

Hildenbrand, Andreas

Working Paper

Is a firm a firm? A Stackelberg experiment

MAGKS Joint Discussion Paper Series in Economics, No. 29-2012

Provided in Cooperation with:

Faculty of Business Administration and Economics, University of Marburg

Suggested Citation: Hildenbrand, Andreas (2012) : Is a firm a firm? A Stackelberg experiment, MAGKS Joint Discussion Paper Series in Economics, No. 29-2012, Philipps-University Marburg, Faculty of Business Administration and Economics, Marburg

This Version is available at:

<https://hdl.handle.net/10419/73094>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

MAGKS



**Joint Discussion Paper
Series in Economics**

by the Universities of
**Aachen · Gießen · Göttingen
Kassel · Marburg · Siegen**

ISSN 1867-3678

No. 29-2012

Andreas Hildenbrand

**Is a “Firm” a Firm?
A Stackelberg Experiment**

This paper can be downloaded from
http://www.uni-marburg.de/fb02/makro/forschung/magkspapers/index_html%28magks%29

Coordination: Bernd Hayo • Philipps-University Marburg
Faculty of Business Administration and Economics • Universitätsstraße 24, D-35032 Marburg
Tel: +49-6421-2823091, Fax: +49-6421-2823088, e-mail: hayo@wiwi.uni-marburg.de

Is a “Firm” a Firm?

A Stackelberg Experiment*

Andreas Hildenbrand[†]

May 20, 2012

*For detailed comments, I am grateful to Max Albert, Matthias Greiff, Michael Hesch, and Björn Kuchinke. I thank participants at the 2011/12 Tübingen Economic Workshop and the 2012 Hohenheimer Oberseminar for useful suggestions. I also thank Steffen Huck, Wieland Müller, and Hans-Theo Normann for providing their data.

[†]VWL VI: Behavioral and Institutional Economics, Justus Liebig University Giessen, Licher Strasse 66, 35394 Giessen, Germany. Telephone: +49 641 99-22203. Telefax: +49 641 99-22209. Electronic Mail: andreas.hildenbrand@wirtschaft.uni-giessen.de. World Wide Web: wirtschaft.uni-giessen.de.

Abstract

Industrial organization is mainly concerned with the behavior of large firms. Experimental industrial organization therefore faces a problem: How can firms be brought into the laboratory? The main approach relies on framing: Call individuals “firms”! This experimental approach is not in line with modern industrial organization, according to which a firm’s market behavior is also determined by its organizational structure. In this paper, a Stackelberg experiment is considered in order to answer the question whether framing individual decision making as organizational decision making or implementing an organizational structure is more effective in generating profit-maximizing behavior. Firms are either represented by individuals or by teams. I find that teams’ quantity choices are more in line with the assumption of profit maximization than individuals’ choices. Compared to individuals, teams appear to be less inequality averse.

Keywords: industrial organization, Stackelberg game, individual behavior, team behavior, framing, experimental economics.

Journal of Economic Literature classification codes: C720, C910, C920, D430, L130.

1. Introduction

For Tirole (1988, p. 3), industrial organization (IO) “certainly begins with the structure and behavior of firms”. So does it for Scherer and Ross (1990, p. 1). Firms are typically assumed to be large, especially when it comes to oligopoly theory (see also Shapiro, 1989). Holmstrom and Tirole (1989) emphasize that organizational structure and market behavior are interdependent. Discussing theories of the firm going beyond the neoclassical production function approach, they show how a firm’s market behavior is affected by its organizational structure (see also Furubotn and Richter, 2005, pp. 361–469).¹ In these theories, the firm is seen as an organizational structure, a nexus of contracts (contractual view). Nevertheless, in IO, firms are usually treated in the neoclassical way: They are regarded as single decision makers that maximize profits (see, e.g., Tirole, 1988, p. 4; Scherer and Ross, 1990, pp. 38, 52).² In the light of the contractual view, the objective of profit maximization can be interpreted as a consequence of the organizational structure transforming individual behavior within the firm into profit maximization on the market.³ If this view is taken, the question is: What organizational structure, if any, results in profit maximization? This question is of special interest if firms are assumed to be large, because large firms always have a nontrivial organizational structure.

Although firms are particularly assumed to be large in oligopoly theory, a quite different question is posed in most oligopoly experiments: How do individuals behave on an experimental market? The market structure is derived from an oligopoly game. No organizational structure is implemented. Instead, individuals are called “firms” (IO framing). Typically, participants are found to exhibit other-regarding preferences. There-

¹For a presentation of the neoclassical theory of the firm, see Nadiri (1982).

²More precisely, firms are assumed to be profit maximizers with rational expectations.

³This position is methodologically sound (see Albert and Hildenbrand, 2012).

fore, framing individual decision making as organizational decision making appears to be ineffective in generating profit-maximizing behavior: An IO framing does not turn individuals into profit maximizers. For example, Huck et al. (2001) proceed in this way. In a Stackelberg experiment (“STACKRAND” treatment, henceforth StackRand), they examine a market for a homogeneous product on which two firms sequentially compete in quantities.⁴ Firms are represented by individuals, and an IO framing is used. They find that, on average, leaders choose a lower quantity, and followers choose a higher quantity than predicted by the unique subgame-perfect Nash equilibrium (Stackelberg equilibrium). For the most part, these deviations can be explained by inequality aversion (see Huck et al., 2001, pp. 758–761; Lau and Leung, 2010).⁵

What is surprising is the fact that it is regularly not controlled for the effect supposed to be triggered by the IO framing: There are no control treatments in which participants are neutrally instructed. Maybe this is because framing is expected to be ineffective (see also Normann and Ruffle, 2011, p. 1). However, if that were expected, there would be no reason for an IO framing at all. If the IO framing is expected to affect individual behavior, two questions arise: First, what does the effect look like? Second, is the effect of framing individual decision making as organizational decision making different from the effect of implementing an organizational structure?

Hoffman et al. (1994) give an answer to the first question on the basis of an ultimatum game. Individuals are either called “seller” and “buyer”, or they are neutrally instructed. They find that proposers offer less under an IO framing than under a neutral framing. Responders reject offers under an IO framing as frequent as under a neutral framing: Rejection rates are the same for both treatments. Thus, at least proposer behavior is

⁴Huck et al. are regarded as the first to implement a Stackelberg market in the laboratory.

⁵In the recent past, many more experiments of that kind were run (see, e.g., Huck and Wallace, 2002). Often, the standard duopoly games were extended by a pre-play stage in order to endogenize the sequence of play. For example, see Huck et al. (2002), Fonseca et al. (2005), Fonseca et al. (2006), and Müller (2006). Hildenbrand (2010) reviews and discusses these experiments.

directed towards profit maximization.

On the basis of a Stackelberg game, an answer to the second question is given by Müller and Tan (2011).⁶ In their Stackelberg experiment, they study a market for a homogeneous product on which two firms sequentially compete in quantities. Firms are either represented by individuals or by three-member teams, and an IO framing is used in both treatments. Team members exchange electronic messages via a chat box in order to come to a unanimous agreement on their collective quantity. They find that, on average, individuals or teams in the leader role choose a lower quantity, and individuals or teams in the follower role choose a higher quantity than predicted by the Stackelberg equilibrium. Compared to individuals' quantity choices, teams' choices are not found to be more in line with the assumption of profit maximization. The same answer is given by Raab and Schipper (2009) on the basis of a Cournot game. In their Cournot experiment, they examine a market for a homogeneous product on which three firms simultaneously compete in quantities. Firms are either represented by individuals or by three-member teams, and an IO framing is also used. Individuals directly choose their quantities. Team members simultaneously choose efforts: The quantity of each team is the sum of its members' efforts. They find no difference between the market behavior of individuals and teams. In contrast to Müller and Tan's findings, both individuals' and teams' mean quantity choices are close to the unique Nash equilibrium (Cournot equilibrium) prediction.

Neither Müller and Tan nor Raab and Schipper find a difference between the market behavior of individuals and teams. However, a difference between the qualities of the predictions emerges: The asymmetric Stackelberg prediction fails for individuals under an IO framing and teams organized according to Müller and Tan's ad-hoc theory of

⁶The first oligopoly experiment taking account of an organizational structure was run by Sauermann and Selten (1959). In contrast to recent experiments, it was more of exploratory nature. More experiments of this kind were published by Selten (1967b,a).

the firm, whereas the symmetric Cournot prediction is always corroborated. In other Cournot experiments, individuals' quantity choices are also found to be in line with the assumption of profit maximization (see, e.g., Holt, 1985; Huck et al., 2004).

In other experiments on team decision making, teams are found to behave more in line with the assumption of profit maximization than individuals (see, e.g., Bornstein and Gneezy, 2002; Bornstein et al., 2008). The same is true for many experiments on group decision making, but there are opposite results, too.⁷ The term “team” is used if groups of participants collectively represent firms in a market experiment or if participants' collective decisions are framed as firms' decisions. Otherwise, the term “group” is used. For example, in their ultimatum experiment, Bornstein and Yaniv (1998) study a situation in which two players sequentially bargain over the division of a sum of money. Players are either represented by individuals or by three-member groups. Group members have face-to-face discussions to make their collective decision: A specific decision rule is not predetermined by the experimenters. Groups' decisions are found to be more in line with the assumption of profit maximization than individuals' decisions: groups in the proposer role offer less than individuals in this role, and groups in the responder role are willing to accept less than individuals in that role.

Two central differences between the experiments can be identified in order to explain the mixed results: the organizational structure and the market structure (see also Raab and Schipper, 2009, pp. 698–700). If the market structure is derived from a Cournot game, both individuals and teams appear to be profit maximizers. Their market behavior seems to be independent of the organizational structure. If the market structure is derived from an ultimatum game or a Stackelberg game, neither individuals nor teams or groups maximize profits. However, depending on the organizational structure, teams

⁷Bornstein (2008) and Engel (2010) characterize, review and discuss experiments on group decision making, and they also survey some experiments on team decision making.

or groups come closer to profit-maximizing behavior than individuals. An IO framing may direct individual behavior towards profit maximization, too. Therefore, an analysis of the effect of an IO framing or the implementation of an organizational structure requires a decision for a market structure first.

In order to find the limits of the profit maximization assumption, I conduct a Stackelberg experiment using Huck et al.'s (2001) market structure, which is also used by Müller and Tan (2011). Firms are either represented by individuals or by two-member teams. Individuals are either called "firms", or they are neutrally instructed. Teams are organized according to a parsimonious version of Alchian and Demsetz's (1972) contractual model of the firm. Teams each consist of a decision maker and a non-decision maker. The non-decision maker is completely passive. The decision maker is active and gets half of the team's monetary payoff. Specifically, I ask whether profit-maximizing behavior on Huck et al.'s Stackelberg market is generated by a type of organization which is in accordance with Alchian and Demsetz's model of the firm. I find that teams' quantity choices are more in line with the assumption of profit maximization than individuals' choices. Compared to individuals, teams appear to be less inequality averse. However, neither the applied IO framing nor the implemented organizational structure generates profit-maximizing behavior: Both individuals called "firms" and teams organized according to Alchian and Demsetz's contractual model appear not to be profit-maximizing firms.

The paper is organized as follows. Section 2 introduces the experimental design, the hypotheses, and the experimental procedures. Section 3 presents the experimental results, which are summarized and discussed in the final section. The appendix contains English translations of the instructions (originally written in German) and the payoff bimatrix used in all treatments.

2. Experimental design, hypotheses, and procedures

2.1. Design

The experimental market structure is derived from Huck et al.’s (2001) Stackelberg game. On a market for a homogeneous product, two firms compete in quantities. Firms A and B face a linear inverse demand function: $p(q) = \max\{30 - q, 0\}$, $q = q_A + q_B$. Each cost function is linear in output: $c_j(q_j) = 6q_j$, $j = A, B$. Hence, marginal costs are constant and identical. Firms sequentially decide how much to supply to the market: Firm A moves first (leads), and firm B moves second (follows). Because B observes A ’s quantity choice, B ’s action is a reaction to A ’s decision. Once market supply is determined, the market clears. Firms’ profits are equal to their revenues minus their production costs: $\pi_j(q_A, q_B) = (30 - (q_A + q_B))q_j - 6q_j$, $j = A, B$.

In the experiment, each firm is either represented by an individual or by a two-member team. Overall, there are three treatments: LOADED, NEUTRAL, and TEAM (see appendices A.1, A.2, and A.3 for English translations of the instructions). In LOADED and NEUTRAL, firms are each represented by an individual. In LOADED, participants are called “firms” competing in “quantities” for “profits” (IO framing). In NEUTRAL, participants are neutrally instructed: individuals choose “numbers” and receive “payments”. LOADED serves as a control treatment. It is identical to Huck et al.’s StackRand treatment. In TEAM, participants are also neutrally instructed. Each team consists of a decision maker and a non-decision maker. The non-decision maker is completely passive. The decision maker is active and gets half of the team’s monetary payoff. The non-decision maker receives the other half. An overview of the treatments, including StackRand, is given in Table 1.

In all treatments, each active participant plays ten (one-shot) Stackelberg games. In

	Instructions	Sessions	Participants
StackRand	loaded	2	$24 + 20 = 44$
LOADED	loaded	1	28
NEUTRAL	neutral	2	$22 + 22 = 44$
TEAM	neutral	1	48

Table 1: Own treatments and Huck et al.’s StackRand treatment.

each game, the active participant in the leader role chooses a number from a (13×13) payoff bimatrix, and the active participant in the follower role is informed about the leader’s choice. Being aware of the leader’s decision, the follower chooses his number from the same payoff bimatrix, and the leader is informed about the follower’s choice. Depending on the resulting combination of choices, both individuals, or both teams, receive monetary payoffs. In the following, I always speak of “choosing a quantity” and ”receiving a profit“ in order to keep the text as simple as possible.

In the payoff bimatrix, all possible combinations of quantity choices and the corresponding profits were shown (see Appendix A.4).⁸ Each active participant could choose a quantity from the set $\{3, 4, \dots, 15\}$. The profits were quoted in ECU (experimental currency unit). In order to make monetary incentives for active participants in TEAM identical to those of participants in LOADED or NEUTRAL, the worth of 1 ECU in TEAM was twice as high as in the other treatments. In LOADED and NEUTRAL, 20 ECU were worth 1 EUR. In TEAM, 10 ECU were worth 1 EUR. The payoff bimatrix was derived from the Stackelberg game described above. In order to ensure the uniqueness of the Stackelberg equilibrium, 14 of the 169 profit pairs were slightly manipulated by subtracting 1 ECU.

⁸The same payoff bimatrix was not only applied by Huck et al. (2001) and Müller and Tan (2011). It was also used by Huck et al. (2002) and Fonseca et al. (2006).

2.2. Hypotheses

In IO, firms are typically assumed to maximize their profits under rational expectations. In other words, it is assumed that firms behave according to a subgame-perfect Nash equilibrium prediction in the Stackelberg game: B maximizes its profit given A 's quantity choice, and A correctly anticipates B 's reaction to all possible quantity choices and maximizes its profit in the light of these anticipations. The subgame-perfect Nash equilibrium quantities $q_A^S = 12$ and $q_B^S = 6$ are given by $q_B(q_A^S) = 12 - \frac{q_A^S}{2}$ with $q_B(q_A) = \arg \max_{q_B} \pi_B(q_A, q_B)$ (best-response function) and $q_A^S = \arg \max_{q_A} \pi_A(q_A, q_B(q_A))$. This quantity combination is called Stackelberg outcome.

If A and B do not maximize their profits in the Stackelberg game, other outcomes can result from other kinds of behavior. If A prefers equality to inequality (fairness) and B maximizes its own profit (selfishness), the Cournot outcome ensues: $q_A^C = 8$ and $q_B^C = 8$. The joint profit is maximized if q_A and q_B add up to a total quantity of 12 (collusion). The symmetric joint profit-maximizing quantities are $q_A^J = 6$ and $q_B^J = 6$ (collusive outcome). The collusive outcome results if B is fair and A is either selfish or fair and interested in profit. An overview of the predicted outcomes and associated profits is provided in Table 2.

	Stackelberg	Cournot	Collusion
Quantities	$q_A^S = 12, q_B^S = 6$	$q_A^C = q_B^C = 8$	$q_A^J = q_B^J = 6$
Total quantity	$q^S = 18$	$q^C = 16$	$q^J = 12$
Profits	$\pi_A^S = 72, \pi_B^S = 36$	$\pi_A^C = \pi_B^C = 64$	$\pi_A^J = \pi_B^J = 72$
Total profit	$\pi^S = 108$	$\pi^C = 128$	$\pi^J = 144$

Table 2: Predicted outcomes and associated profits.

In StackRand, Huck et al. (2001) find that the Stackelberg prediction fails: On average, leaders choose a lower quantity, and followers choose a higher quantity than predicted. Leaders and followers choose the Cournot quantity in 17 and 28 percent of all choices. Lau and Leung (2010) show that the observed deviations from the Stackelberg prediction

can be explained by Fehr and Schmidt’s (1999) model of inequality aversion. In their parsimonious version of Fehr and Schmidt’s model, there are two kinds of followers: those who are selfish and those who have other-regarding preferences. All followers with other-regarding preferences have the same disadvantageous inequality aversion parameter and the same advantageous inequality aversion parameter. The proportion of followers with other-regarding preferences and the inequality aversion parameters are estimated using a maximum likelihood approach. About 40 percent of the followers are found to be averse to disadvantageous inequality, and the advantageous inequality aversion parameter is not statistically different from zero.⁹ These findings motivate

Hypothesis 1. (a) Individuals playing the Stackelberg game will not behave according to the Stackelberg prediction: leaders will choose lower quantities, and followers will choose higher quantities. (b) Leaders’ choices and followers’ responses in LOADED will be in line with Huck et al.’s (2001) findings.

On the basis of an ultimatum game, Hoffman et al. (1994) analyze the effect of an IO framing. In the treatment group, participants are called “seller” and “buyer”, and the proposer in the seller role “chooses the selling PRICE”. In the control group, participants are neutrally instructed, and the proposer makes a “proposal”. Hoffman et al. find that proposers in the treatment group offer less than proposers in the control group. Responders’ rejection rates do not differ significantly. Thus, proposers’ actions are more in line with the assumption of profit maximization. These results lead to

Hypothesis 2. (a) Framing individual decision making as organizational decision making will matter. (b) Leaders in LOADED will choose higher quantities than leaders in NEUTRAL. Compared to followers’ responses in NEUTRAL, followers’ responses in LOADED will be closer to their best responses.

⁹Lau and Leung do not analyze leader behavior.

Hence, the neoclassical theory of the firm as a theory of all firms is expected to be falsified again: single-person firms do not maximize their profits if the market structure is given by a Stackelberg game of duopolistic quantity competition with homogeneous products. A context is supposed to matter: The IO framing is expected to direct individual behavior towards profit maximization.

In TEAM, a multi-person firm is brought into the laboratory using a parsimonious version of Alchian and Demsetz's (1972) contractual view of the firm. For Alchian and Demsetz (p. 783), a firm is a contractual structure "with (a) joint input production, (b) several input owners, (c) one party who is common to all the contracts of the joint inputs, (d) who has rights to renegotiate any input's contract independently of contracts with other input owners, (e) who holds the residual claim, and (f) who has the right to sell his central contractual residual status." Moreover, (g) individuals within the firm maximize (expected) utilities on the basis of utility functions increasing in income and leisure, and (h) the firm is organized in a way that individual utility maximization within the firm is transformed into profit maximization on the market.

Item (g) is interpreted as a restricted version of the *homo oeconomicus* model: While it is known from many experiments that individuals are not rational, egoistic, and materialistic, the hypothesis here is that they will behave in that way in a specific context. Specifically, if leisure is maintained, firm members maximize their income. In TEAM, requirements (a) to (f), and (h) are met, and the duration of the experiment is independent of the actions taken by the firm members. The non-decision maker can be interpreted as an owner, that is, the residual claimant. The decision maker can be seen as a manager, who is motivated by an incentive-compatible contract granting him half of the firm's profit. Hence, teams are expected to maximize their profits because decision makers are predicted to maximize their incomes.

These considerations motivate

Hypothesis 3. (a) Leaders in TEAM will choose higher quantities than leaders in (a1) LOADED or (a2) NEUTRAL. Compared to followers' responses in (a3) LOADED or (a4) NEUTRAL, followers' responses in TEAM will be closer to their best responses. (b) Followers' response functions in TEAM will be closer to the best-response function than those in (b1) LOADED or (b2) NEUTRAL.

Hence, compared to framing individual decision making as organizational decision making, an organizational structure is expected to be more effective in generating profit-maximizing behavior.

2.3. Procedures

The experiment was conducted at Justus Liebig University Giessen in May and June 2011. Overall, 120 students from various fields of study, mostly from business administration, economics, and law, participated in four sessions. Participants were randomly recruited from a pool of potential participants. Each participant took part in only one session. Sessions consisted of ten rounds and lasted between 80 and 95 minutes (including the time to read the instructions). At the end of each session, two out of the ten rounds were randomly chosen to be rewarded. Participants' average earnings were €14.72 (including a fixed amount of €9.00).¹⁰ Rewards were paid out in private.

The experiment was run in lecture halls with pen and paper. At the beginning of the LOADED or each NEUTRAL session, participants were randomly assigned to be either a leader or a follower. This assignment remained fixed throughout the entire session. At the beginning of the TEAM session, participants were randomly assigned to be either a member of a team in the leader role (active or passive) or a member of a team in the follower role (active or passive). This assignment and the composition of the teams

¹⁰Without the flat amount, participants could have made losses in the experiment.

remained fixed throughout the entire session, too. Leaders and followers were seated in separate lecture halls. In order to prevent communication among participants located in the same lecture hall, they were seated with sufficient space between them. After having read the instructions, participants were allowed to ask questions in private. In each round, leaders and followers were randomly matched.¹¹ Hence, each active participant played ten (one-shot) Stackelberg games. At the end of all sessions, participants were asked to answer a questionnaire about their choices and the comprehensibility of the instructions.

Before the first round was started, participants were asked to answer a control question in order to make sure that all participants fully understood the payoff bimatrix. The answers were checked immediately (one follower in LOADED and one leader in TEAM answered the question incorrectly), and the question was answered in public. Nevertheless, in the questionnaire or when the rewards were paid out, three participants (two followers in LOADED and one leader in NEUTRAL) reported that they had had problems with the payoff bimatrix. The data from these three participants are excluded from the analyses in the following section.

In all treatments, leaders and followers were labeled A and B. In TEAM, decision makers and non-decision makers were additionally labeled D and N.¹² In LOADED and NEUTRAL, the leaders received a sheet of writing paper on which they noted their identifiers and their quantity choices at the beginning of each round. The sheets were then passed on to the followers. The followers also noted their quantity choices, and the sheets were passed back to the leaders. At the end of a round, each participant knew (i) his choice, (ii) the other participant's choice, (iii) his profit, and (iv) the other participant's profit and wrote down information (i) to (iv) on a sheet of reporting paper. With that, a

¹¹In all treatments, Kamecke's (1997) rotation random-matching protocol was applied.

¹²Actually, they were labeled E and N, because "Entscheider" is German for "decision maker".

round was finished (see also Huck et al., 2001, p. 753). In TEAM, each decision maker additionally reported his choice and the other team’s choice to his team member after each round. A team’s profit was equally shared between its two team members.

3. Experimental results

At a first glance, there is not much difference between individual behavior and team behavior. An overview of the mean quantities and the mean profits for all treatments, including Huck et al.’s (2001) StackRand treatment, is provided in Table 3; standard deviations are given in parentheses. Mean leader quantities are clearly lower and mean follower quantities are clearly higher than the Stackelberg quantities.

	Stackelberg	StackRand	LOADED	NEUTRAL	TEAM
Quantities A, B	12, 6 (0, 0)	10.19, 8.32 (2.45, 2.07)	9.79, 7.48 (2.06, 1.47)	9.22, 7.69 (2.21, 1.60)	9.82, 7.59 (2.16, 1.36)
Total quantity	18 (0)	18.51 (2.86)	17.27 (2.10)	16.91 (2.27)	17.41 (1.86)
Profits A, B	72, 36 (0, 0)	51.03, 42.43 (24.99, 25.89)	62.61, 49.18 (15.75, 17.26)	61.64, 53.14 (14.76, 18.65)	61.60, 49.73 (13.46, 17.11)
Total profit	108 (0)	93.46 (45.58)	111.79 (26.36)	114.78 (27.89)	111.33 (24.42)

Table 3: Mean quantities and mean profits; standard deviations in parentheses.

As explained above, other outcomes than the Stackelberg outcome can result from other-regarding preferences. The Cournot outcome ensues if leaders are fair and followers are selfish. The collusive outcome results if followers are fair and leaders are either selfish or fair and interested in profit. For all treatments, percentage frequencies of Stackelberg, Cournot, and collusive outcomes are shown in Table 4. The Cournot outcome is most frequently observed in each treatment. The Stackelberg outcome is clearly more frequent in TEAM than in the other treatments. Stackelberg outcomes occur with nearly the same frequency in StackRand, LOADED, and NEUTRAL, but Cournot outcomes are

extremely rare in StackRand compared to LOADED and NEUTRAL.

	Stackelberg $q_A^S = 12, q_B^S = 6$	Cournot $q_A^C = q_B^C = 8$	Collusion $q_A^J = q_B^J = 6$
StackRand	6.36	8.18	0.00
LOADED	10.83	19.79	1.67
NEUTRAL	8.57	27.62	1.90
TEAM	16.67	17.50	0.00

Table 4: Percentage frequencies of Stackelberg, Cournot, and collusive outcomes.

According to Hypothesis 1 (a), individuals playing the Stackelberg game do not behave according to the Stackelberg prediction: leaders choose lower quantities, and followers choose higher quantities. According to Hypothesis 1 (b), leaders' choices and followers' responses in LOADED are in line with Huck et al.'s (2001) findings.

Hypothesis 1 (a) is corroborated, because neither leaders nor followers in LOADED or NEUTRAL behave according to the Stackelberg prediction. In LOADED and NEUTRAL, leader quantities are lower and follower quantities are higher on average. These results are similar to Huck et al.'s findings. However, in StackRand, leaders' choices seem to be more in accordance with the Stackelberg prediction than in LOADED and NEUTRAL. On the other hand, followers' responses appear to be farther away from the Stackelberg prediction: on average, followers react with higher instead of lower quantities.

Using a two-sided Mann-Whitney U (MWU) test, I find that leaders' mean quantities in StackRand do not significantly differ from those in LOADED ($p = 0.363$): in the control treatment, leaders behave in the same way as in StackRand. Analyzing followers' choices is more complicated, because a follower's action is a response to a leader's choice. If followers use the same strategy in two treatments while leaders behave differently, different follower quantities will be observed. Therefore, the absolute value of the difference between the actual response of a follower, q_B , and the best response, $q_B(q_A)$,

is considered (see also Subsection 2.1). This deviation is called a follower’s “adjusted quantity”, formally, $q_B^a = |q_B - q_B(q_A)|$. Thus, $q_B^a = 0$ if B behaves as a profit maximizer. Percentage frequencies of the adjusted follower quantities are reported in Table 5 for all treatments.

	0	1	2	3	4	5	6	7	8	9	10
StackRand	50.91	16.82	8.64	5.45	7.27	6.82	0.91	1.36	0.45	0.91	0.45
LOADED	59.17	20.00	10.83	5.83	3.33	0.00	0.00	0.00	0.83	0.00	0.00
NEUTRAL	64.29	15.71	10.48	5.24	1.43	0.95	0.48	0.48	0.48	0.48	0.00
TEAM	76.67	12.50	3.33	1.67	3.33	2.50	0.00	0.00	0.00	0.00	0.00

Table 5: Percentage frequencies of adjusted follower quantities.

What is surprising is that only about half of the followers’ responses in StackRand are best responses. In all other treatments, best responses are more frequent, and the distributions of the adjusted follower quantities are not spread as widely as in StackRand. Followers’ mean adjusted quantities in StackRand significantly differ from those in LOADED (two-sided *MWU* test: $p = 0.008$): in the control treatment, followers do not behave in the same way as in StackRand. This falsifies Hypothesis 1 (b).

Participants’ actions in LOADED are not completely in line with participants’ actions in StackRand. However, they are similar. One reason for the difference in observed behavior might be a lack of control in Huck et al.’s experiment: No post-experimental questionnaire was used to find out whether participants had had problems with, for example, the payoff bimatrix. A closer look at the raw data reveals that some leaders and followers choose strange quantities. For example, there are leaders who decide for 3 or 4, and there are followers who react with a quantity which is so high that they make losses. Of course, these decisions can also be intentional, but this is hard to believe. Therefore, StackRand is no longer taken into consideration.

According to Hypothesis 2 (a), framing individual decision making as organizational decision making matters. According to Hypothesis 2 (b), leaders in LOADED choose

higher quantities than leaders in NEUTRAL. Compared to followers' responses in NEUTRAL, followers' responses in LOADED are closer to their best responses, that is, the IO framing is expected to direct individual behavior towards profit maximization.

An overview of the mean quantities per round is provided in Table 6 for all treatments; standard deviations are given in parentheses. In eight (out of the ten) rounds, mean leader quantities in LOADED are higher than in NEUTRAL. In only four rounds, followers' mean adjusted quantities in LOADED are lower than in NEUTRAL. Thus, framing individual decision making as organizational decision making seems to matter.

	LOADED			NEUTRAL			TEAM		
	q_A	q_B	q_B^a	q_A	q_B	q_B^a	q_A	q_B	q_B^a
Round 1	10.25 (2.49)	6.92 (1.73)	0.50 (1.00)	8.43 (2.42)	7.38 (1.50)	0.71 (0.85)	9.25 (1.71)	8.67 (1.78)	1.25 (1.96)
Round 2	9.83 (2.52)	7.50 (1.24)	1.00 (1.48)	9.52 (2.06)	7.86 (1.80)	0.67 (1.43)	10.08 (1.71)	8.00 (1.78)	0.92 (1.73)
Round 3	9.75 (2.30)	8.08 (2.15)	1.42 (2.27)	9.81 (2.27)	7.10 (1.55)	1.00 (1.41)	9.33 (2.10)	8.17 (1.40)	0.75 (1.48)
Round 4	10.50 (1.83)	7.50 (1.31)	0.92 (1.24)	9.29 (2.49)	7.52 (1.25)	0.81 (1.36)	9.50 (2.15)	7.58 (1.08)	0.25 (0.45)
Round 5	9.92 (1.56)	7.92 (1.00)	0.75 (1.29)	9.19 (1.91)	8.19 (1.94)	1.05 (1.96)	10.42 (2.07)	7.25 (1.36)	0.58 (0.90)
Round 6	9.42 (2.02)	7.58 (1.16)	0.25 (0.62)	10.00 (2.00)	7.76 (1.97)	0.81 (2.02)	10.25 (2.30)	7.25 (0.97)	0.33 (1.15)
Round 7	9.58 (1.93)	7.50 (1.24)	0.83 (1.03)	9.10 (1.95)	7.90 (1.18)	0.57 (1.03)	9.92 (1.83)	7.08 (0.90)	0.00 (0.00)
Round 8	9.58 (2.19)	7.25 (1.66)	0.83 (1.19)	8.90 (2.14)	7.62 (1.43)	0.67 (1.15)	9.00 (2.09)	7.58 (0.79)	0.17 (0.58)
Round 9	9.67 (2.31)	7.17 (1.59)	0.75 (1.06)	9.38 (2.52)	7.71 (1.93)	0.67 (1.56)	10.25 (2.30)	7.17 (1.19)	0.33 (0.65)
Round 10	9.42 (1.78)	7.42 (1.44)	0.67 (0.98)	8.57 (2.23)	7.86 (1.24)	0.76 (1.26)	10.17 (2.29)	7.17 (1.40)	0.42 (0.90)

Table 6: Mean quantities per round; standard deviations in parentheses.

However, in contrast to leaders' choices, followers' responses appear to be inconsistent with Hypothesis 2 (b): In both treatments, follower behavior looks very similar. This is interesting for two reasons. If leaders' motivations were identical in both treatments, leaders' actions could be explained by varying expectations. If, in addition, followers'

motivations were identical in both treatments, leaders' expectations would mistakenly differ. Alternatively, leaders' actions could be explained by varying motivations: Choosing a higher quantity could be interpreted as being more selfish.

Leaders' mean quantities in NEUTRAL are not significantly different from those in LOADED ($p = 0.157$) if a two-sided MWU test is applied. Using a one-sided MWU test, the difference is significant ($p = 0.078$): Leaders' mean quantities in LOADED are higher than in NEUTRAL. Followers' mean adjusted quantities in NEUTRAL do not significantly differ from those in LOADED (two-sided MWU test: $p = 0.549$). Hence, Hypothesis 2 (b) has been falsified. Followers' actual responses are not closer to their best responses: They are not affected by the IO framing. Hypothesis 2 (a) is corroborated, because there is a significant difference in leaders' choices. On the basis of the present data, it cannot be excluded that the difference is only due to differences in expectations about follower behavior.

According to Hypothesis 3 (a), leaders in TEAM choose higher quantities than leaders in (a1) LOADED or (a2) NEUTRAL. Compared to followers' responses in (a3) LOADED or (a4) NEUTRAL, followers' responses in TEAM are closer to their best responses. According to Hypothesis 3 (b), followers' response functions in TEAM are closer to the best-response function than those in (b1) LOADED or (b2) NEUTRAL. That is, the implementation of an organizational structure is expected to be more effective in generating profit-maximizing behavior than the application of an IO framing.

In six rounds, mean leader quantities in TEAM are higher than in LOADED (see Table 6). In even nine rounds, mean leader quantities in TEAM are higher than in NEUTRAL. In eight rounds, followers' mean adjusted quantities in TEAM are lower than in both LOADED and NEUTRAL. Therefore, team choices appear to be closer to the Stackelberg prediction than individual choices. However, using two-sided MWU tests, neither leaders' mean quantities in LOADED nor leaders' mean quantities in NEUTRAL

significantly differ from those in TEAM ($p = 0.918$ and $p = 0.203$). Using a one-sided MWU test, the difference between TEAM and NEUTRAL is significant ($p = 0.101$): Leaders' mean quantities in TEAM are higher than in NEUTRAL. Hence, Hypothesis 3 (a1) has been falsified, and Hypothesis 3 (a2) is corroborated. Nearly the same is true for followers' mean adjusted quantities. Using two-sided MWU tests, neither followers' mean adjusted quantities in LOADED nor followers' mean adjusted quantities in NEUTRAL are significantly different from those in TEAM ($p = 0.191$ and $p = 0.573$). Using a one-sided MWU test, the difference between TEAM and LOADED is significant ($p = 0.096$): followers' mean adjusted quantities in TEAM are lower than in NEUTRAL. Thus, Hypothesis 3 (a3) is corroborated, and Hypothesis 3 (a4) has been falsified.

On the one hand, the weak MWU test results are surprising because of the clear summary statistics results, on the other hand, treating each active participant as one observation by using his mean quantity choice of all rounds is an extremely conservative method to overcome the problem of repeated measurement. A less extreme way to deal with this problem is to apply a dummy variable regression in order to test for treatment effects. TEAM is the treatment group. LOADED and NEUTRAL serve as control groups. Either leaders' quantities or followers' adjusted quantities are regressed on a constant and a binary variable (*team*), which is one for the observations belonging to the treatment group.

Formally, the model can be written as

$$quantity = \beta_0 + \beta_1 team + \sum_{p=1}^P \gamma_p participant_p + \sum_{r=1}^{10} \delta_r round_r + u, \quad (1)$$

$$\sum_{p=1}^P \gamma_p = 0, \quad (2)$$

$$\sum_{r=1}^{10} \delta_r = 0, \quad (3)$$

where *quantity* are either leaders' quantities in TEAM and LOADED or followers' adjusted quantities in TEAM and NEUTRAL. To control for both participant effects and round influences, a dummy variable for each participant ($participant_p$ for participants $p = 1, 2, \dots, P$) and for each round ($round_r$ for rounds $r = 1, 2, \dots, 10$) is included. The error term (u) contains unobserved factors affecting *quantity*.¹³

Following Suits (1984), the sum of the coefficients of each set of dummy variables is constrained to zero (see also Königstein, 2000). The intercept parameter β_0 can, therefore, be interpreted as the mean quantity in LOADED or NEUTRAL, respectively. The slope parameter β_1 is the difference in mean quantities between either LOADED and TEAM or NEUTRAL and TEAM. If β_1 is estimated to be significantly different from zero, a treatment effect is present. The estimates of β_0 and β_1 , $\hat{\beta}_0$ and $\hat{\beta}_1$, are reported in Table 7; standard errors and p -values for the two-sided standard t tests are shown in parentheses. For q_A , the p -value for the two-sided t test $H_0 : \beta_0 = 12$ is also given.

		$\hat{\beta}_0$	$\hat{\beta}_1$
LOADED	q_A	9.75	0.06
		(0.16, $p = 0.000$; $p_{12} = 0.000$)	(0.22, $p = 0.771$)
	q_B^a	0.79	-0.29
		(0.10, $p = 0.000$)	(0.14, $p = 0.038$)
NEUTRAL	q_A	9.22	0.60
		(0.13, $p = 0.000$; $p_{12} = 0.000$)	(0.21, $p = 0.005$)
	q_B^a	0.76	-0.26
		(0.07, $p = 0.000$)	(0.12, $p = 0.032$)

Table 7: Treatment effects.

Except for the difference in leaders' mean quantities between LOADED and TEAM ($p = 0.771$), treatment effects are significant and have the expected sign, that is, Hypothesis 3 (a1) has been falsified, and Hypotheses 3 (a2, a3, a4) are corroborated. Followers in TEAM are more selfish than in LOADED or NEUTRAL. Leaders in TEAM and

¹³Because each participant is observed in each round, the error terms might be correlated. Here, serial correlation is likely to result from an unobserved effect, namely, different (time-constant) preferences. Serial correlation can also result from learning. In order to avoid these two problems, it is controlled for participant effects and round influences.

LOADED behave in the same way: They behave more like profit maximizers than leaders in NEUTRAL. Thus, teams' quantity choices are more in line with the assumption of profit maximization than individuals' choices, or, to put it the other way round, the implementation of an organizational structure is more effective in generating profit-maximizing behavior than the application of an IO framing. Nevertheless, teams do not completely behave as profit maximizers.¹⁴

In order to find out whether the observed behavior can be explained by other-regarding preferences, followers' responses are analyzed in more detail by estimating their response functions, $q_B = \beta_0 + \beta_1 q_A$, for all treatments. If followers behaved as profit maximizers, the estimated response function would be $\widehat{q}_B = 12.09 - 0.49q_A$ in each treatment. Again, a constrained dummy variable regression is used, including intercept and slope dummy variables for participants and rounds. Because the sum of the coefficients of each set of dummy variables is restricted to zero, the intercept parameter β_0 and the slope parameter β_1 represent means. The estimates $\hat{\beta}_0$ and $\hat{\beta}_1$ of β_0 and β_1 are shown in Table 8; standard errors and p -values for the two-sided standard t tests are given in parentheses. In addition, p -values for the two-sided t tests $H_0 : \beta_0 = 12.09$ and $H_0 : \beta_1 = 0.49$ are also shown.

	$\hat{\beta}_0$	$\hat{\beta}_1$
LOADED	10.45	-0.30
	(0.49, $p = 0.000$; $p_{12.09} = 0.001$)	(0.05, $p = 0.000$; $p_{0.49} = 0.000$)
NEUTRAL	9.61	-0.21
	(0.40, $p = 0.000$; $p_{12.09} = 0.000$)	(0.04, $p = 0.000$; $p_{0.49} = 0.000$)
TEAM	10.84	-0.33
	(0.48, $p = 0.000$; $p_{12.09} = 0.012$)	(0.05, $p = 0.000$; $p_{0.49} = 0.000$)

Table 8: Estimated response functions.

In TEAM, the estimated intercept is 10.84 on average, and the estimated intercepts are lower in LOADED and NEUTRAL. In each treatment, the estimated slope of the

¹⁴Participant effects and round influences are estimated to be present for a small number of participants and rounds.

response function is negative on average, and the absolute value of the estimated slope is largest in TEAM. Therefore, followers' response functions in TEAM appear to be closer to the best-response function than those in the other treatments: Hypotheses 3 (b1, b2) are corroborated.

If followers' response functions looked like in Table 8, leaders' profit-maximizing quantities would be 9.68 in LOADED, 9.11 in NEUTRAL, and 9.82 in TEAM. These quantities are very close to leaders' mean quantities: 9.79 in LOADED, 9.22 in NEUTRAL, and 9.82 in TEAM (see Table 3). In TEAM, leaders' mean quantity even corresponds with the profit-maximizing quantity. Thus, here too, the implementation of an organizational structure is found to be more effective in generating profit-maximizing behavior than the application of an IO framing: teams' quantity choices are more in line with the assumption of profit maximization than individuals' choices. However, intercept and slope of the estimated response function in each treatment differ from the parameters of the best-response function significantly, and leaders' mean quantities in LOADED and NEUTRAL are higher than the profit-maximizing quantities.

Given the estimated response function in TEAM, leaders maximize their profits on average. In LOADED and NEUTRAL, leaders' mean behavior is harder to interpret: An interpretation relies on leaders' expectations. If followers' estimated response functions were correct and expected, leaders' mean quantities choices, compared to the profit-maximizing choices, would lead to less profits for both interacting individuals and would increase inequality in profits. Leaders would be inequality loving, and the loss of profit could be interpreted as their willingness to pay for the increase in (positive) inequality. This explanation is not convincing in view of Fehr and Schmidt's (1999) or Bolton and Ockenfels' (2000) experimental results on inequality aversion: Typically, participants are found to be inequality averse instead of inequality loving. Hence, it is more likely that leaders expectations are wrong or followers' actual response functions are different

in LOADED and NEUTRAL.

However, mean follower behavior is in line with Fehr and Schmidt's and Bolton and Ockenfels's findings. In LOADED, NEUTRAL, and TEAM, estimated response functions intersect the best-response function at $q_A = 8.63$, $q_A = 8.86$, and $q_A = 7.81$, implying that, from the leader quantity of 9 upwards in LOADED and NEUTRAL or 8 upwards in TEAM, followers choose more than the profit-maximizing quantity on average. By doing so, negative inequality decreases. From the leader quantity of 7 in TEAM or 8 in LOADED and NEUTRAL, followers choose less than predicted. This mean behavior also decreases inequality, namely, positive inequality. Because of the steepest slope of the response function in TEAM, inequality aversion is weaker in TEAM than in LOADED or NEUTRAL.

On the whole, it can be said that teams are less inequality averse than individuals. If individual decision making is framed as organizational decision making, individual behavior is also directed towards profit maximization, but the implementation of an organizational structure is more effective. Therefore, an individual under an IO framing is less a profit maximizer than a two-member team organized according to Alchian and Demsetz's contractual model of the firm.

4. Conclusion

In this paper, a Stackelberg experiment is considered in order to answer the question whether framing individual decision making as organizational decision making or implementing an organizational structure is more effective in generating the behavior which is mostly assumed in IO, namely, profit-maximizing. The experimental market structure is derived from Huck et al.'s (2001) Stackelberg game of duopolistic quantity compe-

tition with homogeneous products. Overall, there are three treatments: LOADED, NEUTRAL, and TEAM. In LOADED and NEUTRAL, firms are each represented by an individual. In LOADED, participants are called “firms” competing in “quantities” for “profits”. In NEUTRAL, participants are neutrally instructed: Individuals choose “numbers” and receive “payments”. Individuals in LOADED and NEUTRAL can be seen as single-person firms. In TEAM, participants are also neutrally instructed. Teams each consist of a decision maker and a non-decision maker. The non-decision maker is completely passive. The decision maker is active and gets half of the team’s monetary payoff. Teams can be viewed as multi-person firms organized according to a parsimonious version Alchian and Demsetz’s (1972) contractual model of the firm.

I find that neither individuals nor teams are strict profit maximizers in the Stackelberg game. However, if individual decision making is framed as organizational decision making, leaders behave more like profit maximizers: Leaders in LOADED choose higher quantities than leaders in NEUTRAL. Followers motivations are not affected by such a framing: Followers in LOADED behave in the same way as followers in NEUTRAL. Teams are found to come closer to profit maximization: Leaders in TEAM behave in the same way as leaders in LOADED, but followers in TEAM behave more like profit maximizers than followers in LOADED. Therefore, the organizational structure is most effective in generating profit-maximizing behavior. Teams are less inequality averse than individuals. Followers’ response functions in TEAM appear to be closer to the best-response function than those in the other treatments. Given the estimated response function, leaders’ mean quantity in TEAM is a profit-maximizing choice.

What does that mean for IO? First, the theory of the profit-maximizing firm does not apply to all firms. Neither single-person firms nor multi-person firms generally maximize their profits. However, it does not follow that all kinds of multi-person firms do not maximize their profits. In the experiment, two-person firms behave more like profit

maximizers than single-person firms. That is surprising in view of the simplicity of the organizational structure. Each firm only consists of a passive and an active firm member. There is no communication between them, and there is no such thing as a team spirit: The TEAM treatment is done without loaded instructions or team-building activities.

Whether the behavioral difference between single-person and multi-person firms is triggered by other-regarding preferences between firm members or whether the contract of employment makes people more selfish is still an open question. In order to answer it, other employment contracts could be implemented, or firm sizes could be increased. Because firms are typically assumed to be large in oligopoly theory, increasing firm sizes might be the preferred choice. However, because of the rising costs in the laboratory and the great significance of team production problems in the field, a variation of the organizational structure and an implementation of more complex employment contracts might be a promising approach.

A. Appendix

A.1. Translated instructions: LOADED

Welcome to our experiment!

Please read the instructions carefully. Do not talk to your neighbors during the entire experiment. Raise your hand if you have questions. We will come around and answer your questions.

Your participation in the experiment will be rewarded. Depending on your behavior and the behavior of other participants you are matched with, you receive lower or higher monetary rewards in EUR.

You represent a firm selling the same product as another firm on a market. Both firms each make one decision. That is, each firm chooses the quantity it wants to sell on the market. The resulting combination of quantities is associated with a profit for each firm.

The profits associated with each combination of quantities are shown in the attached table.

The quantities which can be chosen by firm A are displayed in the head of each row. The quantities which can be chosen by firm B are displayed in the head of each column. The profits associated with a combination of quantities are shown in the corresponding cell. The entry on the left side of the vertical bar corresponds to firm A's profit. The entry on the right side of the vertical bar corresponds to firm B's profit.

The profits are quoted in ECU (experimental currency unit). The exchange rate between ECU and EUR is 1/20. That is, 20 ECU are exchanged for 1 EUR.

How are the decisions made?

Please take a look at your identifier. If it begins with an A, you represent a firm A. If it begins with a B, you represent a firm B.

Firm A is the first to make a decision. That is, firm A chooses its quantity (picks a row), and firm B will be informed about firm A's choice. Being aware of firm A's decision, firm B chooses its quantity (picks a column), and firm A will be informed about firm B's choice. With that, a round is finished. That is, at the end of a round, each firm knows (i) its quantity, (ii) the other firm's quantity, (iii) its profit, and (iv) the other firm's profit.

Decisions are communicated in writing. Each firm A gets a sheet of writing paper at the beginning of each round.

Please write down information (i) to (iv) on the attached sheet of reporting paper.

In total, there are ten rounds. You do not know the participant you interact with. In each round, you will be matched with a different participant.

Anonymity among participants and towards experimenters is preserved. Your decisions can only be traced back to your identifier. Your personal data will not be associated with your identifier.

At the end of the experiment, two out of the ten rounds will be randomly chosen to be rewarded. The sum of your profits from these two rounds determines the variable part of your monetary reward in EUR. It can be positive, zero, or negative. In addition, you will receive a fixed amount of 9 EUR. Your monetary reward will be paid out in private. That is, the other participants will not learn about the amount of your monetary reward.

A.2. Translated instructions: NEUTRAL

Welcome to our experiment!

Please read the instructions carefully. Do not talk to your neighbors during the entire experiment. Raise your hand if you have questions. We will come around and answer your questions.

Your participation in the experiment will be rewarded. Depending on your behavior and the behavior of other participants you are matched with, you receive lower or higher monetary rewards in EUR.

You interact with another participant. Both participants each make one decision. That is, each participant chooses the number he wants to use. The resulting combination of numbers is associated with a payment for each participant.

The payments associated with each combination of numbers are shown in the attached table.

The numbers which can be chosen by participant A are displayed in the head of each row. The numbers which can be chosen by participant B are displayed in the head of each column. The payments associated with a combination of numbers are shown in the corresponding cell. The entry on the left side of the vertical bar corresponds to participant A's payment. The entry on the right side of the vertical bar corresponds to participant B's payment.

The payments are quoted in ECU (experimental currency unit). The exchange rate between ECU and EUR is 1/20. That is, 20 ECU are exchanged for 1 EUR.

How are the decisions made?

Please take a look at your identifier. If it begins with an A, you are a participant A. If it begins with a B, you are a participant B.

Participant A is the first to make a decision. That is, participant A chooses his number (picks a row), and participant B will be informed about participant A's choice. Being aware of participant A's decision, participant B chooses his number (picks a column), and participant A will be informed about participant B's choice. With that, a round is finished. That is, at the end of a round, each participant knows (i) his number, (ii) the other participant's number, (iii) his payment, and (iv) the other participant's payment.

Decisions are communicated in writing. Each participant A gets a sheet of writing paper at the beginning of each round.

Please write down information (i) to (iv) on the attached sheet of reporting paper.

In total, there are ten rounds. You do not know the participant you interact with. In each round, you will be matched with a different participant.

Anonymity among participants and towards experimenters is preserved. Your decisions can only be traced back to your identifier. Your personal data will not be associated

with your identifier.

At the end of the experiment, two out of the ten rounds will be randomly chosen to be rewarded. The sum of your payments from these two rounds determines the variable part of your monetary reward in EUR. It can be positive, zero, or negative. In addition, you will receive a fixed amount of 9 EUR. Your monetary reward will be paid out in private. That is, the other participants will not learn about the amount of your monetary reward.

A.3. Translated instructions: TEAM

Welcome to our experiment!

Please read the instructions carefully. Do not talk to your neighbors during the entire experiment. Raise your hand if you have questions. We will come around and answer your questions.

Your participation in the experiment will be rewarded. Depending on your behavior and the behavior of other participants you are matched with, you receive lower or higher monetary rewards in EUR.

You and another participant (partner) are a team. As a team, you interact with another team. Both teams each make one decision. That is, each team chooses the number it wants to use. The resulting combination of numbers is associated with a payment for each team.

The payments associated with each combination of numbers are shown in the attached table.

The numbers which can be chosen by team A are displayed in the head of each row. The numbers which can be chosen by team B are displayed in the head of each column.

The payments associated with a combination of numbers are shown in the corresponding cell. The entry on the left side of the vertical bar corresponds to team A's payment. The entry on the right side of the vertical bar corresponds to team B's payment.

The payments are quoted in ECU (experimental currency unit). The exchange rate between ECU and EUR is 1/10. That is, 10 ECU are exchanged for 1 EUR.

How are the decisions made?

Please take a look at your identifier. If it begins with an A, you are a member of a team A. If it begins with a B, you are a member of a team B. Each team consists of a decision maker and a non-decision maker. If your identifier contains a D, you are a decision maker. If your identifier contains an N, you are a non-decision maker.

Decision maker A is the first to make a decision. That is, decision maker A chooses his number (picks a row), and decision maker B will be informed about decision maker A's choice. Being aware of decision maker A's decision, decision maker B chooses his number (picks a column), and decision maker A will be informed about decision maker B's choice. Finally, non-decision maker A and non-decision maker B will be informed about the choices. With that, a round is finished. That is, at the end of a round, each team knows (i) its number, (ii) the other team's number, (iii) its payment, (iv) and the other team's payment.

Decisions are communicated in writing. Each decision maker A gets a sheet of writing paper at the beginning of each round. All decision makers have reporting sheets.

Please write down information (i) to (iv) on the attached sheet of reporting paper. If you are a non-decision maker, please also note whether you are satisfied with your partner's decision and what your decision would have been.

In total, there are ten rounds. You do not know the team you interact with. In each

round, you will be matched with a different team. The composition of the teams does not change during the entire experiment.

Anonymity among participants and towards experimenters is preserved. Your decisions can only be traced back to your identifier. Your personal data will not be associated with your identifier.

At the end of the experiment, two out of the ten rounds will be randomly chosen to be rewarded. The sum of your payments from these two rounds determines the variable part of your team's monetary reward in EUR, which will be equally shared between you and your partner. It can be positive, zero, or negative. In addition, you will receive a fixed amount of 9 EUR. Your monetary reward will be paid out in private. That is, the other participants will not learn about the amount of your monetary reward.

Quantity/ Number	Firm B/Participant B													
	3	4	5	6	7	8	9	10	11	12	13	14	15	
Firm A/Participant A	3	54 54	51 68	48 80	45 90	42 98	39 104	36 108	33 109	30 110	27 108	24 104	21 98	18 90
	4	68 51	64 64	60 75	56 84	52 91	48 96	44 99	40 100	36 99	32 96	28 91	24 84	19 75
	5	80 48	75 60	70 70	65 78	60 84	55 88	50 89	45 90	40 88	35 84	29 78	25 70	20 60
	6	90 45	84 56	78 65	72 72	66 77	60 80	54 81	48 80	41 77	36 72	30 65	24 56	18 45
	7	98 42	91 52	84 60	77 66	70 70	63 72	55 71	49 70	42 66	35 60	28 52	21 42	14 30
	8	104 39	96 48	88 55	80 60	72 63	64 64	56 63	48 60	40 55	32 48	24 39	16 28	8 15
	9	108 36	99 44	89 50	81 54	71 55	63 56	54 54	45 50	36 44	27 36	18 26	9 14	0 0
	10	109 33	100 40	90 45	80 48	70 49	60 48	50 45	40 40	30 33	20 24	10 13	0 0	-10 -15
	11	110 30	99 36	88 40	77 41	66 42	55 40	44 36	33 30	22 22	11 12	0 0	-11 -14	-22 -30
	12	108 27	96 32	84 35	72 36	60 35	48 32	36 27	24 20	12 11	0 0	-12 -13	-24 -28	-36 -45
	13	104 24	91 28	78 29	65 30	52 28	39 24	26 18	13 10	0 0	-13 -12	-26 -26	-39 -42	-52 -60
	14	98 21	84 24	70 25	56 24	42 21	28 16	14 9	0 0	-14 -11	-28 -24	-42 -39	-56 -56	-70 -75
	15	90 18	75 19	60 20	45 18	30 14	15 8	0 0	-15 -10	-30 -22	-45 -36	-60 -52	-75 -70	-90 -90

A.4. Payoff bimatrix: LOADED/NEUTRAL and TEAM

References

- Albert, Max and Andreas Hildenbrand (2012): *Industrial Organization in the Laboratory*. MAGKS Joint Discussion Paper 05-2012, Department of Business Administration and Economics, Philipps University Marburg.
- Alchian, Armen A. and Harold Demsetz (1972): Production, Information Costs, and Economic Organization. *American Economic Review*, 62(5):777–795.
- Bolton, Gary E. and Axel Ockenfels (2000): ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review*, 90(1):166–193.
- Bornstein, Gary (2008): A Classification of Games by Player Type. In: *New Issues and Paradigms in Research on Social Dilemmas*, eds. Andreas Biel; Daniel Eek; Tommy Gärling; and Mathias Gustafsson, Springer, New York, chap. 3, pp. 27–42.
- Bornstein, Gary and Uri Gneezy (2002): Price Competition Between Teams. *Experimental Economics*, 5(1):29–38.
- Bornstein, Gary; Tamar Kugler; David V. Budescu; and Reinhard Selten (2008): Repeated Price Competition Between Individuals and Between Teams. *Journal of Economic Behavior & Organization*, 66(3–4):808–821.
- Bornstein, Gary and Ilan Yaniv (1998): Individual and Group Behavior in the Ultimatum Game: Are Groups More “Rational” Players? *Experimental Economics*, 1(1):101–108.
- Engel, Christoph (2010): The Behaviour of Corporate Actors: How Much Can We Learn From the Experimental Literature? *Journal of Institutional Economics*, 6(4):445–475.
- Fehr, Ernst and Klaus M. Schmidt (1999): A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics*, 114(3):817–868.

- Fonseca, Miguel A.; Steffen Huck; and Hans-Theo Normann (2005): Playing Cournot Although They Shouldn't: Endogenous Timing in Experimental Duopolies with Asymmetric Cost. *Economic Theory*, 25(3):669–677.
- Fonseca, Miguel A.; Wieland Müller; and Hans-Theo Normann (2006): Endogenous Timing in Duopoly: Experimental Evidence. *International Journal of Game Theory*, 34(3):443–456.
- Furubotn, Eirik G. and Rudolf Richter (2005): *Institutions and Economic Theory: The Contribution of the New Institutional Economics*. University of Michigan Press, Ann Arbor, 2nd edn.
- Hildenbrand, Andreas (2010): *Cournot or Stackelberg Competition? A Survey of Experimental Research*. MPRA Paper 26845, Munich Personal RePEc Archive, Ludwig-Maximilians University Munich.
- Hoffman, Elizabeth; Kevin McCabe; Keith Shachat; and Vernon Smith (1994): Preferences, Property Rights, and Anonymity in Bargaining Games. *Games and Economic Behavior*, 7(3):346–380.
- Holmstrom, Bengt R. and Jean Tirole (1989): The Theory of the Firm. In: *Handbook of Industrial Organization*, eds. Richard Schmalensee and Robert D. Willig, North-Holland, Amsterdam, vol. 1, chap. 2, pp. 61–133.
- Holt, Charles A. (1985): An Experimental Test of the Consistent-Conjectures Hypothesis. *American Economic Review*, 75(3):314–325.
- Huck, Steffen; Wieland Müller; and Hans-Theo Normann (2001): Stackelberg Beats Cournot: On Collusion and Efficiency in Experimental Markets. *Economic Journal*, 111(474):749–765.

- Huck, Steffen; Wieland Müller; and Hans-Theo Normann (2002): To Commit or Not to Commit: Endogenous Timing in Experimental Duopoly Markets. *Games and Economic Behavior*, 38(2):240–264.
- Huck, Steffen; Hans-Theo Normann; and Jörg Oechssler (2004): Two Are Few and Four Are Many: Number Effects in Experimental Oligopolies. *Journal of Economic Behavior & Organization*, 53(4):435–446.
- Huck, Steffen and Brian Wallace (2002): Reciprocal Strategies and Aspiration Levels in a Cournot-Stackelberg Experiment. *Economics Bulletin*, 3(3):1–7.
- Kamecke, Ulrich (1997): Matching Schemes that Efficiently Preserve the Best Reply Structure of a One Shot Game. *International Journal of Game Theory*, 26(3):409–417.
- Königstein, Manfred (2000): Measuring Treatment-Effects in Experimental Cross-Sectional Time Series. In: *Equity, Efficiency and Evolutionary Stability in Bargaining Games with Joint Production*, ed. Manfred Königstein, Springer, Berlin, chap. 2, pp. 33–43.
- Lau, Sau-Him Paul and Felix Leung (2010): Estimating a Parsimonious Model of Inequality Aversion in Stackelberg Duopoly Experiments. *Oxford Bulletin of Economics and Statistics*, 72(5):669–686.
- Müller, Wieland (2006): Allowing for Two Production Periods in the Cournot Duopoly: Experimental Evidence. *Journal of Economic Behavior & Organization*, 60(1):100–111.
- Müller, Wieland and Fangfang Tan (2011): *Who Acts More Like a Game Theorist? Group and Individual Play in a Sequential Market Game and the Effect of the Time Horizon*. Working Paper 1111, Department of Economics, University of Vienna.

- Nadiri, M. Ishaq (1982): Producers Theory. In: *Handbook of Mathematical Economics*, eds. Kenneth J. Arrow and Michael D. Intriligator, North-Holland, Amsterdam, vol. 2, chap. 10, pp. 431–490.
- Normann, Hans-Theo and Bradley Ruffle (2011): Introduction to the Special Issue on Experiments in Industrial Organization. *International Journal of Industrial Organization*, 29(1):1–3.
- Raab, Philippe and Burkhard Schipper (2009): Cournot Competition Between Teams: An Experimental Study. *Journal of Economic Behavior & Organization*, 72(2):691–702.
- Sauermann, Heinz and Reinhard Selten (1959): Ein Oligopolexperiment. *Zeitschrift für die gesamte Staatswissenschaft*, 115(3):427–471.
- Scherer, Frederic M. and David Ross (1990): *Industrial Market Structure and Economic Performance*. Houghton Mifflin, Boston, 3rd edn.
- Selten, Reinhard (1967a): Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. In: *Beiträge zur experimentellen Wirtschaftsforschung*, ed. Heinz Sauermann, Mohr, Tübingen, chap. 5, pp. 136–168.
- Selten, Reinhard (1967b): Ein Oligopolexperiment mit Preisvariation und Investition. In: *Beiträge zur experimentellen Wirtschaftsforschung*, ed. Heinz Sauermann, Mohr, Tübingen, chap. 4, pp. 103–135.
- Shapiro, Carl (1989): Theories of Oligopoly Behavior. In: *Handbook of Industrial Organization*, eds. Richard Schmalensee and Robert D. Willig, North-Holland, vol. 1, chap. 6, pp. 329–414.

Suits, Daniel B. (1984): Dummy Variables: Mechanics V. Interpretation. *Review of Economics and Statistics*, 66(1):177–180.

Tirole, Jean (1988): *The Theory of Industrial Organization*. MIT Press, Cambridge.