

FAIRNESS AND THE OPTIMAL ALLOCATION OF OWNERSHIP RIGHTS*

Ernst Fehr, Susanne Krehelmer and Klaus M. Schmidt

We report on several experiments on the optimal allocation of ownership rights. The experiments confirm the property rights approach by showing that the ownership structure affects relationship-specific investments and that subjects attain the most efficient ownership allocation despite starting from different initial conditions. However, in contrast to the property rights approach, the most efficient ownership structure is joint ownership. These results cannot be explained by the self-interest model nor by models that assume that all people behave fairly but they are largely consistent with approaches that focus on the interaction between selfish and fair players.

Small firms in service industries such as law, consulting, accounting, advertising and medicine are often organised as partnerships. In these firms relationship specific investments seem to be important and particularly difficult to contract upon. We often observe that these firms are governed by a very simple governance structure: the partners jointly own the assets of the firm and share profits. This observation is inconsistent with the modern property rights approach pioneered by Grossman and Hart (1986) and Hart and Moore (1990) who argue that joint ownership is rarely optimal:¹ if two parties jointly own an asset they can prevent each other from using it, which minimises their investment incentives. They argue that it would be better to give all the ownership rights to one party, which improves this party's incentives without reducing the incentives of the other party. However, this argument ignores the possibility that a partnership may promote cooperation and investment incentives by appealing to fairness and reciprocity. The purpose of our article is to re-examine the optimal allocation of ownership rights by combining the property rights approach with recent insights into the nature of social preferences and fairness driven behaviour.

The first part of the article reports on several experiments on the optimal allocation of ownership rights. There are two players who can generate a joint surplus with an asset (the 'firm'). At the initial stage they can decide either to have joint ownership of

* Alexander Klein provided excellent research assistance for the experiments discussed in this article. We thank anonymous referees and Leonardo Felli (the editor) for very useful comments and suggestions. We also thank Oliver Hart, Bentley MacLeod and seminar participants at Bocconi University, Boston University, the University of Virginia, the NBER's workshop on organisation theory, the CESifo area conference on micro-economics and the conference on Economics and Psychology in Toulouse for their comments. Part of this research was conducted while the third author visited Yale University and he thanks the Economics Department for its great hospitality. Financial support by Deutsche Forschungsgemeinschaft through SFB-TR 15 and grant SCHM-1196/4-1 and by the EU-Marie Curie RTN ENABLE (MRTM-CT-2003-505223) is gratefully acknowledged. Ernst Fehr also gratefully acknowledges support from the Swiss National Science Foundation (project number 1214-05100.97) and the Network on the Evolution of Preferences and Social Norms of the MacArthur Foundation and the EU-TMR Research Network ENDEAR (FMRX-CTP98-0238).

¹ The earlier literature on property rights comes to somewhat different conclusions. Coase (1960) emphasises that the clear definition of property rights is of crucial importance for economic efficiency. However, the 'Coase Theorem' implies that economic efficiency does not depend on whom ownership rights are allocated to. In the absence of any transaction costs, any (well defined) allocation of ownership rights implements an efficient outcome. Williamson (1985) points out that if the parties can write complete contingent contracts, an appropriate set of incentive contracts can mimic any ownership structure, making the allocation of ownership rights irrelevant.

the firm, or to have one of the parties as the sole owner who hires the other party as an employee. The two parties can then make relationship-specific investments that increase the joint surplus to be generated. Finally the surplus is shared according to the *ex ante* chosen allocation of ownership rights. We are interested in two main questions: first, which ownership structure is (second-best) efficient, in the sense that it induces the most efficient investment decisions by the two parties? Second, do the experimental subjects set up the (second-best) efficient ownership structure *ex ante*, no matter what initial ownership structure they start from?

Our experiments confirm the property rights approach by showing that the ownership structure affects relationship-specific investments and that the large majority of the subjects achieve the most efficient ownership allocation despite starting from different initial conditions. However, in our experiments the most efficient ownership structure turns out to be joint ownership, contrasting with the predictions of the property rights approach.

We offer a theoretical interpretation of the experimental results in the second part of the article. The property rights approach is based on the self-interest model that assumes that all parties are only interested in their own material payoffs. However, both a large number of experiments and systematic field evidence demonstrate that concerns for fairness and reciprocity play an important role in motivating the behaviour of many people.² Several experiments, e.g., Fehr *et al.* (1997) and Fehr *et al.* (2007), indicate that if contracts are incomplete, fairness may act as an enforcement device that complements (and sometimes substitutes for) explicit incentives that are enforced by the courts.

In Section 4, we compare the predictions of the self-interest model to those of two other approaches. The first approach assumes that *all* parties are fair-minded. We show that the allocation of ownership rights does not matter in this case because fairness suffices as an enforcement device for inducing both parties to invest efficiently regardless of who owns the firm. However, this prediction is not confirmed by the experimental evidence. First, fairness concerns apparently did not suffice as an enforcement device to achieve the first-best because neither under joint ownership nor under individual ownership did both parties invest efficiently. Second, this approach fails to explain why relatively high investment levels are induced under joint ownership whereas under individual ownership investments are far less efficient.

The second approach acknowledges that people differ. Some people seem to care quite strongly about fairness, while other people seem to be mainly self-interested. Furthermore, people often do not know whether they are interacting with a fair-minded or a self-interested opponent. Using the Fehr and Schmidt (1999) model of inequity aversion, we show that, as in Hart (1995), no ownership structure implements first-best investments and that the allocation of ownership rights does matter. However, in contrast to Hart (1995), joint ownership is predicted to be the most efficient ownership structure. Thus, the second approach captures the major experimental regularities regarding investment behaviour. Moreover, this approach also predicts that the

² See Camerer (2003) and Fehr and Schmidt (2003) for recent surveys on this literature.

second-best efficient ownership allocation will be implemented regardless of the initial ownership structure, which is what we observed in the experiments.

The question which allocation of ownership rights is optimal is a question about the incentive properties of different governance structures. It is difficult to address this question empirically by using field data. As has been emphasised by Chiappori and Salani (2003), problems of unobserved heterogeneity and endogenous selection often complicate clean inferences about the incentive effects of contracts and governance structures in field data. In our context these problems may result in an ambiguous interpretation of correlations between different ownership structures and different behaviours. Does a particular ownership structure induce a particular type of behaviour or are behavioural differences across ownership rights the result of self-selection of heterogeneous individuals to different allocations of ownership rights? This problem is – in our view – particularly severe in the context of fairness preferences because there is little hope that non-experimental field data allow the control for such preferences. However, our experiments can address this problem because we can observe the behaviour of the same individuals under different ownership structures.

Our article is related to several experimental papers on the hold-up problem. Hackett (1994) was the first to investigate the impact of relationship-specific investments on *ex post* bargaining outcomes experimentally. Similar studies include Gantner *et al.* (2001), Königstein and Tietz (2000) and Oosterbeek *et al.* (2003). All of these papers show that concerns for fairness mitigate the hold-up problem. Ellingsen and Johannesson (2004*a,b*) look at the role of communication as well as that of threats and promises for inducing efficient investments in a hold-up problem with one-sided investments. Ellingsen and Johannesson (2004*a,b*) apply the Fehr and Schmidt (1999) model of inequity aversion to their experiments and show that this model fits the data much better than does the self-interest model. However, none of these papers considers different allocations of ownership rights and how they affect investment incentives. Moreover, these papers do not examine the *endogenous* determination of property rights, which is one of the key questions in our article. Our article is also related to Fehr and Schmidt (2004) and Fehr *et al.* (2007) which also focus on the importance of the interaction between selfish and fair types. However, these papers report on principal–agent experiments in which the principal can choose whether to use an explicit (but incomplete) incentive contract that is enforced by the courts or an implicit contract that uses concerns for fairness as an enforcement device. The current article, in contrast, has a very different focus and examines the conditions under which individual or joint property rights are efficient and whether the parties involved are able to implement the efficient solution regardless of the initial status quo.

There are a few theoretical models showing that joint ownership may sometimes be optimal. Levin and Tadelis (2005) argue (as we do) that the defining feature of a partnership is the (linear) redistribution of profits among partners. Their model complements ours by focusing on the effect of profit sharing on the selection of employees or partners rather than on the incentive effects. Levin and Tadelis consider markets for services where it is difficult to assess service quality. In these markets firms have an incentive to suboptimally hire low ability workers. A profit sharing

partnership mitigates this problem. As was first observed by Ward (1958), the members of a partnership will not admit an additional partner if his contribution to revenues is smaller than the average revenue per partner, even if his marginal contribution is positive. Thus, a partnership is more reluctant to admit additional partners, which is beneficial if it counteracts the incentive to hire too many (low quality) employees.

Rajan and Zingales (1998) consider how the interaction of access to critical resources and ownership rights on assets affects investment incentives. In contrast to Hart and Moore (1990) they show that the ownership of an asset may reduce investment incentives. In their set-up collective ownership may be beneficial because it may distribute the adverse effects of ownership more efficiently.

Halonen (2002) considers a repeated relationship. In the one-shot game, joint ownership is the worst possible ownership structure, as it minimises investment incentives. However, in the infinitely repeated game this ownership structure is desirable because it provides the strongest punishment possibilities if one of the parties deviates. Bar-Isaac (2007) also considers a repeated relationship where two players may want to form a partnership in order to pool their reputations. An individual with an established reputation cannot credibly commit to exerting effort when working alone. However, by hiring and working with a junior partner, he has an incentive to invest in the reputation of his partner that affects the value of the partnership.

None of these papers considers the effects of fairness which we show to be empirically important for the optimal ownership allocation. Furthermore, these papers implicitly assume that the most efficient ownership structure will always prevail. In contrast, we explicitly allow the trading of property rights. Thus, our experiments examine empirically whether efficient property rights will indeed be implemented.

The rest of the article is organised as follows. Section 1 sets up a simple problem of the allocation of ownership rights. Experimental procedures are discussed in Section 2. Section 3 presents the experimental results. We offer a theoretical analysis of the experiment under different assumption about preferences and concerns for fairness in Section 4. Section 5 concludes the article.

1. A Simple Model of the Allocation of Ownership Rights

Consider two players, called A and B , who can generate a joint surplus if they have access to a physical asset (called the ‘firm’). The gross surplus $v(a, b)$ depends on the investments $a \geq \underline{a}$ and $b \geq \underline{b}$ the two players undertake sequentially. Investments are personally costly with investment costs given by $c_A(a)$ and $c_B(b)$, respectively. For simplicity, let us assume that the problem is symmetric in the sense that $c_A(\cdot) = c_B(\cdot) = c(\cdot)$ and $v(a, b) = v(b, a)$. Suppose that B chooses his investment level first, and that A observes B 's investment before she has to invest herself. Let the first-best investment levels be denoted by a^* and b^* ,

$$(a^*, b^*) = \arg \max S(a, b) = \arg \max v(a, b) - c(a) - c(b), \quad (1)$$

and further suppose that they are uniquely defined and satisfy $a^* > \underline{a}$ and $b^* > \underline{b}$, respectively.

The investments are assumed to be unobservable to outsiders, so that any investments above the minimum investment levels \underline{a} and \underline{b} cannot be contracted upon. However, at some initial stage 0, the two parties can contractually determine the allocation of ownership rights on the firm and thus establish the control of the physical assets that are required for production. If one of the parties, say A , is the sole owner of the firm, she then has to hire B at a fixed wage w as an employee. In this case monetary payoffs are given by

$$\begin{aligned} M^A &= v(a, b) - w - c(a) \\ M^B &= w - c(b). \end{aligned} \tag{2}$$

Clearly, a self-interested B player will choose $b = \underline{b}$ because the marginal return on his investment is zero. The A player, however, is a full residual claimant on the margin and will invest efficiently (given B 's investment). The case of B -ownership is symmetric.

If both parties own the firm jointly, then they share the gross returns of the project equally and payoffs are

$$\begin{aligned} M^A &= 0.5v(a, b) - c(a) \\ M^B &= 0.5v(a, b) - c(b). \end{aligned} \tag{3}$$

In this case, each player gets half of the marginal return of his investment, meaning that self-interested players will also not invest efficiently under joint ownership.

This model is reminiscent of Hart (1995) but there are a few notable differences. First of all, in Hart (1995) parties invest simultaneously while we look at the case of sequential investments. From the point of view of the self-interest model this does not make much of a difference.³ However, it helps the parties to coordinate their expectations in the experiment. With simultaneous investments it is impossible to distinguish whether the decision not to contribute to the joint project is driven by selfishness or by the (possibly mistaken) belief that the other player is not going to contribute, too. Because our experiments focus on the effects of fairness, we want to exclude this possible confusion. Furthermore, in the real world parties are often able to condition their own contributions on the contributions of their partners.

Secondly, Hart (1995) considers a three-stage game: at date 0 parties negotiate the ownership structure, at date 1 investments are made and, at date 2, parties bargain on how to split the surplus. If this structure was implemented in the experiment, players would have to bargain twice. They would have to form beliefs at stage 0 on how the surplus will be split at date 2 which in turn affects their behaviour in the negotiations on the allocation of ownership rights at date 0 and their investments at date 1. The problem is that it is difficult, if not impossible, to control for these beliefs in a convincing way. Therefore, we decided to replace the bargaining at date 2 by the fixed

³ If both players invest simultaneously and if investments are complements at the margin, it is easy to show that both players will underinvest; see, e.g., Hart and Moore (1990, Proposition 1). With sequential investments, however, B takes the fact that his actual investment level may affect the optimal investment level of A into account. Thus, in general we cannot rule out B overinvesting; see Noldeke and Schmidt (1998, Proposition 1). However, with the parameterisation of our experiments there is underinvestment in equilibrium under any ownership structure.

payoff functions (2) and (3). These payoff functions capture the idea that the owner of the asset can appropriate a larger share of the surplus than a non-owner.⁴

Parties negotiate the allocation of ownership rights before investments are made. The property rights theory (based on the self-interest model) claims that the parties will always agree on the most efficient ownership structure (i.e. on A or B -ownership) regardless of the initial allocation of ownership rights. We consider two different treatments in order to test this hypothesis:

- *Joint Ownership Design (JOD and JOD-C)*: The two parties start out with joint ownership. At stage 0, A can either choose to retain joint ownership or she can offer to sell her share of the firm to B at price t .
- *A-Ownership Design (AOD)*: Player A is the single owner of the firm when the game starts. At stage 0, A can choose to remain the sole owner of the firm and to hire B as an ‘employee’ at a fixed wage w . Alternatively, A can choose to make B a co-owner by giving him half of the firm.

Our main questions are, first, whether A -ownership is indeed more efficient than joint ownership and, second, whether parties manage to achieve the most efficient ownership structure independent of the initial ownership structure.

2. Experimental Design and Procedures

We used the following parameterisation of this investment problem in our experiments. The two parties choose $a, b \in \{1, \dots, 10\}$. Investments (a, b) yield a gross surplus $v(a, b) = 22(a + b)$, while investment costs are $c_A(a) = 12a$ and $c_B(b) = 12b$. Thus, investments are neither complements nor substitutes at the margin, implying that optimal investment levels are independent of each other.

Given these parameters of the experiment, the efficient investment levels are given by $a^* = b^* = 10$, yielding a joint surplus of 200. However, the self-interest model predicts that no ownership structure implements efficient investments. With joint ownership, each party receives only half of the gross surplus, leaving the private marginal return on investment under its marginal cost and inducing both parties to choose minimum investment levels, $a = b = 1$. If one party is the sole owner of the firm, then this party receives the full gross surplus on the margin and has an incentive to invest efficiently. The other party, however, receives a fixed wage and will therefore choose the minimum investment level of 1. Thus, both possible ownership structures are inefficient, but A or B ownership is more efficient than joint ownership.

⁴ These payoff functions are identical to the endogenously derived payoff functions in Hart (1995) in the case of investments in a physical asset. If the investment is in a physical asset and A is the sole owner of this asset, she has full access to the returns of B 's investment b , even without B 's consent. Thus, she gets $v(a, b)$ on the margin, while B gets 0. If both parties own the asset, each of them can block the other from using the asset, so they share the surplus equally. Of course, investments may also be in human capital. In this case A would not have access to B 's investment (which is embodied in B rather than in the asset) even if she is the sole owner of the asset. How A and B split the surplus now depends on how the investments affect their outside options. This is slightly more difficult to model explicitly. See Hart (1995, p. 68) and Noldeke and Schmidt (1998) for a more detailed discussion of investments in human and in physical capital. Note that joint ownership is never optimal with investments in human assets while it may be optimal with investments in physical assets. However, we have chosen the parameters of the experiment such that A (or B) ownership always dominates joint ownership if parties are self-interested (see Section 3 below).

The time structure of the experiments is as follows. There is an ownership allocation game at stage 0, followed by an investment game at stage 1. The ownership allocation game differs in the different treatments, while the structure of the investment game is always the same. We consider two main treatments and one control treatment.

- *Joint Ownership Design (JOD)*: Initially, both parties own the firm jointly. *A* can choose
 - either to retain joint ownership, in which case each party earns 50% of the joint surplus
 - or to offer to sell her share of the firm to *B* at price t . If *B* accepts this offer, he earns 100% of the gross surplus $v(a,b)$, while *A* receives the fixed payment t . If *B* rejects, the game ends and both parties receive a payoff of zero.
- *A-Ownership Design (AOD)*: *A* is the sole owner of the firm when the game starts. *A* has two options.
 - To remain the sole owner of the firm, in which case she gets 100% of the gross surplus $v(a,b)$, and to hire *B* as an ‘employee’ at a fixed wage w . If *B* accepts the wage offer the game moves to the investment stage; if *B* rejects both players receive a zero payoff.
 - To make *B* a co-owner by giving him half of the firm. In this case, each party gets 50% of the gross surplus if player *B* accepts the offered share; if *B* rejects, both players receive zero payoff.⁵
- *Control Treatment (JOD-C)*: This treatment is identical to the Joint Ownership Design except for one detail. If *B* rejects *A*’s offer, then the game does not end but continues under joint ownership. This design is more natural than treatment JOD because the two parties would want to continue their relationship on the basis of the status quo after an offer has been rejected. The control treatment enhances the chances for sole ownership because player *A* need not worry that the surplus is lost if her offer gets rejected. Therefore, we expect player *A* to offer her ownership share for sale more frequently. We start out with JOD rather than with JOD-C because in the former treatment expectations on what is going to happen if negotiations break down cannot affect behaviour. Therefore, treatment JOD allows for a ‘cleaner’ experimental analysis. By comparing JOD to JOD-C we can then see what the effect of these expectations is.

After the allocation of ownership rights has been determined at stage 0, the investment game is played at stage 1:

- First player *B* chooses his investment level b .
- Then player *A* is informed of b and subsequently chooses her investment level a .

Finally, payoffs are made according to the investment levels and the allocation of ownership rights.

The experiments were conducted at the University of Munich with undergraduate students in law, political science, engineering, etc. We conducted a total of eight experimental sessions. Four sessions (S1–S4) implemented the Joint Ownership Design

⁵ We did not use the expressions ‘employee’ and ‘partner’ in the actual experiments but the neutral terms ‘actor *A*’ and ‘actor *B*’ instead.

(JOD), two sessions (S5–S6) implemented the *A*-Ownership Design, and two additional sessions (S7 and S8) implemented the control treatment (JOD-C). We had 20–24 subjects in each session, half of them in the role of player *A*, the other half in the role of player *B*. The two groups were located in separate but adjacent rooms. All subjects had to read the instructions and to solve several exercises before the experiment started, to make sure that they all understood the rules of the experiment. We had ten rounds in each session; in each round the *A* and *B* players were matched with new partners. Thus, we have 10–12 contracts with 10 different anonymous contracting partners for each subject in each experimental session.

The subjects computed their own payoffs and that of their opponents after each round. In order to rule out the possibility of reputation building, the outcome of each round was strictly confidential, that is, each pair of players observed only what happened in their own relationship. They neither observed the contracts chosen by or offered to the other subjects in the room nor their current partner's past behaviour. Furthermore, the matching was such that each player *A* (*B*) was matched to a different anonymous player *B* (*A*) in each period. Finally, the subjects collected their total monetary payoffs privately and anonymously at the end of the session. Each session lasted for about one and a half hours. A complete set of the instructions for all our experiments can be found on our webpage.⁶

All participants received an initial endowment of €10.00 in each session (\approx US \$12.50 at the time of the experiment). The experimental (token) payoffs were exchanged into money at the rate of 1 token = €0.03. Thus, *A* and *B* could jointly earn a maximum surplus of €6 in each of the ten rounds. On average, the subjects earned €28.42 (\approx US\$35.25), an hourly wage of about € 18.95 (\approx US\$23.70).

3. Experimental Results

3.1. Joint Ownership Design

In the Joint Ownership Design (JOD) each party initially owned 50% of the firm. At stage 0, player *A* could either choose to retain joint ownership or to sell her ownership stake to player *B*. We conducted four sessions (S1–S4) of this design with a total of 470 observations. The major ownership patterns which emerged in this treatment can be summarised as follows:

RESULT 1. *In the large majority of cases, A players stick to joint ownership. If A players offer to sell their ownership stake to B, B rejects these offers in about 30% of all cases.*

Support for Result 1 is provided by Figure 1 and the following statistics: *A* players prefer joint ownership in 300 out of 470 (63.8%) cases. In the remaining 170 cases (36.2%), *A* players attempt to sell their ownership stake to *B* players. The latter reject these offers in almost one-third of the cases (52 times). Therefore, *B*-ownership was established in only 118 cases (25.1% of all observations). Figure 1 shows the evolution of the relative share of cases in which joint ownership prevailed and in which *B*-ownership is either

⁶ The full set of all our experimental instructions, in the original German and translated into English, are available at http://www.et.wl.uni-muenchen.de/forschung/experiments/property_rights/index.html.




Fig. 1. *Share of Ownership Choices (JOD)*

proposed and accepted or proposed and rejected. Recall that if A 's offer to sell her share is rejected, the game ends and both parties earn a payoff of zero. The share of joint ownership varies roughly between 60 and 70% of the cases, whereas accepted B -ownership varies between 20 and 30% of the cases. There is no time trend in the data.⁷

Result 1 clearly contradicts the self-interest model, which predicts the choice of B -ownership because it supposedly induces more efficient investment behaviour. In order to understand this result, we have to analyse the actual investment behaviour of the two players, depending on whether joint ownership or B -ownership prevailed. We first consider those cases where player A preferred joint ownership:

RESULT 2. *If A decides to stick to joint ownership in the JOD, A 's investment is strongly increasing in B 's investment level. There is, however, significant heterogeneity in A 's investment response.*

Even though it is a dominant strategy for a self-interested player A to choose $a = 1$ under joint ownership, many A players reciprocate high investment levels of player B by choosing a high investment level themselves. Figure 2 illustrates this, depicting how the average investment level of player A responds to B 's investment level. The Figure shows that A 's average investment is increasing in B 's investment. This increase is particularly strong if B invests more than $b = 6$. On average player A chooses $a = 6.7$.

Statistical analyses (see Table 1) also support the reciprocal pattern of A 's investment response. Regression (1) is a simple OLS regression of A 's investment on B 's investment.⁸ The standard errors in this regression ignore the possibility that the A players' investment levels may not be independent of each other because each A player invested several times. Therefore, we treat the observations of each individual A player as a

⁷ A random-effects GLS regression shows that the probability of choosing joint ownership goes down by 1% per period, but this effect is not significant at the 5% level ($p = 0.068$).

⁸ We also conducted Tobit regressions to check the robustness of our results. All variables that are significant in the OLS regressions are also significant in the Tobit regression and they exhibit the same sign.




Fig. 2. A's Average Investment Given b Under Joint Ownership (JOD)

Table 1
A's Investment as a Function of b Under Joint Ownership (JOD)

Dependent variable: A's investment a	(1)	(2)	(3)	(4)
	(robust standard errors)	(robust standard errors & clusters)	(robust standard errors)	(robust standard errors & clusters)
	Based on data from JOD		Based on data from AOD	
Constant	-0.11 (0.38)	-0.11 (0.83)	-0.70 (0.51)	-0.70 (0.54)
B's investment b	0.89*** (0.04)	0.89*** (0.09)	0.81*** (0.07)	0.81*** (0.10)
No. of observations	300	300	187	187
Adjusted R ²	0.57	0.57	0.21	0.21

Table 1 reports the coefficients of OLS regressions. Robust standard errors are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

separate cluster in regression (2). The standard errors in this case are based on the assumption that the investment levels are independent across different A players, but allow for dependent observations within each cluster (i.e. for all observations belonging to one A player). The assumption that investments are independent across A players is reasonable because an A player never had the opportunity of observing what other A players did. Both regression (1) and (2) show that the impact of b on a is sizeable and highly significant. In fact, these regressions indicate that an increase in B's investment level by one unit increases A's expected investment level by 0.89 units. Therefore, the expected marginal return of one additional unit of investment for player B is

$$\frac{1}{2}22(1 + 0.89) - 12 = 8.79 > 0. \tag{4}$$

Thus, if player B maximises his expected monetary income, the choice of b = 10 is optimal.

Table 2
Investments (b,a) under Joint Ownership (JOD)

$b \backslash a$	1	2	3	4	5	6	7	8	9	10	Σ
1	37					1				3	41
2	5	3								1	9
3	2	1	1								4
4	3	1		1				1			6
5	10	1		2	4						17
6	7	1				3	1				12
7	3				1	3	5			1	13
8	4					1	1	1		1	8
9	3						1		6		10
10	11	1			3	3	1	2	2	157	180
Σ	85	8	1	3	8	11	9	4	8	163	300

Although *A* players respond, on average, quite strongly to increases in the investment level of player *B*, it is important to notice that there are substantial differences in individual behaviour. In many cases the *A* players exactly matched the investment level of *B* but there are also a substantial number of cases where *A* players did not fully reciprocate *B*'s choice. In fact, in 28.3% of the cases the *A* player chose the minimal investment level of $a = 1$ and in 25% of the cases $a < b$ prevailed. Thus, a sizeable fraction of *A* players was mainly driven by self-interest.

Table 2 shows that most *B* players trusted that *A* players will reciprocate. 60% of all *B* players (180 of 300) under joint ownership chose the efficient investment level $b = 10$ at stage 1. Only 41 *B* players (13.7%) chose the minimum investment level $b = 1$ and the average investment of *B* players amounted to 7.7. This result contrasts sharply with the prediction of the self-interest model that forecasts *B* players will invest nothing.

Let us now compare the outcome of joint ownership to the case where *A* tried to sell her ownership share to *B*.

RESULT 3. *In the JOD, joint ownership is the more efficient allocation of ownership rights. Total investment ($a + b$) is higher if joint ownership prevails and both players receive a higher average payoff if *A* decides to retain joint ownership rather than selling her ownership rights to *B*.*

A's average income is 77.5 under joint ownership, while *B* earned an average of 66.4. If *A* offered to sell her ownership stake to *B*, average payoffs are 67.9 and 12.8 respectively. Thus, it turns out that the payoff difference between *A* and *B* is fairly small (but statistically significant) under joint ownership, but very large if *A* offered her ownership stake for sale. Furthermore, joint ownership Pareto-dominates *B*-ownership. Figure 3 shows that both players were better off with joint ownership rather than with *B*-ownership in all periods (except for period 1). Again, these differences are statistically significant.⁹

⁹ A Mann-Whitney test shows that *A*'s payoff under joint ownership is significantly larger than that of *B* under joint ownership ($p = 0.0003$) and is also significantly larger than *A*'s payoff if she tried to sell her ownership stake to *B* ($p = 0.0001$).




Fig. 3. Payoffs under Joint Ownership and B-Ownership (JOD)

To better understand why joint ownership is more successful, consider the 170 out of 470 cases (36.2%) where *A* tried to sell her ownership stake to *B*. Recall that these offers were rejected quite often (in 52 cases, i.e. 30.6% of all offers). The average price of the rejected offers was 176.9, while the average price of the accepted offers was just 120.8.¹⁰ If the seller accepts the offer and invests $b = 10$ himself while *A* invest $a = 1$ (which are the dominant strategies for self-interested players if *B* becomes the sole owner), then *B*'s payoff is $M^B = 22(10 + 1) - 12 \cdot 10 - t = 122 - t$. Thus, accepting a price offer in excess of 122 only pays off if *B* expects *A* to invest considerably more than $a = 1$. However, in fact, *A* invested $a = 1$ in 85 out of the 118 cases (72%) where she sold her ownership stake successfully to *B*, and her average investment level in these cases was just 1.9, while virtually all *B* players (111 out of 118) chose $b = 10$. In comparison, *B* invested $b = 7.7$ and *A* invested $a = 6.7$ on average under joint ownership. Thus, offering *B*-ownership is less efficient than retaining joint ownership for two reasons: first, the sales offer was frequently rejected. Second, even if the sales offer was accepted total investment was lower than under joint ownership.

3.2. A-Ownership Design

We now turn to the results of the *A*-Ownership design where *A* players could choose between sticking to *A*-ownership or giving away half of the revenues of the project to the other player. We observed a total of 230 contractual choices in sessions S5 and S6. The major ownership patterns under initial *A*-ownership are summarised in

RESULT 4. *The overwhelming majority of A players opted to make B the joint owner of the project. There is no significant time trend. If anything, the share of joint ownership increases over time.*

¹⁰ In fact, a Mann-Whitney Test confirms that the differences between the accepted and rejected offers are highly statistically significant ($p = 0.0000$).




Fig. 4. *Share of Ownership Choices (AOD)*

Figure 4 supports Result 4, showing the relative share of joint ownership and A-ownership over time. Joint ownership is implemented in 80–90% of the cases from period 3 onwards.¹¹ Overall, joint ownership was chosen in 81.3% of all cases (187 of 230 observations). Thus, A players clearly preferred joint ownership.

The next result shows that the reason for this preference is similar to that in the JOD.

RESULT 5. *If A makes B a co-owner, more than two-thirds of all B players trust A players and choose $b = 10$. In fact, A's average investment is strongly increasing in B's investment level, such that $b = 10$ is the choice that maximises B's expected earnings. There is, however, significant heterogeneity in A's investment response.*

In 135 out of 187 (72.2% of the) joint ownership cases player B chose $b = 10$ and in less than 10% of the cases player B chose $b < 6$. The average investment level of player B is, therefore, rather high and amounts to $b = 8.9$. Player A reciprocates B's high investment with $a = 6.5$ on average. Regressions (3) and (4) in Table 1 provide evidence for player A's reciprocal investment response. These regressions have the same structure as regressions (1) and (2) but are based on different data. While regressions (1) and (2) are related to investment behaviour under joint ownership where joint ownership is the initial condition, regressions (3) and (4) relate to behaviour under joint ownership when the initial condition is given by A-ownership. Despite these differences in initial conditions, the A players behave very similarly, as both the size and the significance of the coefficient for b indicate. In fact, a rise in b by 1 unit raises A's investment by 0.81 units in regressions (3) and (4). Therefore, $b = 10$ maximises B's expected monetary payoff.

¹¹ A random-effects GLS regression shows that the probability of choosing JO goes up by 1.2% per year but the effect is hardly significant ($p = 0.056$).

Table 3
Investments (b,a) under Joint Ownership (AOD)

$b \backslash a$	1	2	3	4	5	6	7	8	9	10	Σ
1	9										9
2	1										1
3											0
4				1							1
5	3	1		1	2						7
6	2					1				1	4
7	5	1	1				3				10
8	2			1	1	1		4			9
9	3						1	2	5		11
10	33	2			2		1	1	4	92	135
Σ	58	4	1	3	5	2	5	7	9	93	187

The individual investment behaviour of A players, however, displays a substantial amount of heterogeneity (see Table 3). In 117 out of 187 (63%) cases A players exactly matched B 's investment choices. However, in 31% of the joint ownership cases A chose $a = 1$ and in 37% of the cases $a < b$ prevailed.

Results 4 and 5 indicate that A players successfully elicited high investment levels by making B a co-owner. This raises the question whether A could do equally well by keeping her ownership share. Our next result therefore compares the efficiency and income levels under joint ownership and A -ownership when A -ownership is the initial condition.

RESULT 6. *If A is initially the sole owner, joint ownership is the more efficient allocation of ownership rights. Total investment ($a + b$) is higher if A makes B a co-owner and both players receive a higher average payoff compared to the case where A remains the sole owner and hires B as an employee.*

A decided to retain A -ownership and to hire B as an employee in 43 out of 230 cases (18.7%) B players rejected the wage offers 5 times. The average wage offer was 62.1. Almost all B players chose the minimum effort level (on average $b = 1.3$) under A -ownership, while all A players chose $a = 10$. Thus, average total investment $a + b = 11.3$, which is significantly less than the average total investment of 15.4 if A makes B a co-owner. A 's average income under joint ownership is 91.4, while it is only 55.0 if she chose to retain A -ownership. B 's average income under joint ownership is 62.9, while he only received 44.5 if A tried to hire him as an employee.¹²

While these results are again inconsistent with the self-interest model, they confirm the Grossman and Hart (1986) prediction that parties will always try to achieve the most efficient ownership structure. Our results show that independent of whether the parties originate with joint ownership or with A -ownership, they will always end up with joint ownership, which turns out to be the most efficient ownership structure.

¹² Even if we only consider those cases where B accepted A 's wage offer, both parties receive a lower payoff than under joint ownership (A : 62.2, B : 50.4).

3.3. Control Treatment for Joint Ownership Design (JOD-C)

We implemented a control treatment in Sessions S7 and S8 for the Joint Ownership Design in which the initial condition was also given by joint ownership. The game did not end after *B* rejected *A*'s offer in the control treatment, however, but continued with joint ownership. We have a total of 240 observations for this treatment. Again, joint ownership prevailed in the large majority of cases (196 of 240 observations, 81.7%). However, *A* players tried to sell their ownership stakes more frequently in this case (in 134 out of 240 cases, 55.8%). It seems that *A* players considered making an offer to be less risky because the parties reverted to the status quo of joint ownership if the offer was rejected. In fact, the offers were rejected considerably more often (in 90 out of 134 cases, 67.2%) than in the JOD treatment. Let us consider the three different possibilities in turn:

- If *A* did not make an offer, the investment behaviour is very similar to the investment behaviour in the other designs when joint ownership prevailed. *B* players invested $b = 7.0$ on average, while *A*-players invested $a = 6.1$ on average. The *A* players' reciprocal behaviour is virtually identical to that we observed under joint ownership in JOD and AOD.
- If *A* made an offer that the *B* player accepted, the investment behaviour is very similar to the corresponding case in the JOD. Almost all *B* players invested efficiently ($b = 9.5$ on average), while almost all *A* players chose the minimum investment level ($a = 1.2$ on average). However, it is interesting to note that the average price of the accepted offers was much lower than the average accepted price in the JOD ($t = 88.9$ as compared to $t = 120.8$). This reflects the fact that *B*'s threat point payoff if he rejected the offer was not zero but the payoff resulting from the status quo of joint ownership. In fact, *B*'s average payoff after accepting *A*'s offer is 31.4 which is significantly larger than the 18.5 that he received on average when he accepted *A*'s offer in the original JOD.
- Two-thirds of all offers *A* players made were rejected. The average price of the rejected offers was 161.4, again somewhat lower than the average price of rejected offers in the original JOD (176.9) but still very high. After the offer was rejected, the two players played the investment game under joint ownership. *B* invested $b = 6.3$ in these cases, while *A* invested only 3.9 on average, which is significantly less than the investment levels under joint ownership in JOD or AOD.

However, the basic ownership pattern that emerged in the control treatment is the same as in the two previous treatments. Joint ownership is more efficient because it induces higher total investments and yields higher payoffs for both players. As a consequence, joint ownership prevails in the large majority of cases.

4. Theoretical Interpretation

The predictions of the property rights approach and the self-interest model can be summarised as follows:

PROPOSITION 1. *If both parties are only interested in maximising their own material payoff, then*

- 1 *A and B-ownership are equally efficient. Both are more efficient than joint ownership.*
- 2 *Regardless of the initial allocation of ownership rights, the parties will trade ownership rights ex ante so as to set up the efficient ownership structure and to implement the second best optimal investment decisions.*

While the experimental results clearly refute the first prediction, showing that joint ownership is far more efficient than *A* or *B*-ownership in all three treatments, the experiments largely confirm the second prediction. The parties agreed on the more efficient joint ownership arrangement *ex ante* in the large majority of all cases, regardless of whether they started from *A*-ownership or from joint ownership.

In this Section, we want to discuss whether theories of fairness are consistent with the experimental results. Several recent theories capture concerns for fairness and/or reciprocity in individual decision making. Some of the proposed models, in particular Rabin (1993) and Dufwenberg and Kirchsteiger (2004), adopt the concept of ‘psychological game theory’ that Geanakoplos *et al.* (1989) introduced in order to model ‘intention-based reciprocity’. Players in these models not only have beliefs about their opponents’ actions but also about their intentions. They are willing to reward kind and to punish unkind *intentions*. While these models convey many interesting insights, they are complicated and often difficult to use even in very simple applications. Furthermore, they are plagued by multiple equilibria even in very simple games. Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) follow a different approach. They assume that players only care about outcomes (and not about intentions) but that they have ‘social preferences’ in the sense that they dislike inequitable allocations. These models do not capture ‘reciprocity’ in the intention-based sense but rather ‘distributional fairness’ or ‘inequity aversion’. They use standard game theoretic tools and they can be applied straightforwardly to any game. Furthermore, Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) show that their models are consistent with many observations in a variety of experimental games.¹³

In the following we apply the Fehr and Schmidt model to our treatment conditions. We first consider the case where *all* players are strongly concerned about fairness. We then consider the more realistic case of a heterogeneous population where *some* people are strongly concerned about fairness while others hardly care at all. The Fehr and Schmidt (1999) model can deal with both cases and allows for incomplete information about the types of the players. With heterogeneous agents, the interaction between fair and self-interested types comes into play and gives rise to some interesting new insights. We analyse these two cases in more detail in the rest of this Section.

Before we proceed we would like to stress that we do not regard our experiments as evidence for the claim that social preferences take the form of inequity aversion. In

¹³ There are also a few models that try to model preferences for fair outcomes and fair intentions simultaneously, in particular Falk and Fischbacher (2006) and Charness and Rabin (2002). These models are more general, because they combine social preferences and intention-based reciprocity, but they are even less tractable for applications and again plagued by multiple equilibria. See Camerer (2003) and Fehr and Schmidt (2003) for extensive surveys and critical discussions of this literature.

order to show the existence of inequity averse individuals simpler experiments such as those in Charness and Rabin (2002), Falk *et al.* (2003), Engelmann and Strobl (2004), or Fehr *et al.* (2006) are needed. We also know from these experiments that inequity aversion is but one important aspect of social preferences, i.e., inequity aversion obviously does not provide a complete description of social preferences. Thus, the purpose of this Section is not to show the superiority of inequity aversion over other models of social preferences¹⁴ but to provide a simple, yet tractable, analysis that puts the interaction between selfish and fair-minded players at the centre and, thus, enables a deeper understanding of the forces behind joint and individual ownership.

4.1. *Homogeneous Fair Agents*

Consider first the case where all parties strongly care about fairness. We use a special case of the theory of inequity aversion by Fehr and Schmidt (1999) to model concerns for fairness. The theory assumes that some people are not only concerned about their own material payoff but also care about inequity or, in our context, inequality.¹⁵ The utility function of inequity averse (fair) players in the two-player case is given by

$$U_i(x) = x_i - \alpha_i \max\{x_j - x_i, 0\} - \beta_i \max\{x_i - x_j, 0\}, \quad (5)$$

$i \in \{1, 2\}$, $i \neq j$, where $x = (x_1, x_2)$ denotes the vector of monetary payoffs, $\beta_i \leq \alpha_i$ and $0 < \beta_i < 1$. The term in the utility function weighted with α_i measures the utility loss that stems from inequality to i 's disadvantage, while the term weighted with β_i measures the loss from advantageous inequality.

We assume in this subsection that all people are sufficiently concerned about fairness to make them always try to achieve an equal distribution of monetary payoffs, i.e. $0.5 < \beta_i < 1$. This assumption implies that if player i is better off than his opponent, then he prefers to give one dollar to his opponent (which reduces inequality to his advantage by two dollars) rather than to keep this dollar for himself. However, he would not throw one dollar away (which reduces inequality to his advantage by one dollar) in order to reduce inequality. Of course, it is a very bold assumption that *all* people are so strongly concerned about inequality. This is why we will consider the case of heterogeneous agents in the next subsection.

Let us start with the case of joint ownership. If all players are sufficiently inequity averse, then player A will match B 's investment level and choose $a(b) = b$ because this equalises final payoffs. Anticipating this, player B will invest efficiently at stage 1.

What about A -ownership? A 's investment in the last stage of the investment game increases her own payoff but does not affect that of B , because B receives a fixed wage anyway. We show that A will invest efficiently in this situation: we can distinguish two different ranges of possible investment levels for A (one of which may be empty) for

¹⁴ The Bolton and Ockenfels (2000) model makes the same qualitative predictions regarding the pattern of the ownership structure in the context of our treatments and it may well be that other models that take the heterogeneity of selfish types and social preference types seriously, also explain the major facts in our experiment.

¹⁵ Fairness implies that equals should be treated equally. The subjects enter the laboratory as equals in our experiments. They have no information about their opponents and do not know with whom they trade. Thus, in these very simple environments, it seems natural to define *equality* as the reference point for a fair payoff distribution.

any given wage w and b . Consider first the range of investment levels where A 's final payoff would be smaller than B 's final payoff. A wants to reduce inequality to her disadvantage in this range by maximising her own payoff, i.e., A chooses $a = 10$. Now consider the range of investment levels for which A 's final payoff would exceed that of B 's. A further increase in A 's investment in this range would additionally increase the inequality to A 's advantage. However, if $\beta_A < 1$, A prefers to earn one additional dollar for herself rather than to throw this dollar away in order to reduce the inequality towards B . Thus, she has an incentive to increase her investment up to the efficient amount in both cases.

B 's investment b at the first stage of the investment game depends on the wage that A offered him at stage 0. If w is small, B will choose a low investment level in order to reduce the payoff difference between himself and A . On the other hand, if A offered a generous wage at stage 0 so that $w = \frac{1}{2}v(a^*, b^*)$, then B will choose the efficient investment level b^* because he wants to increase A 's payoff in order to reduce the inequality that is now to his advantage. Thus, at stage 0, a fair player A will make this generous wage offer and both parties will choose the efficient investment levels a^* and b^* , respectively. The analysis of B -ownership is analogous to that of A -ownership.

PROPOSITION 2. *The ownership structure is irrelevant if both parties are sufficiently concerned about distributional fairness. Both parties will invest efficiently no matter whether there is joint, A -, or B -ownership.*

Thus, if concerns for fairness are universal and sufficiently strong, they induce both parties to invest efficiently even if investments cannot be contracted upon. Fairness suffices as an enforcement device and the allocation of ownership rights does not play any role. However, as in the self-interest model, this prediction is clearly refuted by the experimental evidence.

4.2. Interaction of Self-interested and Fair Players

We now consider the case where some people are concerned about fairness while others are mainly self-interested. Furthermore, we will assume that people do not know whether they face a fair and trustworthy opponent or whether their opponent is selfish and going to exploit them. We use the Fehr and Schmidt (1999) model of inequity aversion again but now we allow for heterogeneous fairness preferences. We assume that there are 60% self-interested types ($\alpha_i = \beta_i = 0$) and 40% 'fair' types with $\alpha_i = 2$ and $\beta_i = 0.6$. Thus, fair types are willing to share the surplus of a contract equally but also to reject offers that give them less than 40% of the surplus. This distribution of types is a simplification of the distribution used in Fehr and Schmidt (1999).¹⁶ On the

¹⁶ Note that the heterogeneity of social preferences is a well established fact; see e.g., Charness and Rabin (2002); Fehr and Schmidt (2003); Engelmann and Strobl (2004) and Fehr *et al.* (2006). Thus, the question is not whether theory should incorporate heterogeneous social preferences but which distribution of preferences theory should assume. In Fehr and Schmidt (1999) we used four different types: 30% of the population are assumed to have $\alpha_i = \beta_i = 0$, 30% are assumed to have $\alpha_i = 0.5$ and $\beta_i = 0.3$, 30% are assumed to have $\alpha_i = 1$ and $\beta_i = 0.6$, and 10% are assumed to have $\alpha_i = 4$ and $\beta_i = 0.6$. It turns out to be very tedious to solve the model for four different types. This is why we simplified the model and use only two different types here. We used the same simplified distribution in Fehr and Schmidt (2004) and Fehr *et al.* (2007).

basis of these assumptions it is straightforward to solve the property rights game using standard game theoretic tools. The full analysis is somewhat lengthy and therefore relegated to an appendix that can be found on our webpage.¹⁷ In the following, we report the main predictions for our experiments and give the intuition for them.

4.2.1. Analysis of the joint ownership design

The contract offer made at stage 0 may signal some information about A 's type, due to the asymmetric information about whether A is self-interested or fair-minded. Let p denote the (endogenously determined) probability B assigns to the event that he faces the self-interested type of player A . We first analyse the players' behaviours at the investment stage, after which we examine the entire game.

Consider first the case where A chooses to retain joint ownership at stage 0. At the investment stage, a self-interested type of A chooses $a = 1$ while a fair-minded player A chooses $a = b$. Therefore, a self-interested type of B chooses $b = 10$ if he believes that it is sufficiently likely that he faces a fair-minded player A , i.e. if p is sufficiently small ($p < 10/11 = 0.91$).¹⁸ Otherwise he chooses $b = 1$. The fair-minded player B , on the other hand, will be more careful, because he not only suffers the monetary loss if player A does not reciprocate, he also suffers from the inequity that is generated if he invests while A does not. Therefore, the fair-minded player B will invest only if $p < 10/35 = 0.29$.¹⁹ This result is surprising. It says that if there is uncertainty about A 's type, then a self-interested player B is more likely to invest than the fair-minded type of B .

Now consider the case where A tries to sell her ownership stake to B at price t . The self-interested type of player A finds it optimal to make a greedy offer that is going to be rejected by a fair-minded player B , in which case the surplus is lost. However, if the offer is accepted, B becomes the sole owner of the project and full residual claimant on profits. In this case it is a dominant strategy for both types of B to choose $b = 10$. The intuition is simply that B 's investment under B -ownership does not affect A 's payoff, so B cannot increase A 's payoff by investing less than the efficient amount. The self-interested type of A again chooses $a = 1$. The fair-minded type of A will invest in order to reduce the inequality between herself and B but only if she sold her ownership stake at a sufficiently high price to B . In fact, she chooses the efficient investment level if and only if $t \geq 220$.

We now turn to the analysis of the entire game. First, we can rule out the possibility of a separating equilibrium in which the self-interested type of A chooses one type of contract with probability 1 and the fair-minded player A chooses another contract type. The intuition is that the self-interested player A would always want to mimic the fair player: suppose that the selfish type of A sells her ownership stake while the fair type retains joint ownership. Then B would invest 10 under joint ownership which induces the self-interested player A to deviate and to retain joint ownership as well. A similar

¹⁷ Please visit: http://www.et.vwl.uni-muenchen.de/forschung/experiments/property_rights/index.html.

¹⁸ B 's expected monetary payoff if he believes he faces a selfish A -player with probability p is given by $EU^{Bs} = p11(a^s + b) + (1 - p)11(a^f + b) - 12b$. Substituting $a^s = 1$ and $a^f = b$, this reduces to $EU^{Bs} = 11p + (10 - 11p)b$. Thus, if $p < 10/11$, B 's monetary payoff is increasing in b .

¹⁹ The fair-minded player B is also concerned about the inequality to his advantage generated by the selfish type of player A . So he maximises $EU^{Bf} = p[11(a^s + b) - 12b - \alpha(12b - 12a^s)] + (1 - p)[11(a^f + b) - 12b]$. Substituting $a^s = 1$ and $a^f = b$, this reduces to $EU^{Bf} = 11p + 12\alpha p + (10 - 11p - 12\alpha p)b$. Thus, if $p < 10/(12\alpha + 11)$, B 's utility is increasing in b . Substituting $\alpha = 2$ we get $p < 10/35$.

argument holds for the opposite case, where the self interested player A sticks to joint ownership while the fair-minded player offers to sell. B would choose $b = 1$ in this case if he is offered joint ownership, so the self interested player A is better off by selling her ownership share.

It seems very plausible in the game under consideration that retaining joint ownership will not be interpreted as a signal that player A is selfish. This is captured by the following condition:

CONDITION 1. *If A chooses to retain joint ownership, then B 's updated belief that he faces a self-interested type of A does not increase.*

This condition implies that the game has a unique Perfect Bayesian Equilibrium:

PROPOSITION 3 [*Joint Ownership Design*]. *With incomplete information about the players' types there exists a unique Perfect Bayesian Equilibrium outcome satisfying Condition 1. The equilibrium is a pooling equilibrium in which both types of player A retain joint ownership.*

In the investment game the self-interested type of B chooses $b^s = 10$, while the fair-minded type of B chooses $b^f = 1$. The self-interested type of A chooses $a^s = 1$ and the fair-minded type of A chooses $a^f = b$.

Note that Proposition 3 differs sharply from Propositions 1 and 2. Proposition 1 assumed that the fact that all players are self-interested is common knowledge. No ownership structure implements first-best investments in this case, but A and B -ownership are strictly better than joint ownership. Proposition 2 assumed that all players are fair-minded. Any allocation of ownership rights in this case implements first-best investment decisions and the allocation of ownership rights is indeterminate. Proposition 3 shows, like Proposition 1, that first-best investments cannot be implemented with incomplete information about the players' types, but that the second-best allocation of ownership rights is the unique equilibrium outcome. However, in contrast to Proposition 1, joint ownership is optimal in this case. Thus, the ownership prediction of Proposition 3 is largely consistent with the experimental evidence of the JOD.

In the control treatment *JOD-C* the game did not end when A 's offer to sell her share was rejected. Instead, the game continued with joint ownership. This improves B 's threat point payoff when A chooses to make an offer which makes it less attractive for A to sell her share to B . On the other hand, it is less risky for A to make an offer, because if her offer is rejected, the parties are just back to joint ownership. Nevertheless, the Appendix shows that Proposition 3 still applies on the equilibrium path, so the prediction for this control experiment is the same, again consistent with the experimental evidence of *JOD-C*.

4.2.2. Analysis of the A -ownership design

In this design A is initially the sole owner of the project. At stage 0 she can choose whether to remain the sole owner and hire B as an employee at wage w , or to give half of the ownership rights to B , in which case there is joint ownership.

If A opts for joint ownership, the analysis is the same as in the previous subsection. Let us now suppose that she decides to retain A -ownership and hire B as an employee. A

self-interested type of player A finds it optimal to make a greedy wage offer that is going to be rejected by the fair-minded type of B . However, if the wage offer is accepted, the choice of $a = 10$ is a dominant strategy for both the self-interested and the fair-minded owner at the investment stage (as in the case of B -ownership above). Consider now player B 's investment choice. Anticipating A 's investment, the self-interested player B clearly chooses $b = 1$. The fair-minded player B also chooses $b = 1$ if his wage is sufficiently small ($w < 67$). Otherwise he will choose b so as to equalise payoffs. This parallels the analysis of B -ownership in the previous subsection.

Consider now stage 0. The self-interested player B will accept the contract offered by A if and only if $w \geq 12$. The fair-minded player B accepts any contract where $w \geq 56$. Thus, only the self-interested type of B will accept an offer of $w = 12$. If A offers $w = 56$ both types of B will accept the offer, yielding a higher payoff.²⁰ It is easy to show that offering more than 56 reduces A 's payoff, so a self-interested player A will offer $w = 56$ which is accepted by both types of B . The fair-minded player A wants to equalise payoffs and offers $w = 67$, which is also accepted by both types of player B . Hence, the theory of inequity aversion predicts the same investment levels as the self-interest theory but it differs in the prediction of the wages offered to B .

Let us now turn to the entire game. Again, if we are willing to impose a condition that parallels Condition 1, we get a unique equilibrium prediction.

CONDITION 1'. *If A offers a joint ownership contract and gives half of the firm's revenues to B , then B 's updated belief that he faces the self-interested type of A does not increase.*

PROPOSITION 4 [A-Ownership Design]. *With incomplete information about the players types there exists a unique Perfect Bayesian Equilibrium outcome satisfying Condition 1'.*

The equilibrium is a pooling equilibrium in which both types of A offer a joint ownership contract which both types of player B accept. The equilibrium outcome is the same as under joint ownership described in Proposition 3.

Thus, Propositions 3 and 4 predict that the players will always end up with joint ownership, regardless of the initial allocation of ownership rights. Joint ownership does not implement first best investment decisions. However, it still outperforms A (or B) ownership. Both parties invest with a significant probability under joint ownership, while only the owner invests under A (or B) ownership and the other party opts for the minimum investment. Furthermore, A has to hire B as an employee under A -ownership and there is a significant probability that a fair-minded player B will reject a wage offer that he perceives as unfair. Similarly, if A sells her ownership stake in the joint ownership design there is again some probability that this offer is going to be rejected. Therefore, the model of inequity aversion predicts that joint ownership is more efficient. This is in contrast to the self-interest model that predicts A (or B) ownership to be more efficient. However, both models support the Coase Theorem which suggests that the parties will adopt the ownership structure that is most efficient independent of the initial allocation of ownership rights. Again, the ownership prediction of the theory

²⁰ The minimum wage that both types of player B accept depends on the value of α . However, the A -player will choose joint ownership in equilibrium for all $\alpha > 0.5$.

of inequity aversion with heterogeneous types is largely consistent with the experimental evidence of this treatment.

5. Conclusions

Our experiments confirm the property rights approach by showing that the ownership structure affects relationship-specific investments and that the subjects achieve the most efficient ownership allocation starting from different initial conditions. However, in our experiments the most efficient ownership structure is joint ownership, which contrasts with the property rights approach. Of course, joint ownership is not always optimal. Our experimental and theoretical analysis suggests that joint ownership is more likely to be optimal if the number of partners is small, if the free-rider problem is not too severe and if there is no other way to contract on the relationship specific investments of the partners. This may explain why joint ownership is often observed in small firms in the service industries mentioned in the introduction but much less frequently in larger firms.

Our results are neither consistent with the self-interest model nor with models that assume that all people behave fairly but they are largely consistent with the theory of inequity aversion that focuses on the interaction between selfish and fair players. The theory suggests that the reason for the superiority of joint ownership is that it makes better (but still imperfect) use of fairness as an enforcement device than does A or B-ownership.

In a recent paper, Hart (2001) argues that ‘although norms are undoubtedly very important both inside and between firms, incorporating them into the theory has been very difficult and is likely to continue to be so in the near future’ and that ‘a norm-free theory of the firm and a norm-rich theory of the firm don’t seem to have very different predictions’. He mainly examined models of repeated games that try to capture norms that are based on repeated interactions of self-interested players. In this article, we have shown that the recent advances in modelling fairness in one-shot games provide powerful tools for incorporating norms of fair behaviour into contract theory. This allows us to derive testable predictions on the optimal allocation of ownership rights, some of which differ significantly from the standard predictions of the self-interest model. Thus, our experiments as well as the theoretical analysis suggest that if we want to understand the incentive effects of real institutions on real people, concerns for fairness should be taken into account.

University of Zurich, CESifo and CEPR

University of Munich

University of Munich, CESifo and CEPR

Submitted: 21 July 2005

Accepted: 1 July 2007

References

Bar-Isaac, H. (2007). ‘Something to prove: reputation in teams’, *RAND Journal of Economics*, vol. 38(2), pp. 495–511.

- Bolton, G.E. and Ockenfels, A. (2000). 'A theory of equity, reciprocity and competition', *American Economic Review*, vol. 100, pp. 166–93.
- Camerer, C.F. (2003). *Behavioral Game Theory*, Princeton: Princeton University Press.
- Charness, G. and Rabin, M. (2002). 'Understanding social preferences with simple tests', *Quarterly Journal of Economics*, vol. 117, pp. 817–69.
- Chiappori, P.-A. and Salani, B. (2003). 'Testing contract theory: a survey of some recent work', in (M. Dewatripont *et al.*, eds.), *Advances in Economic Theory, Eighth World Congress of the Econometric Society*, vol. 1, pp. 115–49, Cambridge: Cambridge University Press.
- Coase, R.H. (1960). 'The problem of social cost', *Journal of Law and Economics*, vol. 3, pp. 1–44.
- Dufwenberg, M. and Kirchsteiger, G. (2004). 'A theory of sequential reciprocity', *Games and Economic Behavior*, vol. 47, pp. 268–98.
- Ellingsen, T. and Johannesson, M. (2004a). 'Promises, threats, and fairness', *ECONOMIC JOURNAL*, vol. 114, pp. 397–420.
- Ellingsen, T. and Johannesson, M. (2004b). 'Is there a hold-up problem?', *Scandinavian Journal of Economics*, vol. 106, pp. 475–94.
- Engelmann, D. and Strobel, M. (2004). 'Inequality, efficiency and maximin preferences in simple distribution games', *American Economic Review*, vol. 94, pp. 857–69.
- Falk, A. and Fischbacher, U. (2006). 'A theory of reciprocity', *Games and Economic Behavior*, vol. 54, pp. 293–315.
- Falk, A., Fehr, E. and Fischbacher, U. (2003). 'On the nature of fair behaviour', *Economic Inquiry*, vol. 41, pp. 20–6.
- Fehr, E., Gächter, S. and Kirchsteiger, G. (1997). 'Reciprocity as a contract enforcement device', *Econometrica*, vol. 65, pp. 833–60.
- Fehr, E., Klein, A. and Schmidt, K.M. (2007). 'Fairness and contract design', *Econometrica*, vol. 75(1), 121–54.
- Fehr, E., Naef, M. and Schmidt, K. (2006). 'Inequality, efficiency and maximin preferences in simple distribution games: comment', *American Economic Review*, vol. 96, pp. 1912–7.
- Fehr, E. and Schmidt, K.M. (1999). 'A theory of fairness, competition and co-operation', *Quarterly Journal of Economics*, vol. 114, pp. 817–68.
- Fehr, E. and Schmidt, K.M. (2003). 'Theories of fairness and reciprocity: evidence and economic applications', in (M. Dewatripont *et al.*, eds.), *Advances in Economic Theory, Eighth World Congress of the Econometric Society*, vol. 1, pp. 208–57, Cambridge: Cambridge University Press.
- Fehr, E., and Schmidt, K.M. (2004). 'Fairness and incentives in a multi-task principal-agent model', *Scandinavian Journal of Economics*, vol. 106(3), 453–74.
- Gantner, A., Güth, W., and Königstein, M. (2001). 'Equitable choices in bargaining games with joint production', *Journal of Economic Behavior and Organization*, vol. 46, pp. 209–25.
- Geanakoplos, J., Pearce, D. and Stacchetti, E. (1989). 'Psychological games and sequential rationality', *Games and Economic Behavior*, vol. 1, pp. 60–79.
- Grossman, S. and Hart, O. (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. (1994). 'Is relational exchange possible in the absence of reputations and repeated contact?', *Journal of Law, Economics, and Organization*, vol. 10, pp. 360–89.
- Halonen, M. (2002). 'Reputation and the allocation of ownership', *ECONOMIC JOURNAL*, vol. 112, pp. 539–58.
- Hart, O. (1995). *Firms, Contracts and Financial Structure*, Oxford: Clarendon Press.
- Hart, O. (2001). 'Norms and the theory of the firm', *University of Pennsylvania Law Review*, vol. 149, pp. 1701–15, Reprinted in (E. Brousseau and J.-M. Glachant, eds), *The Economics of Contracts*, Cambridge: Cambridge University Press, 2002.
- Hart, O. and Moore, J. (1990). 'Property rights and the nature of the firm', *Journal of Political Economy*, vol. 98, pp. 1119–58.
- Königstein, M. and Tietz, R. (2000). 'Profit sharing in an asymmetric bargaining game', in (M. Königstein, ed.), *Equity, Efficiency and Evolutionary Stability in Bargaining Games with Joint Production*, pp. 5–32, Berlin: Springer.
- Levin, J. and Tadelis, S. (2005). 'Profit sharing and the role of professional partnerships', *Quarterly Journal of Economics*, vol. 120, pp. 131–71.
- Noldeke, G. and Schmidt, K.M. (1998). 'Sequential investments and options to own', *RAND Journal of Economics*, vol. 29, pp. 633–53.
- Oosterbeek, H., Sonnemans, J., and van Velzen, S. (2003). 'The need for marriage contracts: an experimental study', *Journal of Population Economics*, vol. 16, pp. 431–53.
- Rabin, M. (1993). 'Incorporating fairness into game theory and economics', *American Economic Review*, vol. 83, pp. 1281–302.
- Rajan, R. G. and Zingales, L. (1998). 'Power in a theory of the firm', *Quarterly Journal of Economics*, vol. 113, pp. 387–432.
- Ward, B. (1958). 'The firm in Illyria: market syndicalism', *American Economic Review*, vol. 48, pp. 566–89.
- Williamson, O. (1985). *The Economic Institutions of Capitalism*. New York: Free Press.