same pattern is found in our related paper, Sloof *et al.* (2000). This is a difference not anticipated by theory. Given our findings on first offers, this difference is, however, perfectly understandable. In the OO-game, M_2's first offer is typically above the predicted DMO-solution. Hence, according to M_2's first proposal, M_1 is made better off than subgame perfection predicts, and thus has a good reason to accept this first offer immediately. In the TP-game, on the other hand, M_2's first offer is typically below the predicted STD-solution. This may give M_1 a reason to disagree and claim a larger share of the

[24] In the TP-games, on 6 occasions (twice for each level of no-trade pay-offs), the parties did not reach agreement within ten rounds. For these observations, the number of rounds is set equal to 11. Similar results are obtained if we drop these 6 cases.

(round) pie. Furthermore, when M_1 rejects M_2's first offer, she receives no payment during disagreement in the OO-game (unless she opts out), while in the TP-game she receives during each round of disagreement a pay-off equal to the no-trade pay-off (divided by 10). Owing to this compensation, M_1 may perceive it as less costly to disagree in the TP-game than in the OO-game, although our set-up ensures that the *joint* costs of disagreement are the same in all treatments.

Table 5 gives the mean numbers of bargaining rounds conditional on outcome for the first nine and last nine periods separately. These results indicate that the mean number of bargaining rounds is lower in the second half of the experiment than in the first half. In three instances, the reduction is significant. Apparently, subjects learn to avoid costly delay when they play the game. Result 6 is, however, not affected by this, as it is also supported by the outcomes for periods 1 to 9 or periods 10 to 18 separately.

Although Result 6 deviates from equilibrium predictions, it is in line with the results obtained in other experimental studies. Ashenfelter *et al.* (1992), for instance, study (unstructured) bargaining followed by binding arbitration in the case of disagreement. They find that disagreement rates are inversely related to the monetary costs of disputes. These results accord well with our finding that, under the OO-game, delay of agreement is decreasing in the value of the no-trade pay-off. Recall that our experimental set-up is such that the *joint* costs of disagreement are independent of the value of *r*, irrespective of which bargaining game is played, but under the OO-game M_1's *private* costs of delay are not. The higher the value of the no-trade pay-off, the higher the value of the forgone opportunities of M_1. One round delay is therefore more costly to M_1 when *r* is high, and therefore less likely to occur. A second result Ashenfelter *et al.* (1992) obtain is that dispute rates differ significantly across arbitration systems. Here, we obtain a similar result that delay of agreement differs significantly across bargaining regimes.

We end this subsection with a result on the optimality of observed investment levels given actual bargaining behaviour.

RESULT 7 *Observed average investment levels are all within three standard deviations of the 'optimum' levels given M_2's first offers.*

A nice feature of our design is that M_2 makes the first proposal, which M_1 can then accept or reject. Hence, M_2's proposal tells us immediately how much M_1 can earn on her investment if she accepts the first offer. Based on this, we can then calculate how much M_1 would earn if she accepts this proposal. If M_1 rejects the first offer, (at least) one round pie is wasted and at least one of the parties will be worse off than would have been the case with immediate acceptance. M_1 might reject M_2's first offer either because she expects that the subsequent bargaining rounds will result in higher pay-offs for her, or because she wants a larger share of the pies that are eventually divided between the two. In the first case, M_1's pay-off in case of immediate acceptance gives a lower bound on M_1's actual pay-off. In the latter case, M_1 is apparently not only motivated by the absolute level of her pay-offs, but also cares about the relative level of her pay-offs. If she then ends up with lower absolute pay-offs, we might say that she has sacrificed some payoff to satisfy her distributional concerns.[25] For these reasons, we believe that M_1's potential pay-off in the case of immediate acceptance gives the cleanest estimate of the returns to her investment.

We estimated regression equations with M_1's potential pay-offs in the case of immediate acceptance as dependent variable, and the level of investment and investment squared as independent variables (besides a constant term).[26] The 'optimum' level of investment then can be directly obtained from the estimated coefficients. Table 6 reports the estimation results and the optimum investment levels.

Comparing the estimated optimum investment levels from Table 6 with the actual mean investment levels reveals that these two levels are close to one another. In four out of six cases, the point estimates are within two standard deviations while, in the other two cases, the actual mean investment levels are within three standard deviations of the estimated optimum investment levels.

Table 6

Regression Results Explaining M_1's Net Earnings if M_2 First Offer Had Been Accepted

Regime	No-trade pay-offs	Coefficients	'Optimum' i	Actual i	Adj R^2
OO-game	$r = 1,800$	$4713 + 59i - 1.02i^2$	28.9	28.0	0.02
		(323) (29) (0.51)	(3.8)		
	$r = 6,800$	$6731 + 32i - 0.91i^2$	17.6	21.5	0.14
		(162) (16) (0.28)	(3.7)		
	$r = 7,800$	$7715 + 22i - 0.96i^2$	11.5	21.9	0.32
		(173) (14)# (0.23)	(4.8)		
TP-game	$r = 1,800$	$6384 + 66i - 1.06i^2$	31.1	30.9	0.67
		(173) (9) (0.09)	(1.9)		
	$r = 6,800$	$8784 + 92i - 1.40i^2$	32.9	39.2	0.47
		(336) (16) (0.17)	(2.2)		
	$r = 7,800$	$8590 + 117i - 1.49i^2$	39.3	43.5	0.42
		(537) (22) (0.21)#	(2.6)		

Remark: All coefficients are significant at the 5%-level, except when marked with # which indicates no significant difference from zero at conventional levels. Standard deviations in parenthesis. For each regression, $n = 120$.

[25] In the OO-game, there are 120 rejections of the first proposal; in 60, M_1 ends up with higher pay-offs than she would have earned in case of immediate acceptance. In the TP-game, the respective figures are 217 and 111. Thus, in about 50% of the cases, rejection of the first offers result in lower pay-offs than would have been possible.

[26] We have also estimated similar 'fixed-effect' regression equations that incorporated subject-specific dummy variables (intercepts). In contrast to the standard regression results reported in the main text, these fixed-effect regressions control for subject-specific characteristics. However, the differences from the results obtained from the standard regressions are only minor. Therefore, we only report the latter. We also estimated a specification where the linear investment term has a spline at the predicted optimum. These splines were never significant, though.

This leads to the conclusion that, although the observed investment levels are quite different from the theoretically predicted investment levels, they seem fairly optimal given actual bargaining behaviour. This suggests that the participants in the experiment who have the role of M_1 have a pretty good idea of what to expect from the bargaining stage at the time they make their investment decision. Thus, where previous experimental studies establish that bargaining outcomes are affected by a preceding investment stage, this result concerns the relation in the other direction. It shows how (differences in) anticipated bargaining outcomes affect investment decisions.

4. Conclusion

Recent contributions on the property rights theory of the firm derive contrary conclusions concerning the relation between asset ownership and incentives to make specific investments. According to Hart (1995), asset ownership encourages specific investments, while (among others) de Meza and Lockwood (1998) show that asset ownership may discourage investment. This difference is entirely due to different assumptions about the exact form that the no-trade pay-offs take. Asset ownership may discourage investment if no-trade pay-offs have the form of *outside options*, while this does not occur when no-trade pay-offs have the form of *threat points*. This paper reports about an experiment designed to test whether this difference shows up in practice.

The experimental design is based on a simplified version of Hart's model: there is only one player who invests and the investment is completely specific. In that case, subgame perfection predicts that the level of investment is not affected by the level of no-trade pay-offs if these have the form of threat points. If the no-trade pay-offs have the form of outside options, theory predicts that an increase of the outside option from a non-binding low level to a binding high level will result in a decrease of the investment level. With threat points, we find that investment levels increase (rather than remain constant) when the investor's no-trade pay-off goes up. With outside options, investment levels tend to decrease (as predicted) when the value of the no-trade pay-off increases, but this decrease is much smaller than predicted and lacks significance. Taken together, these comparative statics results support the theory in a relative sense. When no-trade pay-offs have the form of outside options rather than threat points, there is indeed a depressing effect on incentives to invest.

In all cases considered, the average investment level exceeds the predicted level. The results from the bargaining stage provide an explanation for the differences between observed and predicted investment behaviour. The non-investor's first proposal, as well as the finally agreed amounts, typically grant the investor a higher return on her investment than predicted. This bargaining behaviour is in line with other recent experimental results which point to the importance of (positive) reciprocity as a motivation. Investments above the equilibrium level can be interpreted by the non-investor as fair behaviour of the investor. This friendly behaviour of the investor warrants a reward in the form of a larger return on the investment than predicted by subgame perfection. Investors anticipate this and therefore invest more than predicted. This explains why, in all cases, actual investment levels exceed predicted investment levels. At the same time, the strategic elements which cause different predictions for the OO-game and the TP-game also remain at work. This explains why the theoretical predictions about investment levels in the two types of bargaining environments are confirmed in a relative sense. Given the outcomes of the bargaining stage, actual investment decisions are close to optimal.

University of Amsterdam, Tinbergen Institute and CREED

University of Amsterdam and Tinbergen Institute

University of Amsterdam and Tinbergen Institute

Date of receipt of first submission : January 2000
Date of receipt of final typescript : January 2001

Appendix A: Brief Summary of Sloof *et al.* (2000)

In the companion paper Sloof *et al.* (2000), we consider an experiment which is in all but two respects similar to the experiment reported in the current paper. The differences are

(i) that, in that experiment, it is the non-investor's no-trade pay-off which is positive and takes the values 1,800, 6,800 or 7,800, and
(ii) that, there, the investor is the player who makes the first offer in the bargaining stage, while, in the experiment in the current paper, the non-investor makes the first offer.

The main motivation for this set-up is that giving the non-investor a binding no-trade pay-off in an outside option bargaining environment theoretically solves the holdup underinvestment problem. With a binding outside option, the non-investor receives precisely his outside option and the investor earns the residual. Everything that is added to this residual as a result of an investment accrues to the investor, who then has the right incentives to invest.

Table A1 gives the mean individual investment levels together with the theoretical predictions. This is the analog of Table 2 in the current paper. For both bargaining treatments, we find the investment level to be (virtually) constant in the value of the

Table A1
Mean Investment Levels

Periods	No-trade payoff	OO-game	TP-game
All (1–18)	$r = 1,800$	38.7 [25]	*29.9 [25]
	$r = 6,800$	37.9 [36]	*32.9 [25]
	$r = 7,800$	40.0 [*50*]	32.5 [25]

Remark: theoretical predictions for individual mean investment levels are in brackets, with the socially efficient level in italics. Superscript * indicates a significant difference (investor level) according to a Wilcoxon sign-rank test at the 5% level.

no-trade pay-off. We also find that, for each level of the no-trade pay-off, the investment level is higher under the OO-game than under the TP-game, although these differences are not significant when tested at the investor's level. In all but one case the mean investment exceeds the level predicted by subgame perfection. The single exception is the OO-treatment with $r = 7{,}800$ for which it is predicted that the investment level equals the socially efficient level.

It is interesting to compare the results of Table A1 with the results of Table 2 in the main text. In the OO-game with $r = 1,800$, and in the TP-game with $r = 1{,}800$ or $r = 6,800$, there are no significant difference between the mean investment levels in both tables. For the other treatments, the differences are significant at the 5% level. Overall, in the OO-game, the mean investment level is higher when the non-investor has a positive no-trade pay-off than when the investor has a positive no-trade pay-off; and, in the TP-game, the mean investment level is lower when the non-investor has a high no-trade pay-off level than when the investor has a high no-trade pay-off level. In terms of the ownership of assets, this means that a transfer of enough assets from the investor to the non-investor has a depressing effect on the investment level when the bargaining situation parallels a TP set-up while it has a stimulating effect on the investment level when the bargaining stage resembles an OO-game. This gives further support to the theory of de Meza and Lockwood (1998).

Appendix B: Detailed Comparative Statics of Investment Behaviour

Table B1 presents the detailed period-to-period comparative statics results of investment choices. The first column gives the two adjacent periods that are considered. The second column gives the change in the no-trade pay-off from period j to period $j + 1$, where L, M and H stand for low, intermediate and high, respectively. The following columns present the numbers of investors that decreased (−), kept constant (0) or increased (+) their investment level, for the OO-game and the TP-game separately. In

Table B1

Detailed Period-to-period Comparative Statics of Investment Behaviour

		OO-game			TP-game				
		−	0	+	−	0	+	Fisher's exact	Rank-sum
1–2	MH	8	*6*	6	3	*5*	12	0.134	0.027**
2–3	HL	9	6	*5*	12	*4*	4	0.702	0.225
3–4	LL	2	*9*	9	7	*7*	6	0.219	0.228
4–5	LH	*7*	8	5	3	*2*	15	0.007***	0.001***
5–6	HM	7	*10*	3	10	*7*	3	0.700	0.265
6–7	MM	2	*10*	8	3	*11*	6	0.824	0.792
7–8	MH	4	*11*	5	2	*9*	9	0.360	0.083
8–9	HL	5	5	*10*	13	*6*	1	0.003***	0.001***
9–10	LL	4	*13*	3	6	*10*	4	0.683	0.892
10–11	LM	*9*	6	5	2	*5*	13	0.018**	0.002***
11–12	MM	1	*14*	5	0	*17*	3	0.451	0.671
12–13	MH	3	*15*	2	6	*10*	4	0.341	0.686
13–14	HH	3	*15*	2	3	*10*	7	0.124	0.250
14–15	HL	4	10	*6*	15	*3*	2	0.002***	0.001***
15–16	LM	*8*	9	3	2	*4*	14	0.002***	0.001***
16–17	ML	3	8	*9*	12	*4*	4	0.018**	0.006***
17–18	LH	*10*	7	3	2	*5*	13	0.004***	0.001***

Note: **/*** indicates significance at the 5%/1% level. Theoretical predictions are in italics.

these columns, numbers in italics refer to theoretical predictions. The last two columns present the p-values of two statistical tests, comparing the OO-game and the TP-game. The first, the Fisher exact test, only considers the signs of the changes in investment levels. This test is completely based on the numbers appearing in the previous columns. The rank-sum test statistic in the last column also takes the (not reported) magnitudes of the observed changes into account.

In 5 of the 17 period-to-period transitions, the no-trade pay-off stays the same. So, in both bargaining treatments, we should not observe changes in individual investment behaviour. Observed comparative statics therefore should also not differ between the two bargaining games. Indeed, for all these five transitions, no significant differences are observed between the period-to-period comparative statics of the OO-game and those of the TP-game. In addition, investment behaviour should also not change for the MH and HM transitions. So, for four additional period to period transitions no changes are predicted. (In fact, in one of these transitions some significant differences are found when the magnitude of differences is taken into account (rank-sum test). This, however, only concerns the first transition from period 1 to period 2.) Therefore, only eight relevant transitions remain. We observe significant differences in 7 out of the 8 relevant transitions. These significant differences are all in the direction that the theory predicts, providing additional support for de Meza and Lockwood (1998). The single exception where no significant difference is found (although predicted) concerns the HL transition from period 2 to period 3.

References

Ashenfelter, O., Currie, J., Farber, H. S. and Spiegel, M. (1992). 'An experimental comparison of dispute rates in alternative arbitration systems', *Econometrica*, vol. 60, pp. 1407–33.

Binmore, K., Shaked, A. and Sutton, J. (1989). 'An outside option experiment', *Quarterly Journal of Economics*, vol. 104, pp. 753–70.

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

Binmore, K., Proulx, C., Samuelson, L. and Swierbinski, J. (1998). 'Hard bargains and lost opportunities', ECONOMIC JOURNAL, vol. 108, pp. 1279–98.

Bolton, G. (1991). 'A comparative model of bargaining: Theory and evidence', *American Economic Review*, vol. 81, pp. 1096–136.

Bolton, P. and Whinston, M. D. (1993). 'Incomplete contracts, vertical integration, and supply assurance', *Review of Economic Studies*, vol. 60, pp. 121–48.

Bolton, P. and Xu, C. (1999). 'Ownership and managerial competition: employee, customer, or outside ownership', mimeo.

Chiu, Y. S. (1998). 'Noncooperative bargaining, hostages, and optimal asset ownership', *American Economic Review*, vol. 88, pp. 882–901.

de Meza, D. and Lockwood, B. (1998). 'Does asset ownership always motivate managers? Outside options and the property rights theory of the firm', *Quarterly Journal of Economics*, vol. 113, pp. 361–86.

Ellingsen, T., and Johannesson, M. (2000). 'Is there a hold-up problem?', Working Paper No. 357, Stockholm School of Economics.

Ellingsen, T. and Robles, J. (2000). 'Does evolution solve the hold-up problem?', mimeo, University of Stockholm.

Fehr, E. and Gächter, S. (2000). 'Fairness and retaliation: The economics of reciprocity', *Journal of Economic Perspectives*, vol. 14, pp. 159–81.

Fehr, E., Gächter, S. and Kirchsteiger, G. (1997). 'Reciprocity as a contract enforcement device: Experimental evidence', *Econometrica*, vol. 65, pp. 833–60.

Gantner, A., Güth, W. and Königstein, M. (1997). 'Equity anchoring in simple bargaining games with production', Working paper.

Grossman, S. J. and Hart, O. D. (1986). 'The costs and benefits of ownership: A theory of vertical and lateral integration', *Journal of Political Economy*, vol. 94, pp. 691–719.

Hackett, S. C. (1993). 'Incomplete contracting: A laboratory experimental analysis', *Economic Inquiry*, vol. 31, pp. 274–97.

Hackett, S. C. (1994). 'Is relational exchange possible in the absence of reputations and repeated contract?', *Journal of Law, Economics and Organization*, vol. 10, pp. 360–89.
Hart, O. (1995). *Firms, Contracts, and Financial Structure*. Oxford: Oxford University Press.
Hart, O. and Moore, J. (1990). 'Property rights and the nature of the firm', *Journal of Political Economy*, vol. 98, pp. 1119–58.
Knez, M. J. and Camerer, C. F. (1995). 'Outside options and social comparison in three-player ultimatum games experiments', *Games and Economic Behaviour*, vol. 10, pp. 65–94.
Königstein, M. (1997). 'Convergence to equitable play in the repeated ultimatum game with advance production', Working paper.
MacLeod, W. B. and Malcomson J. M. (1993). 'Investment, holdup and the form of market contracts', *American Economic Review*, vol. 83, pp. 811–37.
Ochs, J. and Roth, A. E. (1989). 'An experimental study of sequential bargaining', *American Economic Review*, vol. 79, pp. 355–84.
Oosterbeek, H., Sonnemans, J. and Van Velzen, S. (1999). 'Bargaining with endogenous pie size and disagreement points; a holdup experiment', mimeo, University of Amsterdam.
Prasnikar, V. and Roth, A. E. (1992). 'Considerations of fairness and strategy: Experimental data from sequential games', *Quarterly Journal of Economics*, vol. 107, pp. 865–88.
Roth, A. E. (1995*a*). 'Introduction to experimental economics'. In (A. Kagel and A. E. Roth, eds): *Handbook of Experimental Economics*. Princeton NJ: Princeton University Press.
Roth, A. E. (1995*b*). 'Bargaining experiments'. In (A. Kagel and A. E. Roth, eds): *Handbook of Experimental Economics*. Princeton NJ: Princeton University Press.
Roth, A. E., Prasnikar, V., Okuno-Fujiware, M. and Zamir, S. (1991). 'Bargaining and market behaviour in Jerusalem, Ljubljana, Pittsburgh and Tokyo: An experimental study', *American Economic Review*, vol. 81, pp. 1068–95.
Siegel, S. and Castellan, N. J. (1988). *Nonparametric Statistics for the Social Sciences*, 2nd edn. New York: McGraw-Hill.
Sloof, R. (2000). 'Finite horizon bargaining with outside options and threat points', Working paper.
Sloof, R., Sonnemans J. and Oosterbeek, H. (2000). 'Specific investments, holdup, and the outside option principle: An experimental study', *Scholar Working Paper Series* WP 08/00.
Tröger, T. (2000). 'Why sunk costs matter for bargaining outcomes: an evolutionary approach', mimeo, University College London.
Winkler, R. L. and Hays, W. L. (1975). *Statistics: Probability, Inference and Decision*, 2nd edn. New York: Holt, Rinehart and Winston.