

DISCUSSION PAPER SERIES

No. 5369

FAIRNESS AND THE OPTIMAL ALLOCATION OF OWNERSHIP RIGHTS

Ernst Fehr, Susanne Kremhelmer and Klaus
M. Schmidt

INDUSTRIAL ORGANIZATION



Centre for **E**conomic **P**olicy **R**esearch

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP5369.asp

FAIRNESS AND THE OPTIMAL ALLOCATION OF OWNERSHIP RIGHTS

Ernst Fehr, University of Zurich and CEPR
Susanne Kremhelmer, University of Munich
Klaus M. Schmidt, University of Munich and CEPR

Discussion Paper No. 5369
December 2005

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Ernst Fehr, Susanne Kremhelmer and Klaus M. Schmidt

CEPR Discussion Paper No. 5369

December 2005

ABSTRACT

Fairness and the Optimal Allocation of Ownership Rights*

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 are neither consistent with the self-interest model nor with models that assume that all people behave fairly, but they can be explained by the theory of inequity aversion that focuses on the interaction between selfish and fair players.

JEL Classification: C7, C9 and J3

Keywords: double moral hazard, fairness, incomplete contracts, ownership rights and reciprocity

Ernst Fehr
Institute for Empirical Research in
Economics
University of Zurich
Bluemlisalpstrasse 10
CH-8006 Zurich
SWITZERLAND

Susanne Krehmelmer
Department of Economics
University of Munich
Ludwigstrasse 28
D-80539 Muenchen
GERMANY

Tel: (41 1) 634 37 09
Fax: (41 1) 634 49 07
Email: efehr@iew.unizh.ch

Email: susanne.krehmelmer@Lrz.uni-muenchen.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=135748

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=163679

Klaus M. Schmidt
Department of Economics
University of Munich
Ludwigstrasse 28
D-80539 München
GERMANY

Tel: (49 89) 2180 2250
Fax:(49 89) 2180 3510
Email: klaus.schmidt@Lrz.uni-
muenchen.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=118063

*Corresponding Author: Schmidt. Alexander Klein provided excellent research assistance for the experiments discussed in this paper. We would like to thank Oliver Hart, Bentley MacLeod and seminar participants at Bocconi University, Boston University, the University of Virginia, the NBER's workshop on organization theory, the CESifo area conference on microeconomics and the conference on Economics and Psychology in Toulouse for helpful comments and suggestions. Part of this research was conducted while the third author visited Yale University and he would like to thank 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).

Submitted 11 November 2005

1. Introduction

The optimal allocation of ownership rights has interested economists for decades, if not for centuries. In this paper we examine this question by combining the modern property rights approach with recent insights into the nature of human social preferences. The modern property rights approach, pioneered by Grossman and Hart (1986) and Hart and Moore (1990), argues that the allocation of ownership rights matters in a world of incomplete contracts because it assigns residual rights of control. These control rights affect the involved parties' incentives to make relationship specific investments that cannot be contracted upon directly. Grossman, Hart, and Moore show that typically no allocation of ownership rights induces first best investments, but that some ownership structures are more efficient than others. Therefore, the involved parties will trade ownership rights ex ante in order to set up the second best optimal ownership structure.¹ Another prominent result of this literature (see Hart 1995, p. 48, e.g.) is that joint ownership is rarely optimal; if both parties own an asset, they then can prevent each other from using it, minimizing both parties' threat point payoffs. Giving all the ownership rights to one party increases this party's threat point payoff and investment incentives, without affecting the other party's incentives.

In the first part of this paper, we report on several experiments on the optimal allocation of ownership rights, where players first have to bargain over the allocation of ownership rights on a joint project (a "firm"). They can either decide on joint ownership, or one of the parties can be the sole owner and hire 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 this ownership structure ex ante, i.e., do they understand what the most efficient ownership structure is and do they manage to implement it? In order to address the second question we consider two different experimental designs. In the Joint Ownership Design (JOD), joint ownership determines the status quo, but

¹ The earlier literature on property rights comes to somewhat different conclusions. Coase (1960) emphasizes that the clear definition of property rights is of crucial importance for economic efficiency. However, the so called "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.

one of the parties can try to sell her ownership stake to the other. In the A-Ownership Design (AOD), A owns the whole project initially, but she can give half of her ownership rights to B. Thus, we can test both whether the initial allocation of ownership rights affects the final outcome and whether the parties always manage to set up the efficient ownership structure.

The 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, the most efficient ownership structure turns out to be joint ownership, contrasting sharply with the property rights approach predictions.

We offer a theoretical interpretation of the experimental results in the second part of the paper. 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 evidence² demonstrate that concerns for fairness and reciprocity play an important role in motivating the behavior of many people. Several experiments (e.g. Fehr, Gächter, and Kirchsteiger, 1997, and Fehr, Klein, and Schmidt, 2004) 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 5, 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. Fairness suffices as an enforcement device for inducing both parties to invest efficiently regardless of the allocation of ownership rights. 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 A or B-ownership did both parties invest efficiently. Second, this approach fails to explain why relatively high (second best optimal) investment levels are induced under joint ownership whereas under A or B-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-Schmidt (1999) model of inequity aversion, we show that in this

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

case, 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 empirical regularities regarding investment behavior. Moreover, this approach also predicts that the second best efficient ownership allocation will be implemented regardless of the initial ownership structure. This is indeed what we observed in the experiments. The large majority of the subjects seem to have understood that joint ownership is more efficient than A or B-ownership. As a consequence, joint ownership prevailed in the Joint Ownership Design and in the A-Ownership Design.

The question which allocation of ownership rights is optimal is a question about the incentive properties of different governance structures. It is difficult to empirically address this question by using field data. As has been emphasized by Chiappori and Salanié (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 behaviors. Does a particular ownership structure induce a particular type of behavior or are behavioral 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 behavior of the same individuals under different ownership structures.

Our paper 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, Güth and Königstein (2001), Königstein and Tietz (2000) and Oosterbeek, Sonnemans and van Velzen (2003). All of these papers show that concerns for fairness mitigate the hold-up problem. Ellingsen and Johannesson (2004a, 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 (2004a, b) apply the Fehr-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 paper.

A few theoretical models show that joint ownership may occasionally be optimal. Rosenkranz and Schmitz (1999) consider a set-up where parties have to engage in multiple types of investments. They show that joint ownership improves the incentives for engaging in know-how disclosure, possibly rendering joint ownership the optimal ownership structure. Maskin and Tirole (1999) consider a buyer-seller relationship and show that joint ownership combined with an option for selling the owned share of the asset to the other party can implement first best incentives. Halonen (2002) considers an infinitely repeated game. In the one-shot game, joint ownership is the worst possible ownership structure, as it minimizes investment incentives.³ However, in the infinitely repeated game this ownership structure is desirable because it provides the strongest punishment possibilities if one of the parties deviates. None of these papers considers the effects of fairness which we show to be empirically important for the ownership allocation. These papers also do not study the endogenous determination of ownership rights. Instead, they implicitly or explicitly assume that the most efficient ownership structure will prevail. As we explicitly allow the trading of property rights, we examine empirically whether efficient property rights will indeed be implemented.

The rest of the paper is organized as follows. Section 2 sets up a simple problem of the allocation of ownership rights. Experimental procedures are discussed in Section 3. Section 4 presents the experimental results. We offer a theoretical analysis of the experiment under different assumption about preferences and concerns for fairness in Section 5. Section 6 concludes the paper.

2. 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 $v(a,b) = v(b,a)$.

³ See also Hart (2001) for a critical discussion of the repeated games approach to the theory of the firm.

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^* ,

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

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

$$M^A = v(a, b) - w - c(a)$$

$$M^B = w - c(b)$$

Clearly, a self-interested B player will choose $b = \underline{b}$ in this case because the marginal return on his investment is zero. The A player, however, is 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

$$M^A = 0.5v(a, b) - c(a)$$

$$M^B = 0.5v(a, b) - c(b)$$

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.⁴

Note that this is a reduced form of Hart (1995) with investments in physical rather than in human assets.⁵ Note also that joint ownership is never optimal with investments in

⁴ 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, it cannot be ruled out that B overinvests at this general level (see Noldeke and Schmidt, 1998, Proposition 1). However, there is underinvestment in equilibrium under any ownership structure in the parameterization of the experiments that we conducted here.

⁵ 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 in case of A-ownership. If both parties own the asset, each of them can block the other from using the asset, so they share the surplus equally. If the investment were in human capital, 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. Therefore, in order to achieve B's cooperation, she would have to negotiate with B and to give some of the surplus $v(a, b)$ to him. We decided not to implement

human assets, while it may be optimal with investments in physical assets.⁶ However, we chose the parameters of the surplus and cost functions in our experiments in such a way that A (or B) ownership always dominates joint ownership if players are self-interested (see Section 3 below).

The parties bargain on the allocation of ownership rights before the 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)*: 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.

3. Experimental Design and Procedures

We used the following parameterization of this investment problem in our experiments. The two parties choose $a, b \in \{1, \dots, 10\}$. Investments (a, b) yield a gross surplus $v(a, b) = 22 \cdot (a + b)$, while investment costs are $c_A(a) = 12a$ and $c_B(b) = 12b$. Thus,

the negotiation process in the experiment. In our experiment the parties have to negotiate on the allocation of ownership rights already. Thus, a second negotiation game would have complicated the experimental design considerably. Furthermore, the outcome of the first negotiation game (on the allocation of ownership rights) might affect behavior in the second (on how to divide the ex post surplus) which would have been difficult to disentangle. See Hart (1995, p. 68) and Noldeke and Schmidt (1998) for a more detailed discussion of investments in human and in physical capital.

⁶ See Hart (1995), p. 48 and p. 69.

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.

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 always remains constant. 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.⁷
- *Control Treatment (JOD')*: 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. 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. However, if we continue to observe joint ownership predominantly in this treatment, our results become even stronger.

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 (JOD), two sessions (S5–S6) implemented the A-Ownership Design, and two additional sessions (S7 and S8) implemented the control treatment (JOD'). 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

⁷ We did not use the expressions "employee" and "partner" in the actual experiments, but the neutral terms "actor A" and "actor B" instead.

relationship. They neither observed the contracts chosen by or offered to the other subjects in the room nor their current partner's past behavior. 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.⁸

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).

4. Experimental Results

4.1. Joint Ownership Design

Each party initially owned 50 percent of the joint venture in the Joint Ownership Design (JOD). 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 summarized as follows:

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

Support for Result 1 is provided by Figure 1 and the following numbers: A players prefer joint ownership in 300 out of 470 (63.8 percent) cases. In the remaining 170 cases (36.2 percent), the A players attempt to sell their ownership stake to the 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 percent 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 proposed and accepted or proposed and rejected. Recall that if A's offer to sell her

⁸ The full set of all our experimental instructions, in the original German and translated into English, are available at http://www.vwl.uni-muenchen.de/l_s_schmidt/experiments/property_rights/index.htm.

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 percent of the cases, whereas accepted B-ownership varies between 20 and 30 percent of the cases. There is no time trend in the data.⁹




Figure 1: Share of joint ownership, realized B-ownership and rejected B-ownership over time in the JOD

Result 1 clearly contradicts the self-interest model, which predicts the choice of B-ownership because it supposedly induces more efficient investment behavior. In order to understand this result, we have to analyze the actual investment behavior 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

⁹ A random-effects GLS regression shows that the probability of choosing joint ownership goes down by 1 percent per period, but this effect is not significant at the 5 percent level ($p=0.068$).

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 $a = 6.7$.




Figure 2: A's average investment response to b under joint ownership (JOD)

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.¹⁰ 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 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

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

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

¹⁰ 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.

Dependent variable:	(1)	(2)	(3)	(4)
A's investment a	(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.

Table 1: Player A's investment as a function of player B's investment

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 behavior. In many cases the A players exactly matched the investment level of B but there are also a substantial number of cases where the A players did not fully reciprocate B's choice. In fact, in 28.3 percent of the cases the A player chose the minimal investment level of $a = 1$ and in 25 percent 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 percent 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 percent) 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.

b/ 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

Table 2: Investments (b,a) with joint ownership contracts in JOD

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.¹¹

¹¹ 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 ($p=0.0001$).




Figure 3: Payoffs over time 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 percent) 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 percent of all offers). The average price of the rejected offers was 176.9, while the average price of the accepted offers was just 120.8.¹² 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 percent) 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.

4.2. A-Ownership Design

¹² In fact, a Mann-Whitney Test confirms that the differences between the accepted and rejected offers are highly statistically significant ($p=0.0000$).

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 summarized 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.*

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.¹³ Overall, joint ownership was chosen in 81.3 percent of all cases (187 of 230 observations). Thus, A players clearly preferred joint ownership.




Figure 4: Share of joint ownership and A-ownership over time in the AOD

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

¹³ 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 ($p=0.056$).

maximizes B's expected earnings. There is, however, significant heterogeneity in A's investment response.

In 135 out of 187 (72.2 percent of the) joint ownership cases player B chose $b = 10$ and in less than 10 percent 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 to 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) related to investment behavior under joint ownership where joint ownership is the initial condition, regressions (3) and (4) relate to behavior 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, $b = 10$ maximizes B's expected monetary payoff.

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

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.

b/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

Table 3: Investments (b,a) with joint ownership contracts in AOD

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 than 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 percent). B players rejected the wage offers 5 times. The average wage offer was 62.1. Almost all B players chose the minimum effort level (average $b = 1.3$) under A-ownership, while all A players chose $a = 10$. Thus, average total investment $a + b = 11.26$, 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.¹⁴

¹⁴ 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).

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.

4.3. Control Treatment for Joint Ownership Design (JOD')

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 percent). However, A players tried to sell their ownership stakes more frequently in this case (in 134 out of 240 cases, 55.8 percent). 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 percent) than in the JOD treatment. Let us consider the three different possibilities in turn:

- If A did not make an offer, the investment behavior is very similar to the investment behavior 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 behavior 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 behavior 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.

5. Theoretical Interpretation

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

***Proposition 1:** If both parties are only interested in maximizing their own material payoff, then*

1. *A and B-ownership are equally efficient. Both are more efficient than joint ownership and*
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, Pearce and Stacchetti (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 both 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. 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 not only able to explain the qualitative, but also the quantitative, results of many classes of experimental games fairly well.¹⁵ For simplicity, we will therefore apply the Fehr and Schmidt model to our treatment conditions.¹⁶

We will first consider the case where all players are strongly concerned about fairness. We will 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 analyze these two cases in more detail in the rest of this section.

5.1. 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

¹⁵ There are also a few models that try to model preferences for fair outcomes and fair intentions simultaneously, in particular Falk and Fischbacher (2005) 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.

¹⁶ The Bolton and Ockenfels model makes the same qualitative predictions regarding the pattern of the ownership structure in the context of our treatments.

but also care about inequity or, in our context, inequality.¹⁷ The utility function of inequity averse (fair) players in the two-player case is given by

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

$i \in \{1, 2\}$, $i \neq j$, where $x = (x_1, x_2)$ denotes the vector of monetary payoffs, $\beta_i \leq \alpha_i$ and $0 \leq \beta_i < 1$. 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. $0.5 \leq \beta_i < 1$. 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 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 equalizes 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 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 maximizing 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. 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

¹⁷Fairness 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.

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 $w = 1/2 \cdot v(a^*, b^*)$, 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, like in the self-interest model, this prediction contrasts starkly to the experimental evidence.

5.2. 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 don't know whether they face a fair and trustworthy opponent or whether their opponent is going to exploit them. We use again the Fehr-Schmidt (1999) model of inequity aversion, but we now allow for heterogeneous fairness preferences. In accordance with Fehr and Schmidt (1999), we assume that there are 60 percent self-interested types ($\alpha_i = \beta_i = 0$) and 40 percent "fair" types. We assume that fair subjects exhibit $\alpha_i > 0.5$ and $0.5 < \beta_i < 1$. Thus, if inequality is to their disadvantage, they are prepared to engage in costly "punishment" in order to reduce their opponent's payoff. On the other hand, if the inequality is to their advantage, they are willing to spend resources in order to benefit the other player. Note that subjects with $1 > \beta_i > 0.5$ and

$\alpha_i > 0.5$ are willing to share the surplus of a contract equally but are also prepared to reject offers that give them less than 25 percent of the surplus. Fehr and Schmidt demonstrated that the evidence from many experiments is roughly compatible with the assumption that 40 percent of all subjects fall in this category.¹⁸

The property rights game can be analyzed on the basis of these assumptions, using standard game theoretic tools. For concreteness, we assume that a fair type has $\alpha_i = 2$, which is consistent with the results in Fehr and Schmidt (1999), but the qualitative results reported below are independent of this assumption. The full analysis is not difficult but somewhat lengthy and therefore relegated to an appendix that can be found on our webpage.¹⁹ In the following, we report the main predictions for our experiments and give the intuition for them.

5.2.1. 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. Therefore, let p denote the (endogenously determined) probability B assigns to the event that he faces the self-interested type of player A. We first analyze the players' behaviors at the investment stage, after which we will 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 ($p < 10/11 = 0.91$).²⁰ 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 $p < 10/35 = 0.29$.²¹ This result is

¹⁸ See Fehr and Schmidt (1999) for a more extensive discussion of the experimental evidence on the distribution of inequity averse types. When they calibrated their model to explain the quantitative evidence in the different games, they used four different types, but aggregated they also have that 40 percent of subjects exhibit $\alpha_i \geq \beta_i > 0.5$ and that 60 percent exhibit $0.5 > \alpha_i \geq \beta_i \geq 0$.

¹⁹ Please visit: http://www.vwl.uni-muenchen.de/ls_schmidt/experiments/property_rights/index.htm.

²⁰ 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 .

²¹ The fair-minded player B is also concerned about the inequality to his advantage generated by the selfish type of player A. So he maximizes $EU^{Bf} = p[11(a^s+b) - 12b - \alpha(12b - 12a^s)] + (1-p)[11(a^f+b) - 12b]$. 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 $p < 10/35$.

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 the fair-minded types of 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 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 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 .

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 prediction of Proposition 3 is largely consistent with the experimental evidence of the JOD.

In the control treatment *JOD'* 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 exactly the same, again consistent with the experimental evidence of *JOD'*.

5.2.2. Analysis of the A-Ownership Design

A is initially the sole owner of the project in this design. She can choose at stage 0 whether to remain the sole owner and hire B as an employee at wage w , or whether 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 equalize 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.²² 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 equalize 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.*

²² 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 $\alpha > 0.5$.

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

6. 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, the most efficient ownership structure is joint ownership, which contrasts with the property rights approach. These results are neither consistent with the self-interest model nor with models that assume that all people behave fairly, but they can be explained by 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 implicit enforcement device than does A or B-ownership.

In a recent paper, Oliver 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 paper, we have shown that the recent advances in modeling fairness in one-shot games provide powerful tools for incorporating norms of fair behavior into contract theory. This allows us to derive important and testable predictions on the optimal allocation of ownership rights, some of which differ significantly from the standard predictions of the self-interest model. If we want to understand the incentive properties of real institutions on real people, concerns for fairness have to be taken into account.

References

- Bolton, Gary E. and Ockenfels, Axel, 2000. A Theory of Equity, Reciprocity and Competition. *American Economic Review* 100, 166-193.
- Camerer, Colin F., 2003. *Behavioral Game Theory*, Princeton: Princeton University Press.
- Charness, Gary, and Rabin, Matthew, 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117, 817-869.
- Chiappori, Pierre-Andre and Salanié, Bernard, 2003. "Testing Contract Theory : A Survey of Some Recent Work." In: M. Dewatripont, L.P. Hansen, S. Turnovski, *Advances in Economic Theory, Eighth World Congress of the Econometric Society*, Vol. 1, Cambridge: Cambridge University Press, 115-149.
- Coase, Ronald H., 1960. "The Problem of Social Cost." *Journal of Law and Economics* 3, 1-44.
- Dufwenberg, Martin and Kirchsteiger, Georg, 2004. "A Theory of Sequential Reciprocity." *Games and Economic Behavior* 47, 268-298.
- Ellingsen, Tore and Johannesson, Magnus, 2004a. "Promises, Threats, and Fairness." *Economic Journal* 114, 397-420.
- Ellingsen, Tore and Johannesson, Magnus, 2004b. "Is there a Hold-up Problem?" *Scandinavian Journal of Economics* 106, 475-494.
- Falk, Armin and Fischbacher, Urs, 2005. "A Theory of Reciprocity." *Games and Economic Behavior*, in press.
- Fehr, Ernst, Gächter, Simon and Georg Kirchsteiger, 1997. "Reciprocity as a Contract Enforcement Device." *Econometrica* 65, 833-860.
- Fehr, Ernst, Klein, Alexander and Schmidt, Klaus M., 2004. "Contracts, Fairness, and Incentives." *CESifo Working Paper* No. 1215.
- Fehr, Ernst and Schmidt, Klaus M., 1999. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics* 114, 817-868.
- Fehr, Ernst and Schmidt, Klaus 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*, Cambridge: Cambridge University Press, 208-257.

- Gantner, Anita, Güth, Werner, and Königstein, Manfred, 2001: "Equitable Choices in Bargaining Games with Joint Production." *Journal of Economic Behavior and Organization* 46, 209-225.
- Geanakoplos, John, Pearce, David and Stacchetti, Ennio, 1989. "Psychological Games and Sequential Rationality." *Games and Economic Behavior* 1, 60-79.
- Grossman, Sanford and Hart, Oliver, 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy* 94, 691-719.
- Hackett, Steven C., 1994. "Is Relational Exchange Possible in the Absence of Reputations and Repeated Contact?" *Journal of Law, Economics, and Organization* 10, 360-389.
- Halonen, Maija, 2002. "Reputation and the Allocation of Ownership." *The Economic Journal* 112, 539-558.
- Hart, Oliver, 1995. *Firms, Contracts and Financial Structure*, Oxford: Clarendon Press.
- Hart, Oliver, 2001. "Norms and the Theory of the Firm." *University of Pennsylvania Law Review* 149, 1701-15. Reprinted in E. Brousseau and J.-M. Glachant (eds), *The Economics of Contracts*, Cambridge University Press, 2002.
- Hart, Oliver and Moore, John, 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy* 98, 1119-1158.
- Königstein, Manfred and Tietz, Reinhard, 2000. "Profit Sharing in an Asymmetric Bargaining Game." In: Manfred Königstein (ed.), *Equity, Efficiency and Evolutionary Stability in Bargaining Games with Joint Production*, Berlin: Springer, 5-32.
- Maskin, Eric and Jean Tirole, 1999. "Two Remarks on the Property Rights Literature." *Review of Economic Studies* 66, 139-149.
- Noldeke, Georg and Schmidt, Klaus M., 1998. „Sequential Investments and Options to Own.“ *RAND Journal of Economics* 29, 633-653.
- Oosterbeek, Hessel, Sonnemans, Joep, and van Velzen, S., 2003. "The Need for Marriage Contracts: An Experimental Study." *Journal of Population Economics* 16, 431-453.
- Rabin, Matthew, 1993. "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, 83, 1281-1302.
- Rosenkranz, Stephanie and Patrick W. Schmitz, 1999. „Know-how Disclosure and Incomplete Contracts.“ *Economics Letters* 63, 181-185.
- Williamson, Oiver, 1985. *The Economic Institutions of Capitalism*. New York: Free Press.