



Noisy leadership: An experimental approach

Werner Güth^a, Wieland Müller^b, Yossi Spiegel^{c,*}

^a *Max Planck Institute of Economics, Kahlaische Strasse 10, 07745 Jena, Germany*

^b *Department of Economics, PO Box 90153, 5000LE Tilburg, The Netherlands*

^c *Recanati Graduate School of Business Administration, Tel Aviv University, Ramat Aviv, Tel Aviv 69978, Israel*

Received 11 August 2002

Available online 13 June 2006

Abstract

We examine the strategic behavior of leaders and followers in sequential duopoly experiments with errors in communication: followers either perfectly observe the leaders' actions or else they observe nothing. Consistent with the theory, the leaders in our experiments enjoy a greater first-mover advantage when followers observe their actions with higher probability, albeit their advantage is weaker than the theory predicts and is only weakly increasing with the probability that their actions will be observed. Our results also show that (i) when informed, followers hardly ever underreact to the leaders' quantities but tend to overreact slightly, and (ii) when uninformed, followers try to predict leaders' quantities and react optimally. This suggests that followers view the symmetric Cournot outcome as "fair," and when informed, "punish" leaders who try to exploit their first-mover advantage. In turn, such punishments by overreactions induce leaders to behave more softly than the theory predicts.

© 2006 Elsevier Inc. All rights reserved.

JEL classification: C72; C91

Keywords: Commitment; Imperfect observability; Errors in communication; Sequential games; Oligopoly; Experiments

* Corresponding author. Fax: +972 3 640 7739.

E-mail addresses: gueth@mpiew-jena.mpg.de (W. Güth), w.mueller@uvt.nl (W. Müller), spiegel@post.tau.ac.il (Y. Spiegel).

URLs: <http://www.econ.mpg.de/esi/gueth/> (W. Güth), <http://center.uvt.nl/staff/muller/> (W. Müller), <http://www.tau.ac.il/~spiegel> (Y. Spiegel).

1. Introduction

The idea that first movers (leaders) may gain a strategic advantage by committing to a certain course of action is one of the most celebrated insights of non-cooperative game theory and is widely used in such diverse fields as macroeconomics, international trade, and industrial organization. This idea dates back at least to von Stackelberg (1934) and was popularized by Schelling (1960) who emphasized that in order to confer a strategic advantage, the leader's action must be reliably communicated to second movers (followers). In reality however, it is highly likely that the actions of leaders will only be imperfectly observed by followers. Using the terminology of van Damme and Hurkens (1997) one can distinguish at least two types of possible imperfections. First, there could be errors in the followers' perceptions about the action that the leader has taken. That is, followers may receive a noisy signal about the leader's action and may therefore erroneously believe that the leader has played some other action than the one that was actually played. Second, there could be errors in communication: followers may simply fail to observe the leader's action altogether. Since in specific applications either type of error is highly likely, it is clearly important to explore the implications of imperfect observability for the strategic behavior of leaders and followers and in particular for the value and robustness of commitment.

Given its importance for applied work, so far, surprisingly little research was done on noisy-leadership games. This research has focused almost exclusively on the errors in perceptions case under the special assumption that the noise structure has full support (with some probability, followers may observe any action from the leader's strategy set). Bagwell (1995) shows that in this case, the pure strategy equilibrium outcome jumps discretely from the Stackelberg outcome (the subgame perfect equilibrium of the sequential-move game) when there is no noise to the Cournot outcome (the Nash equilibrium of the simultaneous move game) when there is even the slightest amount of noise. Van Damme and Hurkens (1997) show that the same model admits a mixed strategy equilibrium that converges to the subgame perfect equilibrium as the noise vanishes. Moreover, this equilibrium is selected by a "reasonable" selection criterion. Maggi (1999) shows that when the leader has private information about his cost, an increase in the ratio between the noise in the signal and the noise in the leader's type shifts the equilibrium outcome smoothly from the Stackelberg to the Cournot outcome. Oechssler and Schlag (2000) analyze Bagwell's noisy leader's game with a wide variety of evolutionary and learning dynamics and find that almost all of them admit the Cournot equilibrium as a possible outcome, and often select it uniquely. Only the continuous best-response dynamic selects the Stackelberg outcome as the unique long-run outcome if the noise is small.

Experimental evidence by Huck and Müller (2000) and Müller (2001) in the context of a 2×2 game with a full-support noise structure reveals that followers seem to ignore small levels of noise and play a best-response against the observed leader's action even though with some probability this might be the "wrong" action. Leaders quickly learn to exploit this tendency and play the Stackelberg leader's quantity.¹ By contrast, with high levels of noise the Cournot equilibrium is unique and indeed play converges to this equilibrium. Morgan and Várdu (2004) conduct experiments on Várdu's (2004) model in which followers can perfectly observe the leaders' choices at a cost. They find that the value of commitment is almost completely preserved when the cost of observation is small, but is lost when it is large. The general conclusion from

¹ With small noise levels, the game admits beside the Cournot equilibrium also two mixed strategy equilibria, one of which converges to the Stackelberg equilibrium as noise vanishes. Play seems to converge to this equilibrium.

these experiments is that neither small (full support) noise nor small cost of observation eliminate the value of commitment though its value is smaller than that predicted by the theory.

In this paper we report the results of a series of experiments intended to study the strategic behavior of leaders and followers in settings that involve errors in communication. The followers in our experiments received an all-or-nothing signal about the leader's choice so they either perfectly observed the leader's choice, or else they observed nothing.² This signal technology is widely used in the literature (e.g., Laffont and Tirole, 1993, and Rubinstein, 1989). From a modeling point of view, the errors in communication case captures situations in which followers are fully aware of whether they have observed the leader's true action or not. By contrast, the case of errors in perceptions captures situations in which followers are always in doubt as to whether they have observed the leader's true action or some other action.³ Unlike the errors in perceptions case with full-support noise structure, the equilibrium outcome under errors in communication shifts continuously with the probability that followers will observe the leaders' actions (i.e., with the level of noise) from the Cournot outcome to the Stackelberg outcome (Chakravorti and Spiegel, 1993).⁴ Moreover, as intuition suggests, first-movers enjoy a greater strategic advantage when their actions are observed with a higher probability. Therefore, our setting has the advantage of allowing for partial commitment (i.e., commitment is no longer a zero or one variable). From an experimental point of view, this setting has the additional advantage of providing us with a continuous measure of the power of commitment and allowing us to rely in all cases (including Cournot and Stackelberg) on the same verbal instructions (the only change is in the probability with which the leader's action is observed). This feature is particularly desirable given that different verbal instructions can knowingly or unconsciously trigger different demand effects and mental representations.

We ran 15 sessions of 30 rounds each, with 6 leaders and 6 followers in each session who were randomly matched in pairs at the beginning of every round to play a sequential quantity-setting duopoly game. We implemented 5 treatments with each treatment being implemented in 3 different sessions. In the first treatment, NOISE0, followers were always informed about the leaders' quantities. In treatments NOISE25, NOISE50, and NOISE75, respectively, followers were informed about the leaders' quantities with probability 0.25, 0.50, and 0.75, implying that each follower was informed in some rounds but not in others. And, in treatment NOISE100 followers were never informed about the leaders' quantities. An important feature of our design was that each subject had a strategy set that contained 20 possible quantity choices (that is, we used a large 20×20 payoff matrix). In contrast, Huck and Müller (2000) used a small 2×2 payoff matrix; this design might have forced their results to be "extreme" since in a 2×2 game

² Our experiments then differ from those of Morgan and Várdy (2004) mainly in that in their experiments, observability is endogenous (followers decide whether or not to become informed), while in ours it is exogenous.

³ To illustrate this difference, consider for instance Brander and Spencer's (1985) strategic trade model in which a home firm and a foreign firm compete in a third-country by setting quantities. By giving the home firm a subsidy, the home government shifts the home firm's best response function outward. If the foreign firm observes this policy, it cuts its own output in response and this benefits the home firm. In practice however, the foreign firm may sometimes be unaware of the subsidy that the home firm gets; this possibility corresponds to errors in communication. Alternatively, it might be that the foreign firm is always aware of the subsidy but is unsure as to its exact size. This case corresponds to errors in perceptions. Obviously, in real life either case is possible.

⁴ In other words, the Cournot and the Stackelberg models could be reinterpreted as the polar cases of a family of duopoly games that differ only with respect to a single parameter: the probability that the leader's action would be observed by followers. We believe that this observation is of independent interest since there is no reason to believe that specific applications should necessarily correspond to one of the polar cases.

there is little room for gradual adjustment. We believe that our current design is advantageous as it allows subjects to modify their actions in “smaller” steps while better separating the various outcomes from one another.

Our main findings can be summarized as follows:

- **Weak monotonicity of output in the level of noise:** On average, leaders chose 2.35–2.50 units less in treatments NOISE75 and NOISE100 than in treatments NOISE0, NOISE25, and NOISE50, but their behavior did not vary significantly within each group of treatments. Likewise, on average, uninformed followers chose 0.77–1.01 units more in treatments NOISE50, NOISE75, and NOISE100 than in treatment NOISE25, but their behavior did not vary significantly across the first group of treatments. The behavior of informed followers was not significantly affected by the level of noise.

The first set of findings shows that we only have weak rather than strict monotonicity as the theory predicts. That is, what matters for leaders and uninformed followers is only whether there are “large” or “small” levels of noise. Nonetheless, it seems that the leaders have recognized that they can enjoy a first-mover advantage when their actions are observed by followers but not otherwise and that uninformed followers have correctly anticipated the leaders’ behavior and reacted accordingly. And, consistent with the theory, informed followers have reacted only to what they saw irrespective of the ex-ante probability of this event. Fischer et al. (2006) obtain similar experimental results in the context of a two-party bargaining game in which second movers either perfectly observe the first movers’ demands (and either accept or reject these demands) or else observe nothing and need to state their own demands (the payoffs then are determined according to the Nash demand game). They find that the demands of first movers (uninformed second movers) weakly decrease (increase) with the level of noise, while the behavior of informed second movers is independent of the level of noise.

While the first set of findings is broadly consistent with the theory, the next set reveals some important differences:

- **Underproduction by leaders:** On average, leaders chose quantities below their equilibrium quantities in all treatments, although the deviation from the equilibrium quantities was smaller in treatments with higher levels of noise.
- **Overproduction by uninformed followers:** On average, uninformed followers chose quantities above their equilibrium quantities in all treatments, although the deviation from the equilibrium quantities was smaller in treatments with higher levels of noise.
- **Overreaction by informed followers:** On average, informed followers chose quantities above their best-response to the leaders’ quantities. Their overreaction was larger the larger was the gap between a leader’s quantity and the Cournot output.

The leaders’ tendency to underproduce and the informed followers’ tendency to overreact are similar to experimental results of Huck et al. (2001) in the context of quantity leadership games under perfect observability. They find that on average, leaders chose quantities that are almost halfway between the Stackelberg leader’s quantity and the (symmetric) Cournot quantity, while followers overreacted by about 1 unit to the leaders’ quantities.⁵ In the present experiments how-

⁵ Huck et al. (2004) find that the leaders’ tendency to underproduce and the followers’ tendency to overreact diminish but do not disappear when the payoffs of leaders and followers in the Stackelberg outcome are equalized. Weimann et al.

ever, the leaders' decisions are much harder since followers do not always observe the leaders' choices.

The present experiments also generate a much richer set of observations about the followers' behavior. Among other things, the results reveal the following:

- **Best-response was the modal behavior for followers:** The modal behavior of both informed and uninformed followers was to respond optimally to the leaders' actual quantities. Quite naturally, informed followers played a best-response more than twice as often as uninformed followers (56.1 vs. 25% of the cases).
- **Over- and under-reactions:** When informed followers did not play a best-response, they almost always overreacted to the leaders' output. In contrast, uninformed followers underreacted in 33.5% and overreacted in 41.5% of the cases.
- **Time trends:** As sessions progressed, uninformed followers played a best-response more often while informed followers played a best-response less often and instead overreacted more often.
- **Persistence of followers' behavior:** Followers who overreacted (underreacted) in round $t - 1$ also tended to overreact (underreact) in round t , especially if they were uninformed in round t .

Taken together, these observations suggest that the soft behavior of leaders may have been a rational response to the aggressive behavior of followers.⁶ In particular, it seems that followers viewed the symmetric Cournot output as "fair" and whenever they were informed, they "punished" leaders who produced more than the Cournot output by overreacting to the leaders' quantities. Such punishments can be very effective as they entail only a small loss to the follower but hurt the leader substantially.⁷ The increasing frequency with which informed followers overreacted to the leaders' output and the persistence in their behavior suggest that as sessions progressed, informed followers became aware of this fact and/or "acquired a taste" for "punishing" leaders.

As for uninformed followers, it seems that for the most part, they were trying to estimate the leaders' output and play a best-response against it. This hypothesis is consistent with the observation that uninformed followers played a best-response more frequently as sessions progressed after gaining experience. This implies in turn that followers were willing to "punish" leaders who were trying to exploit their first-mover advantage only when they were certain that the leaders deserved to be punished. When uninformed, followers accommodated the leaders' behavior even

(2000) ran sequential rent-seeking experiments in which the unique subgame perfect equilibrium features a first-mover advantage. The results were that leaders not only underproduced, but in fact had a strategic disadvantage vis-à-vis followers. Another related study is Kübler and Müller (2002). They consider a sequential differentiated products duopoly market with price competition. However, unlike with quantity competition, their setting features a second-mover advantage rather than a first-mover advantage.

⁶ This is similar to Harrison and McCabe (1996) who show that soft proposer behavior in ultimatum experiments seems to be a best response against actual aggressive behavior by responders.

⁷ To see why, consider a quantity-setting model with two firms, A and B , producing a homogeneous good. The profits of the two firms are $\pi^A = P(q^A + q^B)q^A - C^A(q^A)$ and $\pi^B = P(q^A + q^B)q^B - C^B(q^B)$, where $P(\cdot)$ is the inverse demand function, q^A and q^B are the outputs of firms A and B and $C^A(\cdot)$ and $C^B(\cdot)$ are their cost functions. Now, fix an equilibrium outcome (\hat{q}^A, \hat{q}^B) . A small deviation of firm B from \hat{q}^B lowers π^A by $\frac{\partial \pi^A}{\partial q^B} = P'(\hat{q}^A + \hat{q}^B)\hat{q}^A$, but due to the envelop theorem, it only has a negligible effect on π^B .

though on average, they correctly predicted that the leaders were exploiting their first-mover advantage. The persistence of deviations from best response by uninformed followers suggests however that on average, they made systematic errors in predicting the leaders' outputs.

Despite the important differences between actual and equilibrium behavior, we believe that by and large, our results support the main quantitative implications of the theory. In particular, leaders seem to enjoy a greater first-mover advantage when there is less noise, albeit their advantage is smaller than the theory predicts and is only weakly monotonic in the level of noise.

The effect of imperfect observability on the ability of players to commit was also studied in the context of strategic delegation both theoretically (e.g., Katz, 1991; Fershtman et al., 1991; and Fershtman and Kalai, 1997) and experimentally (Schotter et al., 2000, and Fershtman and Gneezy, 2001). However, in that context, the imperfection is in the observation of the contract that one player offers to his delegate rather than in the observation of the first-mover's action as in our study.

The remainder of the paper is organized as follows. Section 2 describes the experimental design. The results of the experiments are presented and discussed in Section 3. Concluding remarks are in Section 4. Appendix A contains an English translation of the written instructions that were given to the subjects as well as the payoff matrix that they have used.

2. Experimental design

2.1. A noisy-leader game with errors in communication

Our experiments were based on the following game. Two quantity-setting firms, A and B , produce a single homogeneous good at no cost. The profits of the two firms are

$$\pi^A = (a - (q^A + q^B))q^A, \quad \text{and} \quad \pi^B = (a - (q^A + q^B))q^B, \quad (1)$$

where $a > 0$ and q^A and q^B are the quantity choices of the two firms. The strategic interaction between the two firms evolves in two stages. First, firm A chooses q^A . Then, firm B gets a signal about q^A that perfectly reveals q^A with probability $1 - \varepsilon$ and reveals nothing with probability ε . In the latter case, firm B only knows that firm A already chose q^A but does not observe q^A . Based on the signal, firm B chooses q^B , and the profits of the two firms are realized. The objective of both firms is to maximize their respective profits.

We now characterize the equilibrium of this game in order to establish a benchmark against which we can compare our experimental results. In this equilibrium, firm B plays a best-response against its belief about q^A , firm A chooses q^A to maximize its expected profit given firm B 's strategy, and firm B 's belief about q^A is consistent with q^A . When firm B is informed about q^A , its best-response is $BR^B(q^A) = \frac{a - q^A}{2}$. When firm B is uninformed about q^A , its best-response is $BR^B(b) = \frac{a - b}{2}$, where b is firm B 's belief about q^A . Given $BR^B(q^A)$ and $BR^B(b)$, firm A 's expected profit is

$$E\pi^A = (1 - \varepsilon) \left(a - q^A - \frac{a - q^A}{2} \right) q^A + \varepsilon \left(a - q^A - \frac{a - b}{2} \right) q^A. \quad (2)$$

The equilibrium strategy of firm A is defined implicitly by the following first order condition:

$$(1 - \varepsilon) \left(\frac{a}{2} - q^A \right) + \varepsilon \left(\frac{a}{2} - 2q^A + \frac{b}{2} \right) = 0. \quad (3)$$

But since in equilibrium, $b = q^A$ (firm B 's belief about q^A is correct), this condition implies that the equilibrium strategy of firm A is:

$$\hat{q}^A = \frac{a}{2 + \varepsilon}. \tag{4}$$

Given \hat{q}^A , an uninformed firm B will choose the quantity

$$\hat{q}^B = BR^B(\hat{q}^A) = \frac{a(1 + \varepsilon)}{2(2 + \varepsilon)}, \tag{5}$$

while an informed firm B will choose the quantity

$$BR^B(q^A) = \frac{a - q^A}{2}. \tag{6}$$

The equilibrium then is fully characterized by Eqs. (4)–(6).⁸ Note that there is a fundamental difference between firm B 's strategy when it is informed about q^A and when it is not: in the former case, firm B simply chooses a best-response against the actual value of q^A , whatever it is. In the latter case, firm B does not observe q^A and hence chooses a best-response against the equilibrium value of q^A rather than the actual value. Hence, when $\varepsilon = 0$ (firm B is always informed), the game is exactly like a Stackelberg duopoly model with linear demand and marginal cost: firm A chooses $\frac{a}{2}$ units which is the monopoly output, while Firm B chooses $\frac{a}{4}$ units which is half of the monopoly output. At the other extreme where $\varepsilon = 1$ (firm B is never informed), the game is identical to a (positional order protocol of a) Cournot duopoly model and both firms produce $\frac{a}{3}$ units. As ε increases from 0 to 1, firm A 's quantity falls continuously from the Stackelberg leader's quantity of $\frac{a}{2}$ to the Cournot quantity of $\frac{a}{3}$, whereas firm B 's quantity increases continuously from the Stackelberg follower's quantity of $\frac{a}{4}$ to the Cournot quantity of $\frac{a}{3}$.⁹ Consequently, firm A enjoys a larger first-mover advantage as ε falls from 1 to 0.

At an intuitive level, when $\varepsilon = 0$, firm A commits itself to an aggressive behavior by choosing $\frac{a}{6}$ units more than the Cournot output ($\frac{a}{2}$ units instead of $\frac{a}{3}$ units). This gives firm A a strategic advantage vis-à-vis firm B because it induces firm B to cut its quantity by $\frac{a}{12}$ units below the Cournot level ($\frac{a}{4}$ units instead of $\frac{a}{3}$ units). Firm A 's aggressive behavior means that it does not play a best response against q^B : given that $q^B = \frac{a}{4}$, firm A would have liked to cut q^A from $\frac{a}{2}$ to $\frac{a - a/4}{2} = \frac{3a}{8}$. However, it is precisely because firm A cannot alter its commitment to produce $\frac{a}{2}$ units that it gains a strategic advantage vis-à-vis firm 2. As ε grows from 0, the probability that firm B will not observe q^A increases. Whenever firm B does not observe q^A , it acts according to its belief about q^A rather than according to the actual value of q^A . Holding firm B 's belief, b , fixed, firm A finds it optimal to play a best response against b , implying for instance that if firm B would expect that $q^A = \frac{a}{2}$ and would choose $q^B = \frac{a}{4}$, firm A would have actually preferred to set $q^A = \frac{3a}{8}$. In equilibrium of course, firm B fully anticipates this and hence, some of firm A 's

⁸ The game does not admit mixed strategy equilibria. To see why, suppose that the leader plays a mixed strategy. Clearly, the best-response of an informed follower is unique and is independent of how the leader chooses q^A . Since π^B is linear in q^A , an uninformed follower cares only about the expected value of q^A , where the expectation is generated using the follower's belief about the leader's mixed strategy. Hence, an uninformed follower has a unique best response against the expected value of q^A . Consequently, in equilibrium, both the informed and the uninformed followers play pure strategies. Finally, taking the best responses of the informed and uninformed followers as given, π^A has a unique maximizer and hence, the leader also plays a pure strategy in equilibrium.

⁹ In equilibrium, firm B correctly anticipates q^A and hence, q^B is the same irrespective of whether firm B is informed or uninformed about q^A .

commitment power is lost. Firm *A* does not lose its commitment power entirely because with probability $1 - \varepsilon$, firm *B* still observes q^A , in which case it is beneficial for firm *A* to commit to a large output level.

2.2. Experimental implementation

We implemented 5 treatments of the noisy-leader game using the software tool kit *z-Tree* (Fischbacher, 1999). We ran 3 sessions for each of the 5 treatments: one set of the 5 treatments at Humboldt University in Berlin and two more sets at Friedrich-Schiller University in Jena (Germany). Each session included 12 different subjects who were graduate and undergraduate students from various departments and were either randomly recruited from a pool of potential participants or invited to participate by leaflets distributed around the university campus. Sessions lasted between 60 and 75 minutes and the average earnings were about 17 Euros.

The 5 treatments differed with respect to the underlying noise level. In treatment NOISE0 (Stackelberg treatment), we set $\varepsilon = 0$, so all followers were informed about the quantity chosen by the leader with whom they were matched. In treatments NOISE25, NOISE50 and NOISE75, we set ε at 0.25, 0.50, and 0.75, respectively, so followers were informed in some rounds but not in others. When a follower was uninformed, the computer screen displayed the message “You don’t get an information about the quantity produced by firm *A*.” Finally, in treatment NOISE100 (Cournot treatment), we set $\varepsilon = 1$, so after the leader made a choice, the follower’s screen displayed the message “Firm *A* has decided, please make your decision now!”¹⁰ At the end of each round, subjects were told about q^A , q^B , whether or not firm *B* was informed about q^A , own profit in the last round, and own cumulative profit.

Upon arrival in the lab, subjects were assigned a computer screen and received written instructions in German (an English translation appears in Appendix A). After reading the instructions, subjects were allowed to ask clarifying questions that were answered in private. In the instructions, subjects were told that they were to act as a firm and will be randomly matched in each of 30 rounds with another firm and that both firms will choose output levels and will earn profits that were specified in a payoff matrix. At the beginning of each session, 6 subjects were randomly assigned the role of firm *A* (a leader) and 6 were assigned the role of firm *B* (a follower). Players’ roles were kept fixed during the entire session.

Apart from written instructions, subjects also received a payoff matrix (see Appendix A) in which the profits were expressed in terms of a fictitious currency called “Taler”. In order to ensure that the outcomes were sufficiently separated from one another without making the payoff matrix excessively large, we set $a = 60$ and asked subjects to choose quantities from the set $\{13, 14, \dots, 32\}$.¹¹ With $a = 60$, the Stackelberg leader’s and follower’s quantities are 30 and 15 units, respectively, the Cournot output is 20 units, and the symmetric collusive output is 15 units.

¹⁰ Treatment NOISE100 corresponds to a Cournot game with a Positional Order Protocol (POP), where followers know that the leaders took actions but not what these actions are. Güth et al. (1998) and Müller (2001) provide experimental evidence that the POP does not change the behavior of subjects in games with a unique Nash equilibrium (like the Cournot game in treatment NOISE100). In games with multiple Nash equilibria, however, POP may affect behavior (see e.g., Cooper et al., 1993; Camerer et al., 2004).

¹¹ In order to use a more “natural” range of numbers, we shifted those numbers 12 positions to the left such that the possible quantity choices in the payoff matrix were $\{1, 2, \dots, 20\}$. In addition, we ensured that the best-response of informed followers was single-valued by subtracting one Taler in several cells in the matrix.

Since in equilibrium leaders get a higher profit than followers (except in treatment NOISE100 where the equilibrium profits are the same), followers might feel that leaders have an “undue advantage” and may therefore overproduce in order to “punish” leaders who try to exploit their first-mover advantage. To neutralize such fairness considerations as much as possible, we assigned to each subject an individual and confidential exchange rate from Talers to real currency (DM in Berlin and Euro in Jena). We felt that if subjects would not know each other’s exchange rates, fairness considerations would play a smaller role. The exchange rates were randomly selected from the set {300, 320, 330, 340, 350} (e.g., 300 Talers = DM 1 in the Berlin sessions and 300 Talers = 0.5 Euros in the Jena sessions) and subjects got personal messages before the session started that informed them about their own exchange rates but *not* about their rivals’ exchange rates. In the 5 sessions held at Berlin and 5 of the sessions held at Jena (each corresponding to one of the 5 treatments), subjects were not informed about the range of possible exchange rates; to ensure that all subjects were aware of this feature, the personal message that each subject received included the line “Keep in mind that other participants do not necessarily have the same exchange rate.” In 5 of the sessions at Jena (one session for each treatment), subjects were informed about the distribution of exchange rates in their session but not the actual exchange rates of other participants.^{12,13}

3. Results

With 15 sessions (3 for each treatment) of 30 rounds each, and 6 leaders and 6 followers in each session, we have $15 \times 6 \times 30 = 2700$ leaders’ quantity choices, and 2700 followers’ quantity choices. After an initial review of the data, we decided to exclude subjects 47 and 68 who played as followers in treatments NOISE75 and NOISE0, respectively, from the analysis of followers’ behavior. Both subjects adopted highly idiosyncratic, self-explanatory strategies that require no further analysis.¹⁴ This left us with 2640 observations on followers’ choices, of which

¹² The exact message that each subject got was as follows: “Your exchange rate was randomly drawn from the set {300, 310, 320, 330, 340, 350} of possible exchange rates. As a result of the individual random draw for each participant, there is (including you) 1 participant with exchange rate of 300, 3 participants with exchange rate 310, 3 participants with exchange rate 320, 2 participants with exchange rate 330, 1 participant with exchange rate 340, and 2 participants with exchange rate 350.” The distribution of exchange rates was randomly generated once and was then used in all sessions.

¹³ Prior to the 15 sessions reported here, we ran 2 pilot sessions for treatment NOISE0. The design of these sessions was exactly as described above, except that in the first pilot, all participants had the same exchange rate from Talers to DM, and this was commonly known, and in the second pilot, subjects switched roles in every round. In both sessions the estimated reaction function of followers against q^A had a *positive* rather than a *negative* slope, implying that followers behaved as if their strategies were *strategic complements* to the leaders’ strategies rather than *strategic substitutes*. We believe that these results were due to the fact that followers felt that leaders had an undue advantage and therefore “punished” leaders who tried to exploit their first-mover advantage by overproducing. We therefore decided to use confidential exchange rates to neutralize these fairness considerations as much as possible and focus on other aspects of noisy leadership.

¹⁴ Subject 47 adopted a predatory strategy and chose 32 units in each of the first 4 rounds and 30 units in all other rounds. In the post-experimental questionnaire he wrote “The decisive thought was that in a competition with two contestants the aim must be to weaken the rival in the short run and to take him over in the long run in order to then gain a maximal payoff as a monopolist.” Follower subject 68 (who participated in treatment NOISE0 and hence was informed in all rounds) employed a “tit-for-tat” strategy and always matched the leader’s quantity. In the post-experimental questionnaire he wrote: “I was not concerned with payoff maximization, but with a just balance regarding payoffs. Therefore, in all of the 30 rounds I have produced the same number of units as firm A despite the risk of earning 0. My goal as a B-firm was to ‘educate’ the A-firm. That is, the optimal outcome for me would have been the combination 15:15.” Obviously, the strategies of both subjects are self-explanatory and are at odds with the theory. Including these subjects in the analysis of the followers’ behavior would only bias the results and make it harder to infer what guided the other followers whose behavior is not immediately clear.

Table 1

Actual and equilibrium behavior of leaders and followers (Robust standard errors in parentheses)

Treatment	mean q^A ($N = 2700$)	\hat{q}^A	mean $q^{B, \text{uninfo}}$ ($N = 1357$)	$\hat{q}^{B, \text{uninfo}}$	mean $\Delta_{i,t}^{\text{info}}$ ($N = 1283$)
NOISE0 (Stackelberg)	22.74 (4.33)	30	–	–	1.50 (3.10)
NOISE25	22.72 (3.14)	26.7	19.40 (3.47)	16.7	2.10 (3.86)
NOISE50	22.55 (3.25)	24	20.26 (2.41)	18	1.05 (2.29)
NOISE75	20.38 (1.91)	21.8	20.15 (2.36)	19.1	0.60 (2.53)
NOISE100 (Cournot)	20.23 (2.73)	20	20.45 (3.25)	20	–

1283 were made by informed followers who saw q^A before choosing their own quantities, and 1357 were made by uninformed followers who were only told that the leader had already chosen q^A but not what q^A was.

In Table 1, we report for each treatment the means and standard deviations of the leaders' quantities, q^A , the uninformed followers' quantities, $q^{B, \text{uninfo}}$, as well as the equilibrium quantities of leaders, \hat{q}^A , and uninformed followers, $\hat{q}^{B, \text{uninfo}}$. For informed followers we report the mean and standard deviation of $\Delta_{i,t}^{\text{info}} \equiv q_{i,t}^{B, \text{info}} - BR^B(q_{i,t}^A)$ which is the gap between the actual quantity of informed follower i in round t , and the follower's best-response to q^A in that round. According to the theory, we should have $\Delta_{i,t}^{\text{info}} = 0$ for all i and all t . Note that in treatment NOISE0, all followers were informed, whereas in treatment NOISE100, all followers were uninformed. For obvious reasons, we have in treatment NOISE25 more observations on informed followers than in treatments NOISE50 and NOISE75 (400 vs. 261 and 112, respectively) but fewer observations on uninformed followers (140 vs. 279 and 398, respectively).

Table 1 shows that on average, informed followers chose quantities that exceeded their best-responses by 0.6–2.1 units. In addition, leaders chose on average smaller quantities, while uninformed followers chose larger quantities than the theory predicts. By and large however, and consistent with the theory, the average quantity of leaders (uninformed followers) is decreasing (increasing) with the level of noise.

In what follows we use regression analysis to study the behavior of leaders and followers in more detail and uncover some of the factors that were affecting their behavior. Since the same subjects were matched in each session for 30 rounds, the error terms in our regressions may be interdependent within each session. To deal with this potentially non-spherical error structure, we estimated the following regressions using the cluster option provided by 'STATA.' This option does not affect the estimated coefficients but estimates the standard errors using robust variance matrix calculations that relax the assumption of independence of errors within each session (cluster).¹⁵

¹⁵ Let G_1, \dots, G_M be the M clusters specified in the cluster option. Then, the formula for robust variance calculation used by the cluster option is $V_{\text{cluster}} = \frac{(N-1)M}{(N-k)(M-1)} (X'X)^{-1} (\sum_{i=1}^M u_i' u_i) (X'X)^{-1}$, where N is the number of observations, k is the number of independent variables in the regression, X is the $k \times M$ matrix of independent variables, and $u_i = \sum_{j \in G_\ell} e_j x_j$, where e_j is the residual for observation j in cluster G_ℓ , and x_j is a row vector of independent

3.1. Leaders' behavior

To assess leaders' behavior, we estimate the following OLS regression:

$$q_{j,t}^A = \alpha_0 + \alpha_1 D_{25} + \alpha_2 D_{50} + \alpha_3 D_{75} + \alpha_4 D_{100} + \alpha_{ER} DER + \gamma_1 T_t^2 + \gamma_2 T_t^3 + \delta_1 \Delta_{j,t-1}^- + \delta_2 (\Delta_{j,t-1}^{+,info}) + \delta_3 (\Delta_{j,t-1}^{+,uninfo}) + \sum_{j \in \Lambda} \tau_j D_j^A + \eta, \quad (7)$$

where $q_{j,t}^A$ is leader j 's quantity in round t , Λ is the set of index numbers of all subjects who played as leaders (firm A), and η is an error term. The independent variables in the regression are defined as follows:

- D_{25} , D_{50} , D_{75} , and D_{100} are treatment dummies equal to 1 if leader j participated in treatment NOISE25, NOISE50, NOISE75, or NOISE100, respectively, and equal to 0 otherwise. Thus, treatment NOISE0 (the Stackelberg treatment) serves as the reference group, and the treatment dummies measure the effect of noise relative to this benchmark.
- DER is a dummy variable equal to 1 if leader j participated in one of the 5 sessions in which the subjects were informed about the distribution of exchange rates in their session (in the remaining 10 sessions, subjects knew only their own exchange rate but not the range of possible exchange rates of other subjects in their session). We include this dummy variable to examine whether this design feature affected the leaders' behavior.
- $\Delta_{j,t} \equiv q_{j,t}^B - BR^B(q_{j,t}^A)$ is the gap between the actual quantity of the follower with whom leader j was matched in round t and the best-response to leader j 's quantity in the same round. The variable $\Delta_{j,t-1}^-$ is the 1-period lagged value of $\Delta_{j,t}$, conditional on it being negative, and $\Delta_{j,t-1}^{+,info}$ ($\Delta_{j,t-1}^{+,uninfo}$) is the 1-period lagged value of $\Delta_{j,t}$ conditional on it being positive and conditional on the follower being informed (uninformed). In other words, $\Delta_{j,t-1}^-$ includes underreactions by followers in the previous round (cases in which followers chose quantities below their best-response), while $\Delta_{j,t-1}^{+,info}$ and $\Delta_{j,t-1}^{+,uninfo}$, respectively, include overreactions by informed and uninformed followers in the previous round. We include these variables to examine whether leaders modified their behavior in a given round based on their experience in the previous round. We distinguish between $\Delta_{j,t-1}^{+,info}$ and $\Delta_{j,t-1}^{+,uninfo}$ because leaders are likely to interpret $\Delta_{j,t-1}^{+,info}$ as deliberate attempts by informed followers to "punish" them while interpreting $\Delta_{j,t-1}^{+,uninfo}$ as reflecting the uninformed followers' difficulty to predict the leaders' choices. We do not make a similar distinction for underreaction because almost all underreactions came from uninformed followers.
- T_t^2 and T_t^3 , respectively, are dummies equal to 1 if round t is in the 2nd third of a session (rounds 11–20) or the 3rd third (rounds 21–30), and equal to 0 otherwise. Therefore, rounds 1–10 serve as the reference group, with T_t^2 and T_t^3 capturing possible time trends in leaders' behavior.
- D_j^A where $j \in \Lambda$, are leader-specific dummies that control for idiosyncratic behavior of leaders. We restrict the sum of the coefficients of these dummies to zero for each treatment.¹⁶

variables for observation j , including the intercept. For more details, see STATA Corp. (1999, vol. 3, pp. 156–158 and 178–179), White (1980), and Rogers (1993).

¹⁶ This restriction was first proposed by Suits (1984). For a very useful discussion on the use of this approach in experimental economics, see Königstein (2000).

This restriction does not affect the results but makes it easier to interpret the results by ensuring that the estimated coefficient α_0 represents the average quantity chosen by leaders, while the estimated coefficient of each dummy, τ_j , measures the gap between leader j 's quantity and the average quantity selected by all leaders.

The results of leaders' regressions, L1–L4, are shown in Table 2. We do not report the coefficients of the leader-specific dummies since we are only interested in general tendencies rather than in the individual behavior of specific subjects. We note however that most of the leader-specific dummies were significant. Since regression L4 includes lagged variables ($\Delta_{j,t-1}^-$ and $\Delta_{j,t-1}^+$), we lose $15 \times 6 = 90$ observations on the leaders' behavior in round 1 of each treatment.

Regression L2 shows that informing leaders about the distribution of exchange rates in their session did not have a significant effect on their behavior. Since the coefficients γ_1 and γ_2 are insignificant, it appears from regressions L3 and L4 that the leaders' behavior was stable over the course of the sessions. Regressions L3 and L4 suggest that whenever followers were soft and underreacted, leaders tended to "reciprocate" by lowering q^A in the immediately following round by an average of 0.281 units for each unit of underreaction (since $\Delta_{j,t-1}^-$ is negative by

Table 2
Results of the leaders' regressions (Robust standard errors in parentheses)

	Regression			
	L1	L2	L3	L4
α_0	22.737*** (0.568)	22.832*** (0.439)	22.845*** (0.471)	22.909*** (0.535)
$\alpha_1 (D_{25})$	-0.013 (0.693)	-0.013 (0.500)	-0.013 (0.692)	-0.038 (0.760)
$\alpha_2 (D_{50})$	-0.187 (0.611)	-0.187 (0.578)	-0.187 (0.611)	-0.134 (0.621)
$\alpha_3 (D_{75})$	-2.357*** (0.577)	-2.357*** (0.471)	-2.357*** (0.577)	-2.345*** (0.621)
$\alpha_4 (D_{100})$	-2.507*** (-0.691)	-2.507*** (0.555)	-2.507*** (0.691)	-2.275*** (0.739)
$\alpha_{ER} (DER)$		0.286 (0.165)		
$\gamma_1 (T_t^2)$			-0.064 (0.270)	-0.189 (0.235)
$\gamma_2 (T_t^3)$			-0.260 (0.259)	-0.383 (0.229)
$\delta_1 (\Delta_{j,t-1}^-)$				0.281*** (0.048)
$\delta_2 (\Delta_{j,t-1}^{+,info})$				0.086*** (0.027)
$\delta_3 (\Delta_{j,t-1}^{+,uninfo})$				0.082** (0.037)
R^2	0.370	0.377	0.371	0.413
No. of obs.	2700	2700	2700	2610

Note. Parameter estimates for subject dummies not shown.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3
p-values of (two tailed) pairwise cross-treatment differences in the means of q^A

	NOISE0	NOISE25	NOISE50	NOISE75
NOISE25	0.961	–	–	–
NOISE50	0.842	0.847	–	–
NOISE75	0.002	0.000	0.000	–
NOISE100	0.008	0.001	0.000	0.856

definition, $\delta_1 > 0$ implies a reduction in q^A .¹⁷ On the other hand, if followers were aggressive and overreacted, leaders tended to become aggressive as well and raised q^A in the immediately following round by an average of 0.08 units for each unit of overreaction. Overall, the estimated values of δ_1 , δ_2 , and δ_3 suggest that leaders tended to play soft after being “nicely” treated and aggressive after being “mistreated.”¹⁸

Turning to the main issue which is the effect of noise on leaders’ behavior, we present in Table 3 the *p*-values associated with Wald tests for equalities between pairs of the treatment dummy coefficients in regression L4 (since treatment NOISE0 serves as the reference group, the coefficient of this treatment is identically equal to 0).

Table 3 reveals that there is a highly significant difference between the leaders’ behavior in treatments with large noise levels (NOISE75 and NOISE100) and treatments with relatively small noise levels (treatments NOISE0, NOISE25, and NOISE50), but no significant difference between the leaders’ behavior within each of the two groups of treatments. In particular, the estimated treatment dummy coefficients reported above indicate that q^A is more than 2 units larger when there are small noise levels than when there are large noise levels. These results imply that q^A falls when there is more noise, as the theory predicts, though this occurs in steps rather than smoothly.

Observation 1. The behavior of leaders had the following features:

- (i) On average, leaders underproduced relative to their equilibrium quantities in all treatments, although the deviation from the equilibrium quantities is smaller in treatments with higher levels of noise.
- (ii) Leaders tended to lower their output following underreactions by followers and raise it following overreactions by followers.
- (iii) Leaders chose significantly smaller quantities in treatments NOISE75 and NOISE100 than in treatments NOISE0, NOISE25, and NOISE50, but there were no significant across-treatment differences within each of the two groups of treatments.

3.2. Informed followers’ behavior

We now turn to the behavior of informed followers’ and estimate their reaction function using the following OLS regression:

¹⁷ It is important to bear in mind, however, that since leaders and followers were randomly matched in each round, there was only a 1/6 chance that a leader would meet the same follower again in the next round.

¹⁸ Underreactions can be interpreted as a “nice” behavior because they benefit the leader at a personal cost to the follower who produces less than the payoff maximizing output. An overreaction can be interpreted as “mean” or “unkind” behavior because the follower sacrifices a monetary payoff by overproducing in order to hurt the leader.

$$q_{i,t}^{B,\text{info}} = \alpha_0 + \beta q_{i,t}^A + \alpha_{ER} DER + \beta_{ER} DER \times q_{i,t}^A + \gamma_1 T_t^2 + \gamma_2 T_t^3 + \sum_{i \in INF} \tau_i D_i^B + \sum_{i \in INF} \theta_i D_i^B q_{i,t}^A + \eta, \quad (8)$$

where $q_{i,t}^{B,\text{info}}$ is follower i 's quantity in round t provided the follower was informed in that round; $q_{i,t}^A$ is the quantity of the leader with whom follower i was matched in round t ; DER , T_t^2 , and T_t^3 are as in the leaders' regressions; D_i^B are follower-specific dummies, INF is the set of index numbers of all subjects who played as followers (firm B) and were informed at least in some rounds¹⁹; and η is an error term. The follower-specific dummies affect both the intercept and slope of the best-response function and are intended to control for the idiosyncratic behavior of individual followers. As in the leaders' case, we restrict the sum of the τ_i 's and the sum of the θ_i 's to 0 for each treatment in order to ensure that the estimated coefficients α_0 and β will represent the intercept and slope of the "average" response of informed followers, while the coefficients τ_j and θ_j will measure the deviation of follower j 's behavior from this average. We do not include treatment dummies in the regression since according to the theory, informed followers should play a best-response against q^A irrespective of the likelihood of observing it. In Section 3.4 we will show evidence that indeed there were no across-treatments differences in the behavior of informed followers.

As in the leaders' case, we used the cluster option provided by 'STATA' to deal with potential independence of the error terms within each session (cluster). The regression results are presented in Table 4. Again, we do not report the coefficients of the follower-specific dummies but note that most of them were significant.

The table shows that, relative to the equilibrium best-response function, the estimated reaction function of informed followers has a lower intercept (22.66–22.83 compared with 30) and its slope is much flatter (−0.1 compared with −0.5). Regressions INF2 and INF3 show that the response function of informed followers did not change significantly when the followers were informed about the distribution of exchange rates in their session and remained stable over the course of the sessions (the coefficients α_{ER} , β_{ER} , γ_1 , and γ_2 are all highly insignificant).

Observation 2. The estimated reaction function of informed followers had a smaller intercept and was flatter than the equilibrium best-response function and remained stable over time.

3.3. Uninformed followers' behavior

To assess the behavior of uninformed followers, we estimate the following OLS regression:

$$q_{i,t}^{B,\text{uninfo}} = \alpha_0 + \alpha_1 D_{50} + \alpha_2 D_{75} + \alpha_3 D_{100} + \alpha_{ER} DER + \beta q_{i,t-1}^A + \gamma_1 T_t^2 + \gamma_2 T_t^3 + \sum_{i \in UNF} \tau_i D_i^B + \sum_{i \in UNF} \theta_i D_i^B q_{i,t-1}^A + \eta, \quad (9)$$

where $q_{i,t-1}^A$ is the quantity of the leader with whom follower i (who is uninformed in round t) was matched in round $t - 1$, and UNF is the set of index numbers of all subjects who played as followers (firm B) and were uninformed at least in some rounds. We include $q_{i,t-1}^A$ in the regression to examine how uninformed followers adjusted their behavior on the basis of their

¹⁹ Recall however that subjects 47 and 68 were omitted from the data.

Table 4
Results of the informed followers' regressions (Robust standard errors in parentheses)

Independent variable	Regression		
	INF1	INF2	INF3
α_0	22.837*** (0.924)	22.659*** (0.868)	22.844*** (0.890)
$\beta (q_{i,t}^A)$	-0.109** (0.047)	-0.103** (0.044)	-0.109** (0.046)
$\alpha_{ER} (DER)$		-0.252 (0.868)	
$\beta_{ER} (DER \times q_{i,t}^A)$		0.006 (0.044)	
$\gamma_1 (T_t^2)$			0.086 (0.142)
$\gamma_2 (T_t^3)$			-0.085 (0.211)
R^2	0.571	0.573	0.572
No. of obs.	1283	1283	1283

Note. Parameter estimates for subject dummies not shown.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

most recent observation on q^A (recall that after each round subjects saw a screen that summarized what happened in that round, so even uninformed follower knew their leader's choice at the end of the round). The other variables are defined as in the leaders' and the informed followers' regressions. The treatment dummies D_{50} , D_{75} , and D_{100} measure the effect of noise on the behavior of uninformed followers relative to treatment NOISE25 which serves as a reference group (recall that in treatment NOISE0 there were no uninformed followers). The regressions' results are presented in Table 5.

Regressions UNINF2–UNINF3 show that uninformed followers were not significantly affected by the knowledge about the distribution of exchange rates in their session or by the behavior of leaders in the previous round, and their behavior remained stable over the course of the sessions (the coefficients α_{ER} , β , γ_1 , and γ_2 are all highly insignificant).

Turning to the effect of noise, Regression UNINF3 shows that uninformed followers chose 0.77–1.01 units less in treatment NOISE25 than in the other 3 treatments in which the level of noise was higher. Table 6 below shows the p -values associated with Wald tests for equalities between pairs of the treatment dummy coefficients in regression UNINF3 (since treatment NOISE25 serves as the reference group, the coefficient of this treatment is identically equal to 0).

The table shows that the difference between treatment NOISE25 and the other 3 treatments is significant. The table also shows however that there are no significant differences in the behavior of uninformed followers between treatments NOISE50, NOISE75, and NOISE100. Thus, we can conclude that while more noise induces uninformed followers to raise their quantities as the theory predicts, the effect occurs, as in the leaders' case, in steps.

Observation 3. Uninformed followers chose significantly larger quantities in treatments NOISE50, NOISE75, and NOISE100, than in treatment NOISE25, but there was no significant cross-treatment differences in their behavior within the former group of treatments.

Table 5
Results of the uninformed followers' regressions (Robust standard errors in parentheses)

Independent variable	Regression		
	UNINF1	UNINF2	UNINF3
α_0	19.500*** (0.331)	18.692*** (0.872)	18.802*** (0.841)
$\alpha_{50} (D_{50})$	0.805** (0.345)	0.908*** (0.265)	0.910*** (0.273)
$\alpha_{75} (D_{75})$	0.638 (0.413)	0.767* (0.360)	0.771* (0.366)
$\alpha_{100} (D_{100})$	0.949*** (0.342)	1.010*** (0.294)	1.011*** (0.302)
$\alpha_{ER} (DER)$		-0.081 (0.063)	-0.082 (0.063))
$\beta (q_{i,t-1}^A)$		0.035 (0.033)	0.034 (0.033)
$\gamma_1 (T_t^2)$			-0.191 (0.237)
$\gamma_2 (T_t^3)$			-0.137 (0.180)
R^2	0.512	0.587	0.587
No. of obs.	1357	1310	1310

Note. Parameter estimates for subject dummies not shown.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6
 p -values of (two tailed) pairwise cross-treatment differences in the means of $q^{B,uninfo}$

	NOISE25	NOISE50	NOISE75
NOISE50	0.007	–	–
NOISE75	0.059	0.551	–
NOISE100	0.007	0.347	0.311

3.4. Followers' over- and under-reactions

In this subsection we study the followers' behavior in greater detail by looking at $\Delta_{i,t}$ which is the gap between the actual quantity of follower i in round t and the follower's best-response to q^A in the same round. We begin by looking at the distribution of $\Delta_{i,t}$ for informed and uninformed followers across all treatments.

Figure 1 shows that the modes of both distributions are equal to 0: whether informed or uninformed, the modal behavior of followers was to play a best-response against q^A . Not surprisingly however, the mode for informed followers is more than twice as large as the mode for uninformed followers, with informed followers playing a best-response in 56.1% of the cases, compared with only 25% for uninformed followers. The figure also shows that the two distributions are very similar for $\Delta > 0$, implying that informed and uninformed followers tended to overreact to q^A at about the same frequency (37.8% for informed followers and 41.5% for uninformed followers). The main difference between the two distributions is that while informed followers almost never

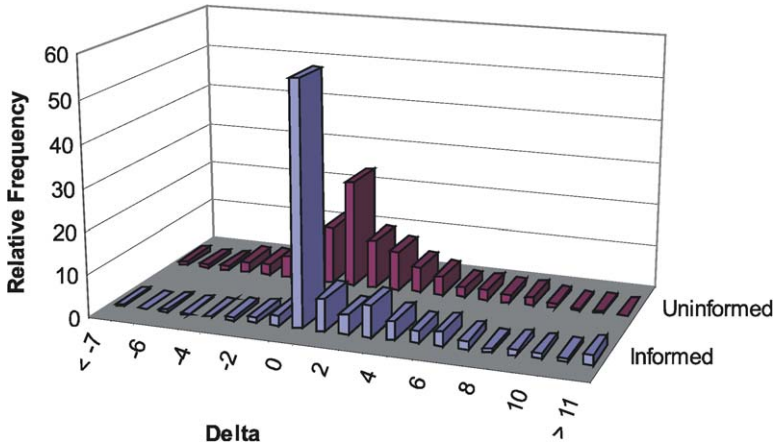


Fig. 1. The distribution of $\Delta_{i,t}$ for informed and uninformed followers, pooled across treatments.

underreacted to q^A (only 6.1% of the cases), uninformed followers ended up underreacting to q^A in 33.5% of the cases.²⁰

In Fig. 2 we examine the time trend in $\Delta_{i,t}$. The top panel in Fig. 2 shows that as sessions progressed, the distribution of $\Delta_{i,t}$ for uninformed followers became more concentrated around 0, with uninformed followers playing a best-response against q^A , (i.e., choosing $\Delta = 0$) in 13.7% of the cases in rounds 1–10, 26% in rounds 11–20, and 35.3% in rounds 21–30. This suggest that uninformed followers improved their predictions about q^A and “learned” to play best-responses against these predictions. The lower panel in Fig. 2 shows an opposite trend for informed followers: they played a best-response against q^A in 60.1% of the cases in rounds 1–10, 54.4% in rounds 11–20, and 53.9% in rounds 21–30. Instead of playing a best-response, informed followers tended to overreact to q^A . The frequency of overreactions has increased from 33.3% in rounds 1–10 to 39.8% in rounds 11–20 and 40.3% in rounds 21–30. Bearing in mind that small overreactions by followers hurt leaders substantially at a small personal loss to followers (see footnote 7), the decreasing time trend of $\Delta_{i,t}$ suggests that as sessions progressed, informed followers “learned” that small overreactions were sufficient to “discipline” leaders and induce them to play soft.

A final breakdown of the distribution of $\Delta_{i,t}$ according to treatments appears in Fig. 3. The figure shows that the mode of the distribution of $\Delta_{i,t}$ for both informed and uninformed followers is equal to 0 in all treatments. The lower panel reveals that there are virtually no cross-treatment differences in the behavior of informed followers, suggesting that they responded to what they observed but not to the ex-ante probability of this event. And, although the top panel shows several cross-treatment differences in the behavior of uninformed followers, we shall see shortly that these are not statistically significant.

²⁰ There is a big variance however in the behavior of individual followers. For instance, out of 72 followers who were informed (followers in all treatments but NOISE100), 16 always played a best-response and another 6 played a best-response in more than 90% of the cases. On the other hand, 9 followers played a best-response in less than 10% of the cases.

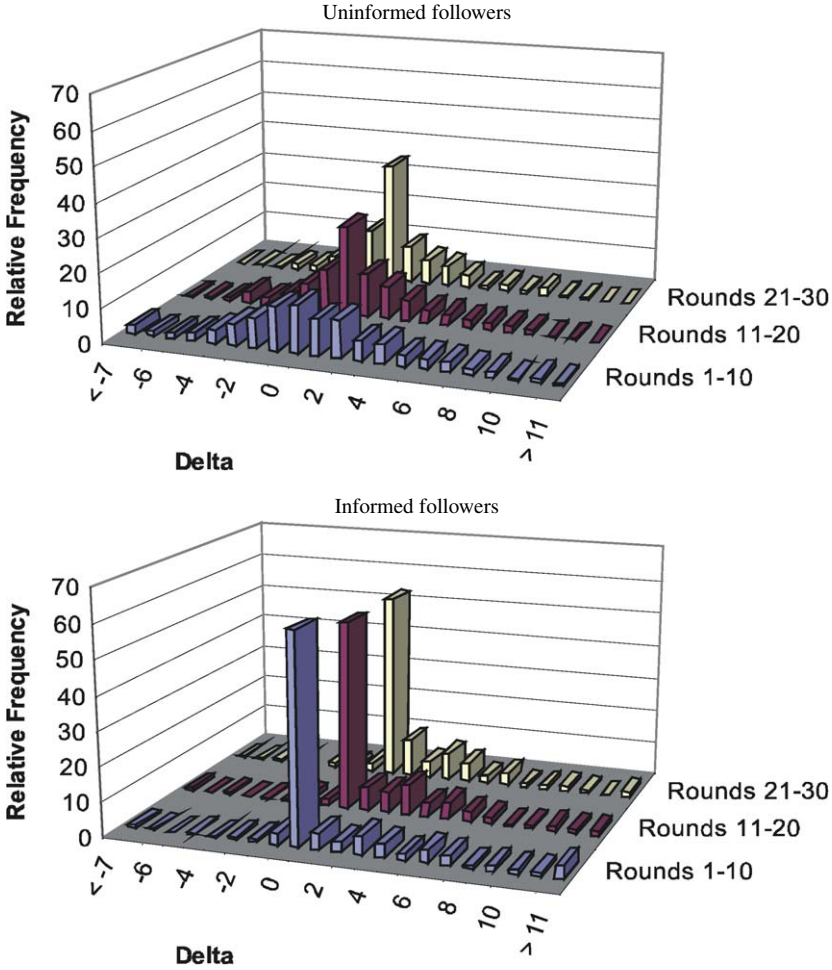


Fig. 2. The time trend in the distribution of $\Delta_{i,t}$.

To study the behavior of followers further, we estimated the following OLS regression:

$$\Delta_{i,t} = \alpha_0 + \alpha_1 \Delta_{i,t}^{A,info} + \alpha_2 L_{i,t}^{info} + \alpha_3 \Delta_{i,t-1}^{A,uninfo} + \alpha_4 L_{i,t-1}^{uninfo} + \rho_1 \Delta_{i,t-1}^{info} + \rho_2 \Delta_{i,t-1}^{uninfo} + \sum_{i \in F} \tau_i D_i^B + \eta, \tag{10}$$

where F is the set of index numbers of all subjects who played as followers (firm B), D_i^B are the follower-specific dummies, and η is an error term. The definitions of the new independent variables and the reasons for including them are as follows:

- $\Delta_{i,t}^{A,info} \equiv q_{i,t}^{A,info} - \hat{q}_i^A$ is the gap between the actual quantity of the leader with whom informed follower i was matched in round t and the equilibrium quantity of that leader. We include $\Delta_{i,t}^{A,info}$ in the regression to test the hypothesis that whenever leaders do not fully exploit their first-mover advantage and choose quantities below their equilibrium quantities, informed followers reward them by underreacting.

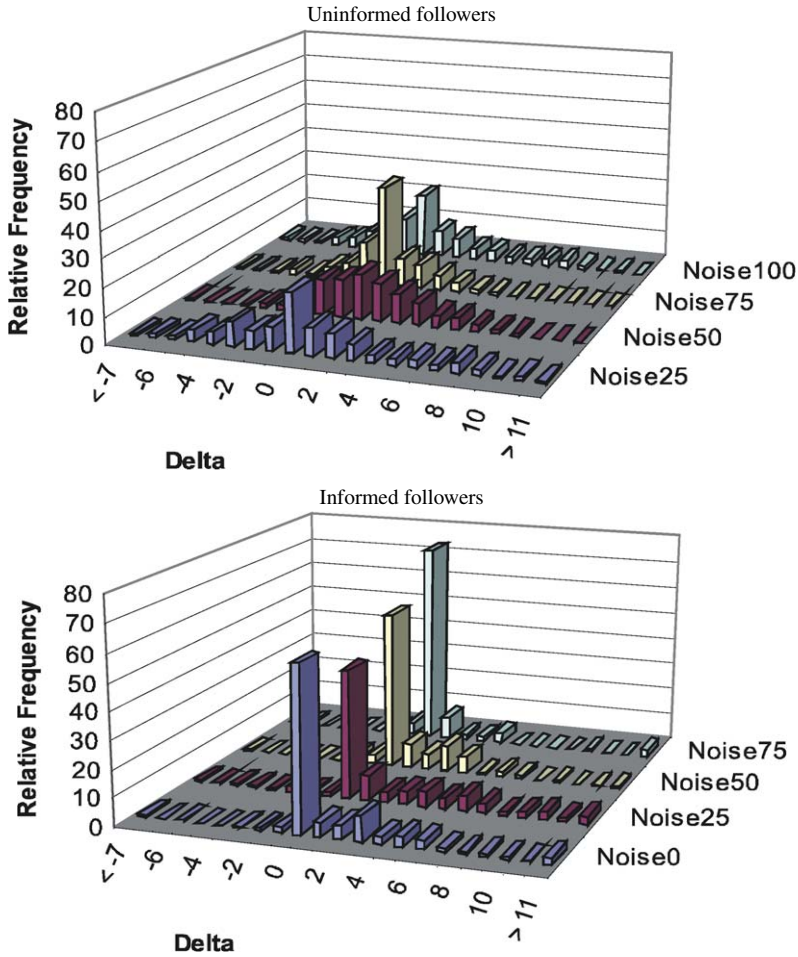


Fig. 3. The distribution of $\Delta_{i,t}$ according to treatments.

- $L_{i,t}^{\text{info}} \equiv q_{i,t}^{A,\text{info}} - 20$ is the gap between the actual quantity of the leader with whom informed follower i was matched in round t and the Cournot output, conditional on follower i being informed about $q_{i,t}^{A,\text{info}}$. We include $L_{i,t}^{\text{info}}$ in the regression to test the hypothesis that informed followers view the symmetric Cournot outcome, which in our design is $(20, 20)$, as “fair” (it gives leaders and followers equal payoffs) and “punish” leaders when $q^A > 20$.
- $\Delta_{i,t-1}^{A,\text{uninfo}}$ and $L_{i,t-1}^{\text{uninfo}}$ are the 1-period lagged values of $\Delta_{i,t}^{A,\text{uninfo}}$ and $L_{i,t}^{\text{uninfo}}$ which are defined similarly to $\Delta_{i,t}^{A,\text{info}}$ and $L_{i,t}^{\text{info}}$, but are conditional on follower i being uninformed about $q_{i,t}^{A,\text{info}}$. We include these variables to examine whether the behavior of uninformed followers in round t was affected by leaders’ deviations from either their equilibrium or their Cournot quantity in round $t - 1$ (note that $q_{i,t-1}^{A,\text{info}}$ is the most recent observation that uninformed follower i has on leaders’ behavior).
- $\Delta_{i,t-1}^{\text{info}}$ and $\Delta_{i,t-1}^{\text{uninfo}}$, respectively, are the 1-period lagged values of $\Delta_{i,t}$, conditional on follower i being either informed or uninformed about $q_{i,t}^A$. These variables are intended to examine whether the behavior of informed and uninformed followers showed persistence.

The results of the Δ regressions are shown in Table 7. Regression D1 shows results for pooled data across treatments and across informed/uninformed followers. Regressions D2 and D3 examine the treatment effects on the behavior of informed and uninformed followers by replacing $\Delta_{i,t}^{A,info}$, $L_{i,t}^{info}$, $\Delta_{i,t-1}^{A,uninfo}$, and $L_{i,t-1}^{uninfo}$ with their respective breakdowns by treatments.²¹

Table 7
Results of the Δ regressions (Robust standard errors in parentheses)

Independent variable	Regression					
	D1		D2		D3	
α_0	0.350*	(0.181)	0.365	(0.255)	0.302	(0.246)
$\alpha_1 (\Delta_{i,t}^{A,info})$	-0.020	(0.028)				
$\alpha_2 (L_{i,t}^{info})$	0.435***	(0.054)				
$\alpha_3 (\Delta_{i,t-1}^{A,uninfo})$	-0.085	(0.069)				
$\alpha_4 (L_{i,t-1}^{uninfo})$	0.176**	(0.077)				
$\alpha_5 (\Delta 0_{i,t}^{A,info})$			-0.005	(0.030)	0.003	(0.030)
$\alpha_6 (\Delta 25_{i,t}^{A,info})$			-0.013	(0.054)	0.003	(0.058)
$\alpha_7 (\Delta 50_{i,t}^{A,info})$			-0.041	(0.070)	-0.024	(0.069)
$\alpha_8 (\Delta 75_{i,t}^{A,info})$			-0.181	(0.225)	-0.176	(0.193)
$\alpha_9 (L 0_{i,t}^{info})$			0.403***	(0.031)	0.390***	(0.030)
$\alpha_{10} (L 25_{i,t}^{info})$			0.622***	(0.131)	0.599***	(0.122)
$\alpha_{11} (L 50_{i,t}^{info})$			0.251***	(0.073)	0.235***	(0.070)
$\alpha_{12} (L 75_{i,t}^{info})$			0.365	(0.441)	0.337	(0.396)
$\alpha_{13} (\Delta 25_{i,t-1}^{A,uninfo})$			0.020	(0.072)	0.002	(0.064)
$\alpha_{14} (\Delta 50_{i,t-1}^{A,uninfo})$			-0.257***	(0.065)	-0.249***	(0.061)
$\alpha_{15} (\Delta 75_{i,t-1}^{A,uninfo})$			0.125	(0.176)	0.088	(0.156)
$\alpha_{16} (\Delta 100_{i,t-1}^{A,uninfo})$			-0.005	(0.053)	-0.129**	(0.045)
$\alpha_{17} (L 25_{i,t-1}^{uninfo})$			0.115	(0.040)	-0.046	(0.043)
$\alpha_{18} (L 50_{i,t-1}^{uninfo})$			0.342***	(0.080)	0.245***	(0.080)
$\alpha_{19} (L 75_{i,t-1}^{uninfo})$			0.046	(0.265)	-0.025	(0.225)
$\rho_1 (\Delta_{i,t-1}^{info})$					0.105**	(0.040)
$\rho_2 (\Delta_{i,t-1}^{uninfo})$					0.254***	(0.043)
R^2	0.448		0.473		0.492	
No of obs.	2552		2552		2552	

Note. Parameter estimates for subject dummies not shown.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

²¹ For instance, $\Delta 0_{i,t}^{A,info}$, $\Delta 25_{i,t}^{A,info}$, $\Delta 50_{i,t}^{A,info}$, $\Delta 75_{i,t}^{A,info}$ are equal to $\Delta_{i,t}^{A,info}$ if follower i participated in treatment NOISE0, NOISE25, NOISE50, or NOISE75, respectively, and are equal to 0 otherwise (e.g., $\Delta 25_{i,t}^{A,info} = \Delta_{i,t}^{A,info}$, if follower i participated in treatment NOISE25 and $\Delta 25_{i,t}^{A,info} = 0$ otherwise). We do not define the variables $\Delta 100_{i,t}^{A,info}$, $L 100_{i,t}^{A,info}$, $\Delta 0_{i,t}^{A,uninfo}$, and $L 0_{i,t}^{A,uninfo}$ as in treatment NOISE100 all followers were uninformed while in treatment

Regressions D1–D3 show that the estimated coefficient α_1 is virtually 0, implying that informed followers did not react to leaders' deviations from their equilibrium quantities. This remains the case even when we break down the data by treatments (the coefficients α_5 – α_8 are all insignificant). Hence, it appears that informed followers did not interpret leaders' quantities below the equilibrium quantities as a “nice” behavior and did not feel compelled to “reward” it by underreacting. On the other hand, the coefficient α_2 is significant and equal to 0.43, implying that for each leader's quantity above the Cournot quantity of 20 units, informed followers overreacted by an average of 0.43 units. This is consistent with the hypothesis that informed followers viewed the leaders' first-mover advantage as “unfair” and “punished” leaders who were trying to exploit it. Regressions D2–D3 show that such overreactions were present in all treatments but NOISE75.

As for uninformed followers, it appears from regression D3 that in treatment NOISE50 they tended to underreact to q^A by an average of 0.249 units for each unit that q^A was below the equilibrium quantity in the immediately previous round and overreact by an average of 0.245 units for each unit that q^A was above the Cournot quantity in the immediate previous round. In all other treatments, uninformed followers do not seem to react to past deviations of leaders from either their equilibrium or Cournot quantities (an exception is the coefficient α_{16} which is negative and significant in D3).

Regression D3 also shows that the behavior of informed, and especially uninformed, followers was persistent as both ρ_1 and ρ_2 are significant and positive: other things being equal, followers who overreacted in round $t - 1$ also tended to overreact in round t . This tendency was more than twice as large when followers were uninformed than when they were informed. As discussed above, the persistence of informed followers could be due to inertia or might indicate that they “acquired a taste” for “punishing” leaders. For uninformed followers, the persistence could indicate systematic errors in predicting q^A .²²

Observation 4. The followers' tendency to over- and under-react can be summarized as follows:

- (i) Irrespective of whether followers were informed or uninformed, their modal behavior was to play a best-response against the leaders' output. Not surprisingly, however, informed followers played a best-response more than twice as often as uninformed followers (56.1% of the cases vs. 25% for uninformed followers).
- (ii) As sessions progressed, uninformed followers played a best-response against the leaders' outputs more often, whereas informed followers played a best-response less often (though the latter effect is small).
- (iii) Informed followers almost never underreacted. Their tendency to overreact was stronger the larger was the gap between the leader's quantity and the Cournot quantity of 20 units but was not affected by leader's deviations from their equilibrium quantities.
- (iv) With the exception of Treatment NOISE50, uninformed followers did not react to past deviations of leaders from either their Cournot or their equilibrium quantities.

NOISE0 all followers were informed. We also do not include the variable $L100_{i,t-1}^{\text{uninfo}}$ in regressions D2 and D3 since in treatment NOISE100, $\hat{q}_i^A = 20$, implying that $\Delta 100_{i,t-1}^{\text{A,uninfo}} = L100_{i,t-1}^{\text{uninfo}}$.

²² We also tested for time trends in the evolution of $\Delta_{i,t}^{\text{info}}$ and $\Delta_{i,t}^{\text{uninfo}}$ by including the “third” dummies, T_t^2 and T_t^3 , in the regression. However, the coefficients of these dummies were highly insignificant.

- (v) Followers' behavior showed persistence as followers who overreacted (underreacted) in round $t - 1$ also tended to overreact (underreact) in round t . The level of persistence was more than two times larger for uninformed followers than for informed followers.

3.5. Why were leaders soft?

Having examined the followers' behavior in detail, we now return to the leaders' behavior and briefly discuss possible reasons for why they underproduced relative to their equilibrium quantities.

The first reason might be that leaders—despite the random-matching scheme—were trying to induce collusive outcomes by choosing low quantities and thereby invite followers to behave similarly. But since only informed followers are aware of the fact that the leader has played soft, we should expect leaders to be more inclined to play soft in treatments with less noise. Yet, the leaders' choices do not show this pattern. Therefore it appears that the soft behavior of leaders was not motivated by collusion considerations.

Second, it could be that leaders were reluctant to fully exploit their first-mover advantage because they do not like inequality. But then it is reasonable to expect that this concern for inequality will greatly diminish if followers do not reciprocate and take advantage of the soft behavior of leaders. That is, it seems reasonable that leaders would not feel bad about exploiting their first-mover advantage if they expect followers to exploit their soft behavior. Table 8 shows for each treatment how followers (including subjects 47 and 68) responded when leaders chose quantities below the Cournot quantity of 20 units. These quantities can be interpreted as attempts by leaders to induce outcomes that give both players higher profits than they can get at the Cournot outcome.

As Table 8 shows, in 141 out of 167 cases (84.4%) in which q^A was strictly below 20 units, informed followers responded with $q^B > q^A$ and therefore ended up getting larger payoffs than the leaders. It might be thought that this lack of reciprocity by followers would induce leaders to exploit their first-mover advantage more often. Yet, as we saw earlier, the coefficients γ_1 and γ_2 in the leaders' regressions are insignificant, thereby implying that such pattern did not emerge in our experiments.

Table 8
Followers' average response to below-Cournot leader quantities

q^A	$BR^B(q^A)$	Mean reactions of informed and uninformed followers								
		NOISE0		NOISE25		NOISE50		NOISE75		NOISE100
		$q^{B,info}$	$q^{B,uninfo}$	$q^{B,info}$	$q^{B,uninfo}$	$q^{B,info}$	$q^{B,uninfo}$	$q^{B,info}$	$q^{B,uninfo}$	
13	24	22.57	–	–	19.00	–	–	–	22.00	
14	23	19.20	–	27.00	–	–	–	–	18.50	
15	23	21.90	20.00	22.00	19.00	21.43	21.00	23.00	20.09	
16	22	20.00	19.00	19.00	21.50	22.00	25.60	21.00	21.38	
17	22	22.00	17.25	20.50	20.38	22.00	19.80	24.67	20.86	
18	21	20.47	17.71	20.20	16.60	21.00	20.78	21.89	19.66	
19	21	20.84	19.50	20.50	21.22	20.83	19.39	20.38	20.26	
20	20	20.26	18.77	20.47	20.44	20.62	20.41	21.00	20.71	
Cases with $q^B > q^A$ when $q^A < 20$		72/82	–	29/42	–	18/20	–	22/23	–	

This brings us to a third possible reason which is that the soft behavior of leaders was a rational response to the aggressive behavior of informed followers. According to this hypothesis, leaders were reluctant to fully exploit their first-mover advantage because they wanted to avoid costly punishments by followers (recall from footnote 7 that the follower's punishments are proportional to the leader's action). This hypothesis is consistent with the observations that followers overreacted (i.e., chose $\Delta > 0$) in 40.4% of the cases (1090 out of 2700 cases) and underreacted (i.e., chose $\Delta < 0$) in only 19.9% of the cases (536 out of 2700 cases) and that informed followers tended to overreact more the farther q^A was from the Cournot output of 20 units.

4. Conclusion

Sequential decisions in markets are probably the rule rather than the exception. But in practice, early choices are not always perfectly revealed to rivals. This begs the question of how players behave in sequential strategic situations with imperfect observability. Studying such strategic situations empirically is extremely difficult, however, due to obvious limitations of available data sets. In this paper we study this issue with a controlled experiment under the assumption that followers either perfectly observe the leaders' choices or else they observe nothing.

The decision problems of leaders and uninformed followers are rather complex in our experimental design, especially given the fact that we used a large 20×20 payoff matrix. Despite that, by and large, the qualitative results we get are consistent with the theory. In particular, leaders seem to respond correctly to their ability (or lack thereof) to gain a first-mover advantage, albeit not to the extent predicted by the theory, while uninformed followers seem to correctly anticipate the leaders' behavior and respond accordingly. Therefore, our results seem to support the proposition that errors in communications erode the first-mover advantage of leaders but do it in a way which is related to the level of noise.

Our experiments also reveal some important differences between the predictions of the theory and the subjects' behavior. First, punishments by followers are very effective since a small overreaction to the leaders' choice entails only a negligible loss to the follower while inflicting a large loss on the leader which is proportional to the leader's quantity. Consequently, as leaders choose larger quantities, they become more vulnerable to overreactions by followers. This property, which has been almost completely neglected in the Industrial Organization literature, suggests that in sequential games in which strategies are strategic substitutes (like the duopoly game that we have considered), it is reasonable to expect leaders to play more cautiously than the theory predicts. This is particularly so in noisy-leadership games in which followers may remain uninformed about the leaders' choices and may therefore overreact to the leaders' choices inadvertently.²³ Our results suggest that leaders took this possibility into account and shaded their quantities below their equilibrium quantities in all treatments.

Second, informed followers are willing to sacrifice small amounts of their own payoffs in order to hurt leaders who try to exploit their first-mover advantage. In particular, their willingness to "punish" leaders is greater the farther away is the leader's choice from the symmetric Cournot outcome.

Third, it seems that followers do not try to "punish" leaders when they are uninformed, even if on average, they seem to correctly predict that leaders' quantities exceed the symmetric Cournot

²³ Although uninformed followers may end up underreacting (and thereby benefit leaders), risk aversion implies that leaders should be more concerned with potential overreactions than with underreactions.

output. Instead, they seem to simply try to play a best-response against their prediction of the leader's choices. This suggests that followers punish leaders who try to exploit their first-mover advantage only when they are certain that the leaders deserve to be punished. When uninformed, followers accommodate the leaders' behavior even if they would have punished it had they observed it.²⁴ In other words, it seems that followers punish only "what they see" but do not punish "what they do not see," even if on average they correctly anticipate the leaders' choices. A similar behavior has been observed in ultimatum experiments in which only proposers know the actual size of the pie: when the pie turns out to be large, most proposers offer exactly one half of the small pie and are never punished, even though the large pie is twice as likely implying that with a high probability the proposer's offer is in fact "unfair" (see Güth et al., 1996).

Acknowledgments

The second author acknowledges financial help from the German Science Foundation, DFG. Part of this research was done while the second author was visiting the Economics Department of New York University. For helpful comments and discussions we thank Chaim Fershtman, Itzhak Gilboa, Eilon Solan, an associate editor and a referee, and seminar participants at Tel Aviv University.

Appendix A

A.1. *Translated instructions (from German)*

Welcome to our experiment! Please read these instructions carefully! Do not talk to your neighbors and be quiet during the entire experiment. If you have a question, give notice. We will answer them privately.

In our experiment you can earn different amounts of money, depending on your behavior and that of other participants who are matched with you.

You act in the role of a firm which produces the same product as another firm in the market. Both firms always have to make a single decision, namely which quantities they want to produce. In the attached table, you can see the resulting profits of both firms for all possible quantity combinations.

The table is read as follows: the head of the row represents one firm's quantity (*A*-firm) and the head of the column represents the quantity of the other firm (*B*-firm). Inside the little box where row and column intersect, the *A*-firm's profit matching this combination of quantities is up to the left and the *B*-firm's profit matching these quantities is down to the right. The profit is denoted in a fictitious unit of money which we call Taler.

How do you make your decision? When the experiment starts, the computer screen will indicate whether you are an *A*-firm or a *B*-firm. You keep this role in the entire experiment. The procedure is that the *A*-firm always starts.

[The following paragraph only in treatment NOISE0.] This means that the *A*-firm chooses its quantity first (selects a line in the table) and the *B*-firm will be informed about the *A*-firm's choice. Knowing the quantity produced by the *A*-firm, the *B*-firm then decides on its quantity (selects a column in the table).

[The following paragraph only in treatments NOISE25, NOISE50 and NOISE75.] This means that the *A*-firm chooses its quantity first (selects a line in the table). Then a random move takes place that decides whether the *B*-firm will or will not be informed about the decision of the *A*-firm: With a probability of 25% [50%, 75%] the *B*-firm will be informed about the quantity chosen by firm *A*. With a probability of 75% [50%, 25%] the *B*-firm will not be informed about the quantity chosen by firm *A*. Then the *B*-firm decides on its quantity (selects a column in the table) either knowing or not knowing the quantity chosen by firm *A* before.

²⁴ A case in point is subject 31 who played as a follower in treatment NOISE50 and wrote in the post-experimental questionnaire: "As a *B*-firm one can only try to push down *A*'s profit if one knows its quantity and to try to optimize the own profit if *A*'s quantity is not known."

[The following paragraph only in treatment NOISE100.] This means that the *A*-firm chooses its quantity first (selects a line in the table) and the *B*-firm will not be informed about the *A*-firm's choice. Not knowing the quantity produced by the *A*-firm, the *B*-firm then decides on its quantity (selects a column in the table).

This procedure is repeated over thirty rounds. You do not know the participant with whom you serve the market. In each round you will be randomly matched with another participant such that always one *A*-firm and one *B*-firm will meet. That is, if you are an *A*-firm you will always be matched with a *B*-firm and vice versa.

After each round you will be informed about the quantity of the other firm as well as about your profit in the previous round and your total payoff so far.

The experiment will be conducted at the computer. This guarantees both anonymity between all participants and anonymity between you and the experimenter since your decisions can not be assigned to your person.

Your total payoff will be determined by the sum of your own payoffs in each round.

The exchange rate from Taler to DM valid for you will be displayed on the computer screen.

A.2. Payoff matrix

To save space, the payoff matrix presented here shows only the payoffs of the row player. The matrix that was used in the experiments showed the payoffs of both the row and the column player. This matrix however is too big to fit a standard size page.

Table A.1
Payoff matrix

Quantity	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
13	442	429	416	403	390	377	364	351	338	325	312	299	286	273	260	247	234	221	208	195
14	462	448	434	420	406	392	378	364	350	336	322	308	294	280	266	252	238	224	209	196
15	480	465	450	435	420	405	390	375	360	345	330	315	300	285	270	255	239	225	210	195
16	496	480	464	448	432	416	400	384	368	352	336	320	304	288	271	256	240	224	208	192
17	510	493	476	459	442	425	408	391	374	357	340	323	305	289	272	255	238	221	204	187
18	522	504	486	468	450	432	414	396	378	360	341	324	306	288	270	252	234	216	198	180
19	532	513	494	475	456	437	418	399	379	361	342	323	304	285	266	247	228	209	190	171
20	540	520	500	480	460	440	419	400	380	360	340	320	300	280	260	240	220	200	180	160
21	546	525	504	483	461	441	420	399	378	357	336	315	294	273	252	231	210	189	168	147
22	550	528	505	484	462	440	418	396	374	352	330	308	286	264	242	220	198	176	154	132
23	551	529	506	483	460	437	414	391	368	345	322	299	276	253	230	207	184	161	138	115
24	552	528	504	480	456	432	408	384	360	336	312	288	264	240	216	192	168	144	120	96
25	550	525	500	475	450	425	400	375	350	325	300	275	250	225	200	175	150	125	100	75
26	546	520	494	468	442	416	390	364	338	312	286	260	234	208	182	156	130	104	78	52
27	540	513	486	459	432	405	378	351	324	297	270	243	216	189	162	135	108	81	54	27
28	532	504	476	448	420	392	364	336	308	280	252	224	196	168	140	112	84	56	28	0
29	522	493	464	435	406	377	348	319	290	261	232	203	174	145	116	87	58	29	0	0
30	510	480	450	420	390	360	330	300	270	240	210	180	150	120	90	60	30	0	0	0
31	496	465	434	403	372	341	310	279	248	217	186	155	124	93	62	31	0	0	0	0
32	480	448	416	384	352	320	288	256	224	192	160	128	96	64	32	0	0	0	0	0

References

- Bagwell, K., 1995. Commitment and observability in games. *Games Econ. Behav.* 8, 271–280.
- Brander, J., Spencer, B.J., 1985. Export subsidies and international market share rivalry. *J. Int. Econ.* 18, 83–100.
- Camerer, C., Knez, M., Weber, R.A., 2004. Timing and virtual observability in ultimatum bargaining and ‘weak link’ coordination games. *Exper. Econ.* 7, 25–48.
- Chakravorti, B., Spiegel, Y., 1993. Commitment under imperfect observability: Why it is better to have followers who know that they don't know rather than those who don't know that they don't know. Discussion paper No. 103. Bellcore Economics.
- Cooper, R., DeJong, D.V., Forsythe, R., Ross, T.W., 1993. Forward induction in the Battle-of-the-Sexes games. *Amer. Econ. Rev.* 83, 1303–1316.

- Fershtman, C., Gneezy, U., 2001. Strategic delegation: An experiment. *RAND J. Econ.* 32, 352–368.
- Fershtman, C., Kalai, E., 1997. Unobserved delegation. *Int. Econ. Rev.* 38, 763–774.
- Fershtman, C., Judd, K., Kalai, E., 1991. Observable contracts: Strategic delegation and cooperation. *Int. Econ. Rev.* 32, 551–559.
- Fischbacher, U., 1999. z-Tree, Zurich Toolbox for readymade economic experiments. Working paper 21. Institute for Empirical Research in Economics, University of Zurich.
- Fischer, S., Güth, W., Müller, W., Stiehler, A., 2006. From ultimatum to Nash bargaining: Theory and experimental evidence. *Exper. Econ.* 9, 17–33.
- Güth, W., Huck, S., Ockenfels, P., 1996. Two-level ultimatum bargaining with incomplete information. *Econ. J.* 106, 593–604.
- Güth, W., Huck, S., Rapoport, A., 1998. The limitations of the positional order effect. *J. Econ. Behav. Organ.* 34, 313–325.
- Harrison, G., McCabe, K., 1996. Expectations and fairness in a simple bargaining experiment. *Int. J. Game Theory* 25, 303–327.
- Huck, S., Müller, W., 2000. Perfect versus imperfect observability: An experimental test of Bagwell's result. *Games Econ. Behav.* 31, 174–190.
- Huck, S., Müller, W., Normann, H.-T., 2001. Stackelberg beats Cournot—On collusion and efficiency in experimental markets. *Econ. J.* 111, 749–765.
- Huck, S., Königstein, M., Müller, W., 2004. Fairness in Stackelberg games. In: Huck, S. (Ed.), *Advances in Understanding Strategic Behaviour: Game Theory, Experiments, and Bounded Rationality: Essays in Honour of Werner Güth*. Palgrave, Basingstoke.
- Katz, M., 1991. Game-playing agents: Unobservable contracts as precommitments. *RAND J. Econ.* 22, 367–388.
- Königstein, M., 2000. Measuring treatment effects in experimental cross-sectional time series. In: Königstein, M. (Ed.), *Equity, Efficiency and Evolutionary Stability in Bargaining Games with Joint Production*, c. 2. Springer, Berlin/Heidelberg/New York.
- Kübler, D., Müller, W., 2002. Simultaneous and sequential price competition on heterogeneous duopoly markets: Experimental evidence. *Int. J. Ind. Organ.* 20, 1437–1460.
- Laffont, J.J., Tirole, J., 1993. *A Theory of Incentives in Procurement and Regulation*. MIT Press, Cambridge, MA.
- Maggi, G., 1999. The value of commitment with imperfect observability and private information. *RAND J. Econ.* 30, 555–574.
- Morgan, J., Várdy, F., 2004. An experimental study of commitment and observability in Stackelberg games. *Games Econ. Behav.* 49, 401–423.
- Müller, W., 2001. The quality of the signal matters: A note on imperfect observability and the timing of moves. *J. Econ. Behav. Organ.* 45, 99–106.
- Oechssler, J., Schlag, K., 2000. Does noise undermine the first-mover advantage? An evolutionary analysis of Bagwell's example. *Int. Game Theory Rev.* 2, 83–96.
- Rogers, W.H., 1993. Regression standard errors in clustered samples. *Stata Tech. Bull.* 13, 19–23.
- Rubinstein, A., 1989. The electronic mail game: Strategic behavior under almost common knowledge. *Amer. Econ. Rev.* 79, 385–391.
- Schelling, T.C., 1960. *The Strategy of Conflict*. Harvard Univ. Press, Cambridge, MA.
- Schotter, A., Zheng, W., Snyder, B., 2000. Bargaining through agents: An experimental study of delegation and commitment. *Games Econ. Behav.* 30, 248–292.
- von Stackelberg, H., 1934. *Marktform und Gleichgewicht*. Springer, Vienna and Berlin.
- STATA Corp., 1999. *Reference Manual*, release 6. STATA Press, TX.
- Suits, D., 1984. Dummy variables: Mechanics v. interpretation. *Rev. Econ. Statist.* 66, 177–180.
- van Damme, E., Hurkens, S., 1997. Games with imperfectly observable commitment. *Games Econ. Behav.* 21, 282–308.
- Várdy, F., 2004. The value of commitment in Stackelberg games with observation costs. *Games Econ. Behav.* 49, 374–400.
- Weimann, J., Yang, C.-L., Vogt, C., 2000. An experiment on sequential rent-seeking. *J. Econ. Behav. Organ.* 41, 405–426.
- White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48, 817–830.