Noisy leadership: An experimental approach*

Werner Güth†
Max Planck Institute

Wieland Müller‡
Humboldt University and New York University

Yossi Spiegel§
Tel Aviv University

March 12, 2002

Abstract

We examine the strategic behavior of leaders and followers in sequential duopoly experiments in which followers either perfectly observe the leaders' actions or else observe nothing. Our experiments show that consistent with the theory, leaders enjoy a greater first-mover advantage when followers observe their actions with higher probability. However, the results also show that (i) leaders do not fully exploit their first-mover advantage, (ii) when informed, followers tend to overreact slightly (i.e., choose quantities above their best-response to the leaders' quantities), and (iii) when uninformed, followers try to predict leaders' quantities and react optimally. This suggests that followers view the symmetric Cournot outcome as "fair" and whenever they observe leaders who are trying to exploit their first-mover advantage, they "punish" them by overreacting. Such punishments in turn induce leaders to behave more softly than the theory predicts.

Keywords: commitment; imperfect observability; sequential games; oligopoly; experiments.

JEL classification numbers: C72; C91.

---

*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, and seminar participants at Tel Aviv University.

†Max Planck Institute for Research into Economic Systems, Kahlaische Strasse 10, 07745 Jena, Germany, Fax: +49 3641 686 622, email: *gueth@mpiew-jena.mpg.de*.

‡Department of Economics, Spandauer Strasse 1, 10178 Berlin, Germany, Fax +49 30 2093 5704, email: *wmueller@wiwi.hu–berlin.de*.

§Recanati Graduate School of Business Administration, Tel Aviv University, Ramat Aviv, Tel Aviv, 69978, Israel, email: *spiegel@post.tau.ac.il*

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. 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). Recently however, Bagwell (1995) showed that if the communication channel is not perfectly reliable the leader's strategic advantage vanishes completely. More precisely, he showed that under a full-support noise structure where with a small probability followers may observe any action from the leader's strategy set, the pure strategy outcome of the sequential-move game coincides with the Nash equilibrium outcome of the associated simultaneous-move game.[1] This result casts a doubt on the empirical relevance of the vast literature in such diverse fields as macroeconomics, international trade, and industrial organization, that builds on the idea that commitment may confer a strategic advantage on first-movers. Moreover, it suggests that the widely used concept of subgame perfection may be highly non-robust and appropriate only in extreme situations in which followers perfectly observe the leaders' actions.

In this paper we report the results of a series of experiments intended to study the strategic behavior of leaders and followers in settings in which followers imperfectly observe the leaders' actions. Instead of using the full-support noise structure that Bagwell (1995) assumes, we use an alternative, "all-or-nothing" noise structure, whereby followers either perfectly observe the leaders' actions, or else they observe nothing. The main difference between the two noise structures is that under the all-or-nothing structure, followers are fully aware of whether they have observed the leader's true action or not, whereas under the full-support structure, followers always doubt whether they have observed the leader's true action or some other action. The all-or-nothing noise structure is widely used in the literature (e.g., Laffont and Tirole, 1993, and Rubinstein, 1989) and has the advantage of restoring, at least partially, the first-mover advantage (Chakravorti and Spiegel,1993).[2]

[1] Intuitively, the follower does not observe the leader's action directly and therefore plays a best-response against his belief about the leader's action. In equilibrium, the follower's belief is correct so the follower plays a best-response against the leader's equilibrium action. Since this best-response is independent of the leader's *true action*, the leader has an incentive to play a best-response against the follower's action and hence the outcome coincides with the equilibrium outcome of the simultaneous move game.

[2] Several papers show that the first-mover advantage of leaders can be restored even under the full-support noise structure. van Damme and Hurkens (1997) show that the noisy-leader game also admits a mixed strategy equilibrium that converges to the subgame perfect equilibrium as the noise vanishes. Moreover, this equilibrium is selected by

Specifically, under the all-or-nothing noise structure, the equilibrium outcome shifts continuously with the probability that followers will observe the leaders' actions from the Cournot outcome (the Nash equilibrium of the simultaneous move game) to the Stackelberg outcome (the subgame perfect equilibrium of the sequential-move game). This continuity stands in stark contrast to the situation under the full-support noise structure where the equilibrium outcome jumps discretely from the Stackelberg outcome when there is no noise to the Cournot outcome when there is even the slightest amount of noise. The continuity property implies that as intuition suggests, first-movers enjoy a greater strategic advantage when their actions are observed with less noise. Moreover, from an experimental point of view it has the advantage of providing us with a continuous measure of the impact of noise on the ability of leaders to gain a first-mover advantage.

We ran 5 sessions, each with 6 leaders and 6 followers who were randomly matched to play a sequential quantity-setting duopoly game. Sessions lasted for 30 rounds, with leaders and followers being randomly matched in pairs at the beginning of every round. The 5 sessions differed only in the probability with which followers observed the leaders' quantity choices (i.e., the noise level). We conducted one "Stackelberg" session in which followers were always informed about the leaders' quantities (Noise0), one "Cournot" session in which followers were never informed about the leaders' quantities (Noise100), and three sessions in which followers were informed about leaders' quantities with probabilities 0.25, 0.50, and 0.75 (Noise25, Noise50, and Noise75, respectively). In the latter three sessions, each follower was informed in some rounds but uninformed in others.

Our main findings can be summarized as follows:

- **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. The tendency of leaders to underproduce became more pronounced as sessions progressed.

- **Overproduction by uninformed followers:** On average, uninformed followers chose quantities above their equilibrium quantities in all treatments, although the deviation from

a "reasonable" selection criterion. Maggi (1999) considers a noisy-leader game in which the leader has private information about his own payoff. He shows that since the leader has an incentive to choose a "large" action to signal his type, the equilibrium outcome must be close to the subgame perfect equilibrium outcome when the noise level is small. Oechssler and Schlag (1997) examine an evolutionary model of the noisy-leader game and show that continuous best-response dynamics select the subgame perfect equilibrium outcome, although under alternative selection dynamics, the Nash equilibrium outcome is selected.

the equilibrium quantities was smaller in treatments with higher levels of noise. This tendency to overproduce remained stable throughout a session.

- **Overreaction by informed followers:** On average, informed followers chose quantities above their best-response to the leaders' quantities. Their estimated reaction function had a smaller intercept and was flatter than the equilibrium best-response function. The overreaction of informed followers was stronger the bigger was the gap between a leader's quantity and the Cournot output.

The tendency of leaders to underproduce and the tendency of informed followers to overreact are similar to the findings in Huck, Müller, and Normann (2001). They ran experiments on sequential duopoly games with quantity competition and found that on average, leaders choose 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.[3]

A closer look at the followers' behavior reveals the following:

- **Best-response was the modal behavior for followers:** The modal behavior of both informed and uninformed followers is to respond optimally to the leaders' quantities. As expected, informed followers played a best-response more than twice as often as uninformed followers (54% of all cases vs. 25% for uninformed followers).

- **Time trends:** As sessions progressed, informed followers played a best-response less often, whereas uninformed followers played a best-response more often.

- **Over- and underreactions:** When informed followers did not play a best-response, they almost always overreacted to the leaders' output. In contrast, uninformed followers underreacted in roughly 30% and overreacted in 45% of all cases.

- **Persistence of followers' behavior:** Followers who overreacted (underreacted) in round $t-1$ also tended to overreact (underreact) in round t. The level of persistence was even higher if a follower was uninformed in period t.

[3] Weimann, Yang, and Vogt (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-a-vis followers.

Taken together, these observations suggest that the soft behavior of leaders may have been a rational response to the aggressive behavior of followers.[4] In particularly, 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.[5] 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 "learned" that small overreactions are not very costly for them 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 deserve to be punished. When uninformed, followers accommodated the leaders' behavior even 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.

Our final findings are about the effect of noise on the behavior of leaders and followers:

- **Monotonicity of output in the level of noise:** With the exception of the Noise50 treatment, the leaders' quantities were monotonically decreasing with the level of noise, and with the exception of the Noise100 treatment (the Cournot treatment), the uninformed followers' quantities are monotonically increasing with the level of noise.

- **Impact of noise on the tendency of followers to over- and underreact:** The behavior of informed followers was not affected by the level of noise whereas the behavior of uninformed followers was affected by the noise level.

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

[5] To see why, consider a quantity-setting model with two firms, A and B, producing a homogenous 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 $(\widehat{q}^A, \widehat{q}^B)$. A small deviation of firm B from $\widehat{q}^B$ lowers π^A by $\frac{\partial \pi^A}{\partial q^B} = P'(\widehat{q}^A + \widehat{q}^B)\widehat{q}^A$, but due to the envelop theorem, it only has a negligeble effect on π^B.

Both results are consistent with the theory. Monotonicity is consistent with the fact that under the all-or-nothing noise structure, the equilibrium outcome is a convex combination of the Stackelberg outcome and the Cournot outcome. The result that the behavior of informed followers was not affected by the noise level is consistent with the fact that in equilibrium, informed followers should play a best-response against the leaders' choices irrespective of the ex-ante probability with which they become informed.

Apart from the Huck, Müller and Normann (2001) paper mentioned above, we are aware of only two other experiments on sequential oligopolies. Kübler and Müller (2000) consider a sequential differentiated products duopoly markets with price competition. Unlike with quantity competition, their setting features a second-mover advantage. Moreover, unlike in our experiments, they asked followers to specify a complete response function (strategy method) rather than a single action. The other study is Huck and Müller (2000), who like us, consider a noisy-leader game. The difference is that while we consider a game with a large strategy set and an all-or-nothing noise structure, they consider a 2×2 game with a full-support noise structure. With small noise levels, followers seemed to ignore the noise and played a best-response against the observed leader's action even though with some probability this may have been the "wrong" action. Leaders quickly learned to exploit this tendency and played the Stackelberg leader's quantity. With high levels of noise, however, leaders played their Stackelberg quantities only half of the time. 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, 1990; Fershtman, Judd, and Kalai, 1991; and Fershtman and Kalai, 1997) and experimentally (Schotter, Zheng, and Snyder, 2000, and Fershtman and Gneezy, 2001). However, in that context, the imperfection is in the observation of the contract that one player offers to another who chooses an action on the player's behalf 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. An Appendix contains the written instructions (translation) that were given to the subjects as well as the payoff matrix that the subjects were using to make decisions.

2 Experimental design

2.1 A noisy-leader game with all-or-nothing noise structure

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

$$\pi^A = \left(a - (q^A + q^B)\right) q^A, \qquad \text{and} \qquad \pi^B = \left(a - (q^A + q^B)\right) q^B, \tag{1}$$

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 an all-or-nothing signal about q^A. With probability $1 - \varepsilon$ the signal perfectly reveals q^A. With probability ε the signal reveals nothing (firm B then 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

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

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

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

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:

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

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

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

while an informed firm B will choose the quantity

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

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 q^A whatever the value of q^A is. In the latter case, firm B does not observe q^A and hence it chooses a quantity that depends on firm A's equilibrium quantity but not on firm A's actual quantity. The implication of this behavior is that 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 $a/2$ units which is the monopoly output, while Firm B chooses $a/4$ 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 $a/3$ units. As ε increases from 0 to 1, firm A's quantity falls continuously from the Stackelberg leader's quantity of $a/2$ to the Cournot quantity of $a/3$, whereas firm B's quantity increases continuously from the Stackelberg follower's quantity of $a/4$ to the Cournot quantity of $a/3$.[6] Consequently, firm A enjoys a smaller first-mover advantage as ε increases towards 1.

At an intuitive level, when $\varepsilon = 0$, firm A commits itself to an aggressive behavior by choosing $a/6$ units more than the Cournot output ($a/2$ units instead of $a/3$ units). This gives firm A a strategic advantage vis-a-vis firm B because it induces firm B to cut its quantity by $a/12$ unit below the Cournot level ($a/4$ units instead of $a/3$ units). Firm A's aggressive behavior means that it does not play a best response against firm B's choice: given that firm B chooses $a/4$ units, firm A would have liked to cut its output from $a/2$ units to $\frac{a-a/4}{2} = \frac{3a}{8}$ units. However, it is precisely because firm A cannot alter its commitment to produce $a/2$ units that it gains a strategic advantage vis-a-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 $q^A = a/2$ and would choose $q^B = a/4$ units, firm A would have actually preferred to produce only $3a/8$ units. In equilibrium of course, firm B fully anticipates this and hence, some of the commitment power of firm A 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.

[6] 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.

2.2 Experimental implementation

The computerized experiments on the noisy leader game were conducted at Humboldt University using the software tool kit *z-Tree* (Fischbacher, 1999). We ran 5 sessions, each with 12 different subjects. Subjects were students from various departments at Humboldt University, mainly from Economics, Business Administration and Law. They 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. The average earnings were DM 33.87 which was about $15 or 17.35 Euros at the time of the experiment.

Upon arrival in the lab, subjects were assigned a computer screen and received written instructions in German (an English translation appears in the Appendix). 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 subjects were assigned the role of firm B (a follower). Players' roles were kept fixed during the entire session. We implemented 5 treatments of the noisy-leader game. 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, ε was set 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, 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!"[7] 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.

Apart from written instructions, subjects also received a payoff matrix (see the Appendix)

[7] 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, Huck, and Rapoport (1998) and Müller (2001) provide experimental evidence showing 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, Knez, and Weber, 1996; and Rapoport, 1997).

that specified the profits of firms A and B for each possible outcome of the game. The profits were expressed in terms of a fictitious currency called "Taler" and were then converted to DM according to a prespecified exchange rate (see below). 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 $\{13, 14, ..., 32\}$.[8] 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.

2.3 Pilot sessions and exchange rates from Taler to DM

Prior to the experiments reported below, we ran 2 pilot sessions for treatment NOISE0 (the Stackelberg treatment). Both sessions followed the same rules as described above with two differences. In the first pilot session, all participants had the same exchange rate from Talers to DM, and this was commonly known. The results were such that average quantity of leaders was only 21.32 units, compared with 30 units predicted by theory. More importantly, the estimated reaction function of followers against q^A (using a simple OLS regression), was $BR^B(q^A) = 13.81 + 0.38q^A$. This function differs significantly from the equilibrium best-response function which is $q^B = 30 - 0.5q^A$. The main difference is that the slope is *positive* rather than *negative*, implying that followers behaved as if strategies were *strategic complements* rather than *strategic substitutes*. This is diametrically opposed to the theory since one of the main features of the Cournot/Stackelberg models is that strategies are strategic substitutes. We believe that the result was 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.[9]

In order to neutralize these fairness considerations as much as possible and focus on other aspects of noisy leadership, we ran a second pilot session in which subjects switched roles in every round so that each subject enjoyed a first-mover advantage in 15 out of the 30 rounds.[10] As

[8] 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 $\{1, 2, ..., 20\}$. In addition, we ensured that the best-response of informed followers was single-valued by subtracting one Taler in several entries in the payoff matrix.

[9] In our design, where the Stackelberg outcome is $(30, 15)$, an increase of 1 unit in firm B's quantity lowers the price from $60 - 45 = 15$ to $60 - 46 = 14$. Consequently, firm A's profits falls by $1 \times 30 = 30$ Talers. On the other hand, firm B also produces an extra unit and sells it at the new price of 14 so its profits falls by only $1 \times 15 - 14 = 1$ Taler. Hence, firm B hurts firm A substantially at a small loss to itself.

[10] Prior to odd rounds, 6 subjects were assigned the role of leaders and 6 were assigned the roles of followers; the leaders and followers were then randomly matched in pairs. In even rounds, the roles were reversed and leaders and

in the first pilot session, all subjects had the same exchange rate from Talers to DM and this was commonly known. The results in this session were qualitatively similar to the first pilot session. The average quantity of leaders was merely 20.09 units, close to the Cournot output of 20 units, and the estimated reaction function of followers was (again, using a simple OLS regression) $q^B = 17.98 + 0.11q^A$.

Following the two pilot sessions we decided to assign to each subject an individual and confidential exchange rate from Talers to DM. The exchange rates were randomly selected from the set $\{300, 320, 330, 340, 350\}$ (e.g., 300 Talers = DM 1). Before a session started, subjects saw on computer screens a personal message informing them about their own exchange rates, but *not* about their rivals' exchange rates or the range of possible exchange rates. We felt that if subjects would not know each other's exchange rates, fairness considerations would play a smaller role. To ensure that all subjects were aware of this feature, the personal messages included the following line: "Keep in mind that other participants do not necessarily have the same exchange rate."

3 Results

With 5 sessions (one for each treatment) of 30 rounds each, and 6 leaders and 6 followers in each session, we have $5 \times 6 \times 30 = 900$ leaders' quantity choices, and $5 \times 6 \times 30 = 900$ followers' quantity choices. After an initial review of the data, we decided to exclude subject 47 who played as a follower in treatment NOISE75.[11] This left us with 870 observations on followers' choices, of which 434 were made by informed followers who saw q^A before choosing their own quantities, and 436 were made by uninformed followers who were only told that the leader had already chosen q^A but were not told what the value of q^A was.

In Table 1 below, 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, $\widehat{q}^A$, and uninformed followers, $\widehat{q}^{B,\text{uninfo}}$. In addition, we report for informed

followers were again randomly matched in pairs.

[11] This subject adopted a highly idiosyncratic predatory behavior by choosing a quantity of 32 units in each of the first 4 rounds and a quantity of 30 units in all other rounds. In the post-experimental questionnaire he explained his behavior as follows: "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." Since this behavior is self-explanatory and requires no further analysis, we simply decided to ignore subject 47 when analyzing the *followers'* behavior. Incidently, the suboptimal behavior of subject 47 meant that he ended the experiment with the lowest monetary payoff among all 60 subjects.

Treatment	Actual and equilibrium behavior					
	mean q^A (N = 900)	$\widehat{q}^A$	mean $q^{B,\text{uninfo}}$ (N = 436)	$\widehat{q}^{B,\text{uninfo}}$	mean $\Delta_{i,t}^{\text{info}}$ (N = 434)	$\widehat{\Delta}_{i,t}^{\text{info}}$
NOISE0 (Stackelberg)	22.68 (3.93)	30	—	—	1.26 (2.68)	0
NOISE25	21.98 (2.84)	26.7	19.77 (3.09)	16.7	2.05 (3.28)	0
NOISE50	22.90 (3.10)	24	20.51 (2.69)	18	1.36 (2.55)	0
NOISE75	20.58 (2.17)	21.8	20.89 (2.17)	19.1	1.24 (3.54)	0
NOISE100 (Cournot)	19.32 (2.18)	20	20.36 (2.73)	20	—	—

Table 1: Actual and equilibrium behavior of leaders and followers. Standard deviations appear in parentheses.

followers 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 more observations on informed followers in treatment NOISE25 (127 observations) than in treatments NOISE50 and NOISE75 (86 and 41 observations, respectively), whereas in treatment NOISE75, we have more observations for uninformed followers than in treatments NOISE25 and NOISE50 (139 observations versus 53 and 94).

Table 1 shows that on average, informed followers overreacted to the leaders' quantities and chose quantities that exceeded their best-response by $1.24 - 2.05$ units. The table also shows that on average, leaders chose smaller quantities, while uninformed followers chose larger quantities than the theory predicts. Nonetheless, it appears that the comparative statics properties of the noisy-leader model are by and large reflected in the data: with the exception of treatment NOISE50, the average quantity of leaders decreases with the level of noise, and with the exception of treatment NOISE100, the average quantity of uninformed followers increases with the level of noise. To examine whether the effect of noise on the behavior of leaders and uninformed followers

	Noise0	Noise25	Noise50	Noise75	Noise100
Noise0	—	—	—	—	—
Noise25	0.029	—	—	—	—
Noise50	0.529	0.002	—	—	—
Noise75	0.000	0.000	0.000	—	—
Noise100	0.000	0.000	0.000	0.000	—

Table 2: p-values of (two tailed) pairwise cross-treatment differences in the means of q^A

	Noise25	Noise50	Noise75	Noise100
Noise25	—	—	—	—
Noise50	0.065	—	—	—
Noise75	0.011	0.218	—	—
Noise100	0.131	0.571	0.066	—

Table 3: p-values of (two tailed) pairwise cross-treatment differences in the means of $q^{B,\text{uninfo}}$

is significant, we conducted pairwise cross-treatment comparisons of the means of q^A and $q^{B,\text{uninfo}}$. The resulting (two-tailed) p-values are reported in Tables 2 and 3 (treatment Noise0 is missing from Table 3 because in this treatment all followers were informed).[12]

Table 2 reveals that with the exception of treatment Noise50 vs. treatment Noise0, all pairwise cross-treatment differences in the means of q^A are highly significant. Table 3 shows that the impact of noise on the behavior of uninformed followers is less clear-cut. Differences between the means of $q^{B,\text{uninfo}}$ in treatment Noise50 on the one hand and Noise75, and Noise100 on the other hand (second column in Table 3) are insignificant, while all other cross-treatment differences are weakly significant.

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.[13]

[12]The p-values were obtained by running OLS regressions across observed quantities, using the treatment as a dummy. For example, to test for differences in leaders' behavior in treatments Noise0 and Noise25, we estimated the equation $q_i^A = \beta_0 + \beta_1 TREAT + \varepsilon_i$, where $TREAT$ is a dummy variable equal to 0 in treatment Noise0 and equal to 1 in treatment Noise25, and ε_i is an error term. The estimated value of β_1 then represents the difference in means across the two treatments. We use White's (1980) robust standard errors adjusted for possible non-independence of observations within treatments to estimate the covariance matrix.

[13]Due to the repeated random matchings within each session, only sessions can be regarded as independent ob-

3.1 Leaders' behavior

In order to assess leaders' behavior, we estimate the following OLS regression:

$$\begin{aligned} q^A_{j,t} &= \alpha_0 + \alpha_1 D_{25} + \alpha_2 D_{50} + \alpha_3 D_{75} + \alpha_4 D_{100} + \beta_1 \Delta^-_{j,t-1} + \beta_2(\Delta^{+,\text{info}}_{j,t-1}) + \beta_3(\Delta^{+,\text{uninfo}}_{j,t-1}) \\ &\quad + \gamma_1 T^2_t + \gamma_2 T^3_t + \delta_1 \Delta^2_{j,t-1} + \delta_2 \Delta^3_{j,t-1} + \sum_{j\in\Lambda} \tau_j D^A_j + \eta, \end{aligned} \tag{7}$$

where $q^A_{j,t}$ is leader j's quantity in round t, Λ is the set of leaders,[14] and η is an error term. The independent variables in the regression equation 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 on the leaders' behavior relative to this benchmark.

- $\Delta_{j,t} \equiv q^B_{j,t} - BR^B(q^A_{j,t})$ 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 $\Delta^{+,\text{info}}_{j,t-1}$ ($\Delta^{+,\text{uninfo}}_{j,t-1}$) 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 $\Delta^{+,\text{info}}_{j,t-1}$ and $\Delta^{+,\text{uninfo}}_{j,t-1}$, respectively, include overreactions by informed and uninformed followers in the previous round (cases in which followers chose quantities above their best-response). We include these variables in order to examine whether leaders modified their behavior in a given round based on their experience in the previous round. We make a distinction between $\Delta^{+,\text{info}}_{j,t-1}$ and $\Delta^{+,\text{uninfo}}_{j,t-1}$ because leaders are likely to interpret the former as deliberate attempts by informed followers to "punish" them while interpreting the latter as reflecting the difficulty of uninformed followers to predict the leaders' choices. We do

servations according to rigorous statistical standards. This would have meant to run at least four sessions for each treatment, at least for the purpose of studying the impact of noise on the behavior of leaders and followers. However, we are quite confident that more sessions would not have questioned our main conclusions. In the regression analysis we deal with the possible interdependece of observations within each session by using the cluster option in STATA (see footnote 16 below for more details).

[14] The set Λ includes the index numbers of all subjects who played the role of leaders (firm A): $\Lambda = \{j \mid j = 1, ..., 6, 13, ..., 18, 25, ..., 30, 37, ..., 42, 49, ..., 54\}$.

not make a similar distinction for underreaction because almost all underreactions came from uninformed followers.

- T_t^2 and T_t^3, respectively, are "third" 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 $1-10$ serves as the reference group, with the third dummies 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.[15] This restriction implies that 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 are presented in Table 4. 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 highly significant. Regressions L1-L3 differ in terms of the independent variables that are included. In addition, since regression L3 includes lagged variables ($\Delta_{j,t-1}^-$ and $\Delta_{j,t-1}^+$), we lose $5 \times 6 = 30$ observations on the leaders' behavior in round 1 of each treatment. Regression L4 uses the cluster option provided by the statistical software package 'STATA' to account for possible interdependencies in the error terms within each treatment cluster. These interdependencies might arise because in each treatment, the same subjects were matched for 30 rounds. The cluster 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 cluster.[16]

We already saw in Table 1 that leaders tended to underproduce relative to their equilibrium quantities. Since the coefficients γ_1 and γ_2 are significant, it follows from regressions L2 and L3 that

[15] This restriction was first proposed by Suits (1984). A discussion on the use of this approach in experimental economics appears in Königstein (2000).

[16] Let $G_1, ..., G_M$ be the M clusters specified in the cluster option. Then, the formula for robust variance calculation used by the cluster option is $V_{cluster} = \frac{(N-1)M}{(N-k)(M-1)}(X'X)^{-1}\left(\sum_{i=1}^{M} u_i' u_i\right)(X'X)^{-1}$, 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 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).

	Regression L1	Regression L2	Regression L3	Regression L4 (with cluster option for treatment clusters)
α_0	22.683^{***}	23.092^{***}	23.172^{***}	23.172^{***}
	(113.446)	(98.144)	(94.668)	(142.618)
α_1 (D_{25})	-0.706^{**}	-0.706^{**}	-0.758^{***}	-0.758^{***}
	(-2.495)	(-2.509)	(-2.704)	(-41.554)
α_2 (D_{50})	0.217	0.217	0.331	0.331^{***}
	(0.766)	(0.770)	(1.184)	(23.013)
α_3 (D_{75})	-2.106^{***}	-2.106^{***}	-2.284^{***}	-2.284^{***}
	(-7.446)	(-7.487)	(-8.060)	(-68.854)
α_4 (D_{100})	-3.367^{***}	-3.367^{***}	-3.118^{***}	-3.118^{***}
	(-11.906)	(-11.971)	(-10.763)	(-46.894)
γ_1 (T_t^2)		-0.503^{**}	-0.594^{***}	-0.594^{*}
		(-2.311)	(-2.715)	(-2.432)
γ_2 (T_t^3)		-0.723^{***}	-0.820^{***}	-0.820^{**}
		(-3.321)	(-3.734)	(-3.539)
β_1 ($\Delta^{-}_{j,t-1}$)			0.217^{***}	0.217^{**}
			(2.610)	(3.146)
β_2 ($\Delta^{+,\text{info}}_{j,t-1}$)			0.050	0.050^{**}
			(1.180)	(3.022)
β_3 ($\Delta^{+,\text{uninfo}}_{j,t-1}$)			0.046	0.046
			(1.132)	(1.325)
adj. R^2	0.302	0.310	0.333	0.358
Number of observations	900	900	870	870

Note: Parameter estimates for subject dummies not shown. t-values in parentheses.

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

Table 4: Results of the leaders' regressions

this tendency became stronger as sessions progressed. For instance, while the average quantity of leaders in treatment NOISE0 is around 23.1 units in the first 10 rounds, it falls by about $0.50-0.59$ units in rounds $11-20$ and by $0.72-0.79$ units in rounds $21-30$. Recalling from Table 1 that followers tended to overproduce, the decreasing trend in the leaders' output may have been a rational response to the followers' aggressive behavior. Regressions L3 and L4 reveal however that this response was not immediate as leaders barely reacted to followers' overreactions in the immediately preceding round. If anything, Regression L4 shows that whenever informed followers were aggressive and overreacted, leaders tended to raise their quantities in the immediately following round by an average of 0.05 units for each unit of overreaction by the followers (in regression L3, this effect is not significant).[17] In addition, Regressions L3 and L4 show that whenever followers were soft and underreacted, leaders tended to "reciprocate" by lowering their own output in the immediately following round by an average of 0.217 units for each unit of underreaction by the followers (since $\Delta^-_{j,t-1}$ is negative by definition, $\beta_1 > 0$ implies a reduction in q^A).[18] Overall, the estimated values of β_1, β_2, and β_3 suggest that leaders tended to play soft after being "nicely" treated, but showed only a weak tendency to play aggressively after being "mistreated."[19]

Turning to the effect of noise on leaders' behavior, regressions L1-L4 show that except for α_2, the coefficients of the treatment dummies are negative and highly significant, with $|\alpha_1| < |\alpha_3| < |\alpha_4|$. Hence, with the exception of treatment NOISE50, the leaders' output falls when there is more noise, as the theory predicts. Although Tables 1 and 2 already showed a similar trend, here the result is "cleaner" as we also control for the idiosyncratic behavior of individual leaders, for time trends (regressions L2-L4), and for responses to followers' behavior (regressions L3-L4).[20]

[17] We also tried to examine whether there were any time trends in the the leaders' response to under- and overreactions by followers. This was done by multiplying $\Delta^-_{j,t-1}$ and $\Delta^+_{j,t-1}$ by the "third" dummies, T^2_t and T^3_t, and including the new variables in the regression. However, the coefficients of the new variables were not significant. In addition, we also included a 2-period lagged value of $\Delta_{j,t}$ in the regression in order to detect delayed responses by leaders to the followers' behavior. This variable however was highly insignificant.

[18] 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 will meet the same follower again in the next round.

[19] 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 a "mean" or "unkind" behavior because the follower sacrifices a monetary payoff by overproducing in order to hurt the leader.

[20] The coefficient α_2 is not significant in regression L3 but is significant and positive in regression L4. The large $t - value$ of α_2 in regression L4 could be due to negative intra-cluster correlations which can lead to smaller estimated standard errors with the cluster option relative to the OLS standard errors. See http://www.stata.com/support/faqs/stat/cluster.html

It should also be noted that except for treatment NOISE50, the deviation of the leaders' behavior from equilibrium behavior shrinks when there is more noise. For instance, Regression L2 shows that leaders produced in treatment NOISE0 23% less than the theory predicts (23.1 versus 30 units predicted by theory), 16.1% less in treatment NOISE25 ($23.092 - 0.706 = 22.386$ versus 26.7 units), 3.7% less in treatment NOISE75 ($23.092 - 2.106 = 20.986$ versus 21.8 units), and only 1.3% less in treatment NOISE100 ($23.092 - 3.367 = 19.725$ versus 20 units).

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) The leaders' tendency to underproduce became more pronounced as sessions progressed.

(iii) With the exception of treatment NOISE50*, the leaders' quantities were decreasing with the level of noise.*

3.2 Informed Followers' behavior

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

$$q_{i,t}^{B,\text{info}} = \alpha_0 + \beta q_{i,t}^{A} + \gamma_1 T_t^2 + \gamma_2 T_t^3 + \sum_{i \in F} \tau_i D_i^B + \sum_{i \in F} \theta_i D_i^B q_{i,t}^A + \eta, \tag{8}$$

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, T_t^2 and T_t^3 are "third" dummies defined as in the leaders' regression, D_i^B are follower-specific dummies, F is the set of followers,[21] 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. We restrict the sum of the τ_i's and the sum of the θ_i's to 0. This restriction implies that the estimated coefficients α_0 and β represents the intercept and slope of

[21] The set F includes the index numbers of all subjects who played the role of followers (firm B), excluding subject 47 whose quantity choices were omitted from the data (see the discussion at the begining of Section 4): $F = \{i \mid i = 7, ..., 12, 19, ..., 24, 31, ..., 36, 43, ..., 46, 48, 55, ..., 60\}$.

Independent variable	Regression INF1	Regression INF2	Regression INF3 (with cluster option for treatment clusters)
α_0	23.947***	23.957***	23.957***
	(37.745)	(36.639)	(48.56)
β ($q^A_{i,t}$)	−0.159***	−0.160***	−0.160**
	(−5.64)	(−5.655)	(−3.891)
γ_1 (T^2_t)		0.1	0.1
		(−0.439)	(−0.393)
γ_2 (T^3_t)		−0.056	−0.056
		(−0.243)	(−0.156)
adj. R^2	0.525	0.523	0.568
Number of observations	870	870	870

Note: Parameter estimates for subject dummies not shown. t-values in parentheses.

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

Table 5: Results of the informed followers' regressions. t-values in parentheses

the "average" best-response of informed followers, while the coefficients τ_j and θ_j 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 the observed leaders' quantities irrespective of the ex ante likelihood that they will be informed. In Subsection 3.4 we will show evidence that indeed there were no cross-treatments differences in the behavior of informed followers.

The results of the informed followers' regressions are presented in Table 5. As with the leaders' regression, we do not report the coefficients of the follower-specific dummies but note that most of them were highly significant.

Regressions INF1-INF3 show that relative to the equilibrium best-response function, the estimated reaction function of informed followers has a lower intercept (23.95 compared with 30) and its slope is much flatter (−0.16 compared with −0.5). Regressions INF2 and INF3 show that the estimated reaction function of informed followers remained stable and did not change over the course of the sessions as the coefficients of T^2_t and T^3_t are 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 throughout each session.*

3.3 Uninformed Followers' behavior

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

$$\begin{aligned} q_{i,t}^{B,\text{uninfo}} &= \alpha_0 + \alpha_1 D_{50} + \alpha_2 D_{75} + \alpha_3 D_{100} + \beta q_{i,t-1}^{A} \\ &\quad + \gamma_1 T_t^2 + \gamma_2 T_t^3 + \sum_{i \in F} \tau_i D_i^B + \sum_{i \in F} \theta_i D_i^B q_{i,t-1}^{A} + \eta, \end{aligned} \tag{9}$$

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 $t-1$. We include this variable in the regression to examine how uninformed followers adjusted their behavior on the basis of their 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 the leaders' choice ex post). 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 results of the followers' regressions are presented in Table 6. Regressions UNINF1-UNINF2 show that uninformed followers reacted with a lag to the leaders' choices and became more aggressive as the leader with whom they were matched in the previous round chose a larger quantity (in regression UNINF3 this effect is not significant). Moreover, Regressions UNINF1-UNINF3 show that uninformed followers tended to choose larger quantities in treatments with larger noise although the relationship is not monotonic. For instance, in Regression UNINF3, the estimated value of α_2 is larger than the estimated value of α_3 implying that holding fixed the leaders' quantities in the previous round, uninformed followers chose larger quantities in treatment Noise75 than in Noise100. Regressions UNINF2-UNINF3 show that the behavior of uninformed followers did not change as sessions progressed as the coefficients of T_t^2 and T_t^3 are highly insignificant.

Observation 3 *The behavior of uninformed followers had the following features:*

(i) Uninformed followers become more aggressive the higher the quantity chosen by the leader in the immediately preceding round.

Independent variable	Regression UNINF1	Regression UNINF2	Regression UNINF3 (with cluster option for treatment clusters)
α_0	16.87***	16.848***	16.848***
	(15.357)	(14.801)	(11.326
β ($q^A_{i,t-1}$)	0.141***	0.142***	0.142
	(2.966)	(2.963)	(2.222)
α_{50} (D_{50})	0.526	0.513	0.513***
	(1.288)	(1.249)	(15.156)
α_{75} (D_{75})	1.018***	1.011***	1.011***
	(2.597)	(2.57)	(9.223)
α_{100} (D_{100})	0.63*	0.625*	0.625**
	(1.638)	(1.612)	(3.423)
γ_1 (T_t^2)		-0.046	-0.046
		(-0.179)	(-0.146)
γ_2 (T_t^3)		0.068	0.068
		(0.26)	(0.377)
adj. R^2	0.415	0.412	0.474
Number of observations	841	841	841

Note: Parameter estimates for subject dummies not shown. t-values in parentheses.

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

Table 6: Results of the uninformed followers' regressions. t-values in parentheses

(ii) The behavior of uninformed followers remained stable throughout each session.

(iii) Uninformed followers tended to choose larger quantities in treatments with a higher level of noise although the relationship is not monotonic.

3.4 Followers' over- and underreactions

In this subsection we study the follower's 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 the leader's quantity in the same round. We begin by looking at the distribution of $\Delta_{i,t}$ for informed followers (the solid line) and uninformed followers (the dashed line) across all treatments.

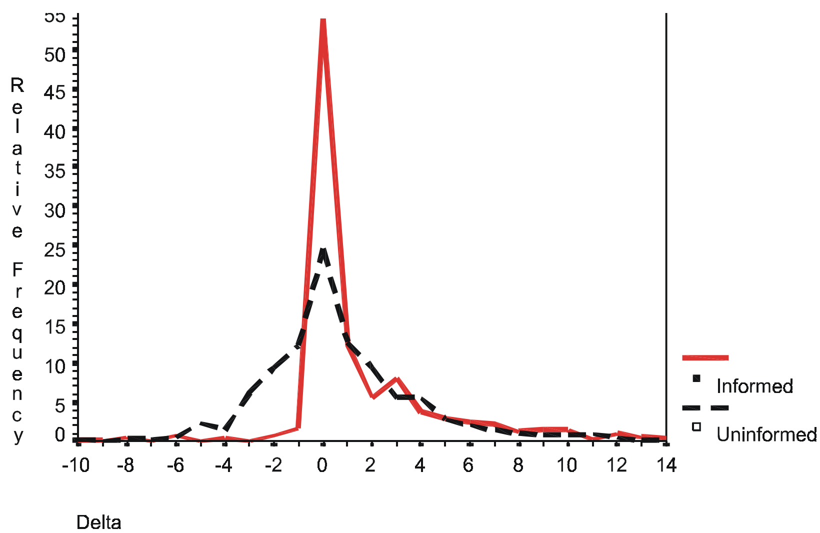

Figure 1: The distribution of $\Delta_{i,t}$ for informed followers (the solid line) and uninformed followers (the dashed line), pooled across treatments.

Figure 1 shows that the modes of both distributions are equal to 0, implying that irrespective of whether followers were informed or uninformed, their modal behavior 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 roughly 54% of all cases, compared with only 25% for uninformed followers. The figure also shows that the two distributions virtually coincide for $\Delta > 0$, implying in particularly that informed and uninformed followers tended to overreact to q^A at about the same frequency (42% of all cases for informed followers and 45% for uninformed). The main difference between the two distributions is

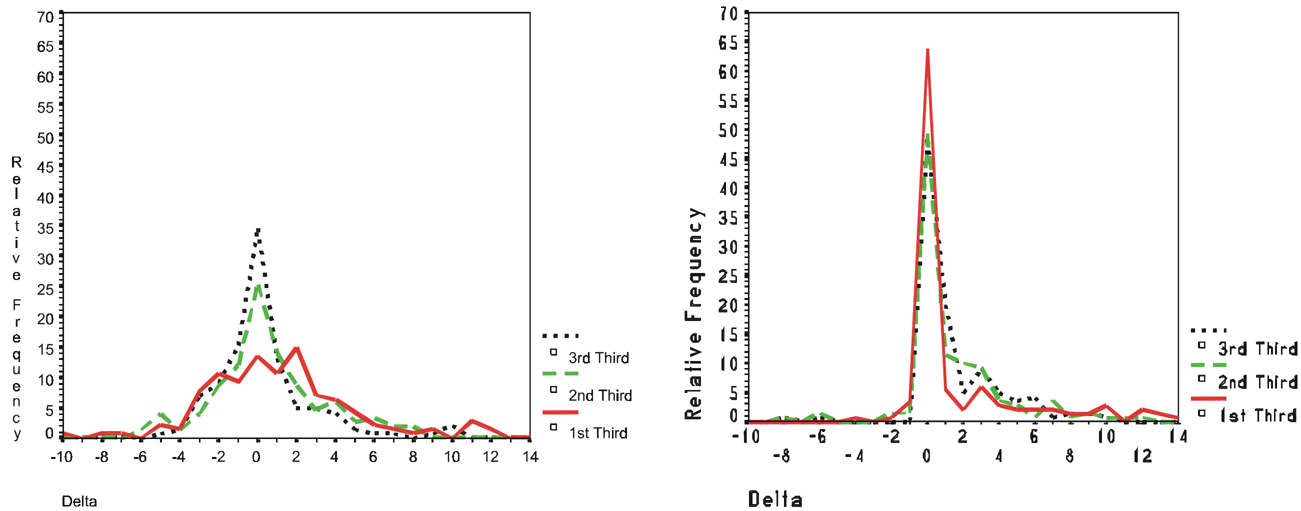

Figure 2: The time trend in the distribution of $\Delta_{i,t}$ for uninformed (left) and informed (right) followers

that while informed followers almost never underreacted to q^A (only 4% of all cases), uninformed followers ended up underreacting to q^A in roughly 30% of all cases.[22]

In Figure 2 we examine the time trend in $\Delta_{i,t}$ by looking at a break down of $\Delta_{i,t}$ to $1st$ third (rounds $1-10$), $2nd$ third (rounds $11-20t$), and $3rd$ third (rounds $21-30$). The left panel in Figure 2 shows that as sessions progressed, the distribution of $\Delta_{i,t}$ for uninformed followers became more concentrated around 0. In particular, uninformed followers played a best-response against q^A, i.e., chose $\Delta = 0$, in about 13% of all cases in rounds $1-10$, 25% in rounds $11-20$, and 34% in rounds $21-30$. These observations suggest that uninformed followers improved their predictions about q^A and "learned" to play best-responses against these predictions. The right panel in Figure 2 shows an opposite trend for informed followers: they played a best-response against q^A less frequently as sessions progressed (about 64% of all cases in rounds $1-10$, but only in 49% in rounds $11-30$). Instead of playing a best-response, informed followers tended to overreact: In rounds $11-20$ they overreacted by 1, 2, or 3 units, each in about 10% of all cases, while in rounds $21-30$ they overreacted by 1, 2, or 3 units, respectively, in 20%, 5%, and 10% of all cases. Bearing in mind

[22] There is a big variance in the behavior of individual followers. For instance, out of 24 followers who were informed (followers in treatments NOISE0, NOISE25, NOISE50, and NOISE75), 4 always played a best-response and another 2 played a best-response in more than 90% of the cases, while 4 played a best-response in less than 10% of all cases.

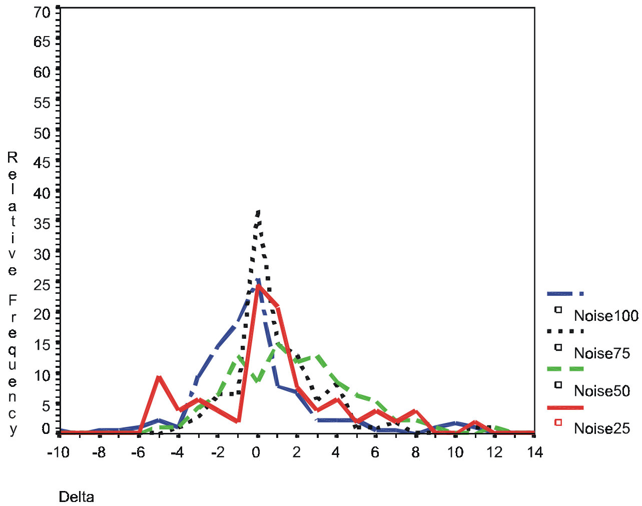

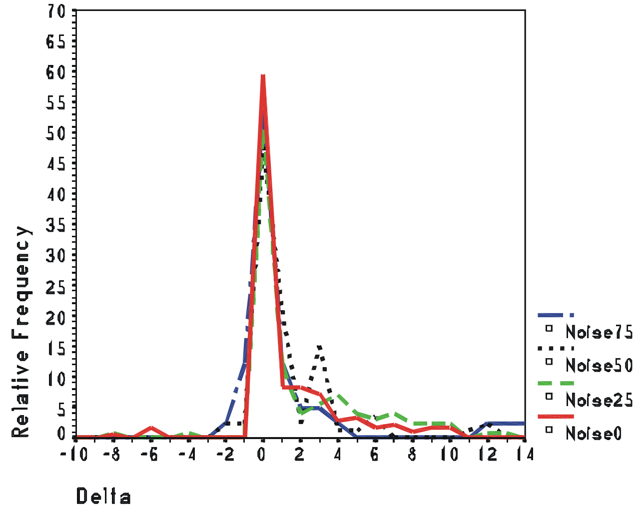

Figure 3: The distribution of $\Delta_{i,t}$ according to treatments for uninformed (left) and informed (right) followers.

that small overreactions by followers hurt leaders substantially at a small personal loss to followers (see Footnotes 5 and 9) and recalling from Subsection 3.1 that the average output of leaders fell as sessions progressed, 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 Figure 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. In addition, the right panel reveals that there are no cross-treatment differences in the behavior of informed followers, suggesting that informed followers were affected only by what they observed but not by the ex-ante probability of this event. And, although the left panel shows several cross-treatment differences in the behavior of uninformed followers, we shall see shortly that these are not statistically significant.

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

$$\begin{aligned}\Delta_{i,t} &= \alpha_0 + \alpha_1 \Delta_{i,t}^{A,\text{info}} + \alpha_2 L_{i,t}^{\text{info}} + \alpha_3 \Delta_{i,t-1}^{A,\text{uninfo}} + \alpha_4 L_{i,t-1}^{\text{uninfo}} \\ &\quad + \rho_1 \Delta_{i,t-1}^{\text{info}} + \rho_2 \Delta_{i,t-1}^{\text{uninfo}} + \sum_{i \in \mathcal{F}} \tau_i D_i^B + \eta, \qquad (10)\end{aligned}$$

where D_i^B, $i \in \mathcal{F}$, are the follower-specific dummies that were defined in the followers' regression, and η is an error term. The definitions of the new independent variables and the reasons for including them in the regression are as follows:

- $\Delta_{i,t}^{A,\text{info}} \equiv q_{i,t}^{A,\text{info}} - \widehat{q}_i^A$ is the gap between the actual quantity of the leader with whom uninformed follower i was matched in round t and the equilibrium quantity of that leader. We include $\Delta_{i,t}^{A,\text{info}}$ 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.

- $L_{i,t}^{\text{info}} \equiv q_{i,t}^{A,\text{info}} - 20$ is the gap between the actual quantity of the leader with whom uninformed follower i was matched in round t and the Cournot output which is 20 units, 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 who produce more than 20 units.

- $\Delta_{i,t-1}^{A,\text{uninfo}}$ 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 quantities or from the 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 presented in the following 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 $\Delta_{i,t}^{A,\text{info}}$, $L_{i,t}^{\text{info}}$, $\Delta_{i,t-1}^{A,\text{uninfo}}$, and $L_{i,t-1}^{\text{uninfo}}$ with their respective breakdowns by treatments.[23]

[23] For instance, $\Delta 0_{i,t}^{A,\text{info}}$, $\Delta 25_{i,t}^{A,\text{info}}$, $\Delta 50_{i,t}^{A,\text{info}}$, $\Delta 75_{i,t}^{A,\text{info}}$ are equal to $\Delta_{i,t}^{A,\text{info}}$ if follower i participated in treatment Noise0, Noise25, Noise50, or Noise75, respectively, and are equal to 0 otherwise (e.g., $\Delta 25_{i,t}^{A,\text{info}} = \Delta_{i,t}^{A,\text{info}}$, if follower i participated in treatment Noise25 and $\Delta 25_{i,t}^{A,\text{info}} = 0$ otherwise). We do not define the variables $\Delta 100_{i,t}^{A,\text{info}}$, $L100_{i,t}^{A,\text{info}}$, $\Delta 0_{i,t}^{A,\text{uninfo}}$, and $L0_{i,t}^{A,\text{uninfo}}$ as in treatment Noise100 all followers were uninformed while in treatment 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, $\widehat{q}_i^A = 20$, implying that $\Delta 100_{i,t-1}^{A,\text{uninfo}} = L100_{i,t-1}^{\text{uninfo}}$.

Independent variable	Regression D1		Regression D2		Regression D3 (with cluster option for treatment clusters)	
α_0	0.401^{***}	(3.109)	0.341^{**}	(2.349)	0.341	(0.673)
α_1 ($\Delta_{i,t}^{A,\text{info}}$)	0.007	(0.285)				
α_2 ($L_{i,t}^{\text{info}}$)	0.358^{***}	(10.384)				
α_3 ($\Delta_{i,t-1}^{A,\text{uninfo}}$)	-0.098^{*}	(-1.927)				
α_4 ($L_{i,t-1}^{\text{uninfo}}$)	0.040	(0.768)				
α_5 ($\Delta 0_{i,t}^{A,\text{info}}$)			0.018	(0.688)	0.018	(0.385)
α_6 ($\Delta 25_{i,t}^{A,\text{info}}$)			-0.057	(-1.183)	-0.057	(-0.884)
α_7 ($\Delta 50_{i,t}^{A,\text{info}}$)			-0.048	(-0.430)	-0.048	(-0.574)
α_8 ($\Delta 75_{i,t}^{A,\text{info}}$)			-0.356	(-1.452)	-0.356^{*}	(-2.624)
α_9 ($L0_{i,t}^{\text{info}}$)			0.350^{***}	(8.262)	0.350^{***}	(7.286)
α_{10} ($L25_{i,t}^{\text{info}}$)			0.598^{***}	(8.402)	0.598^{***}	(8.789)
α_{11} ($L50_{i,t}^{\text{info}}$)			0.232^{***}	(2.625)	0.232^{**}	(2.858)
α_{12} ($L75_{i,t}^{\text{info}}$)			0.869^{***}	(2.937)	0.869^{***}	(5.733)
α_{13} ($\Delta 25_{i,t}^{A,\text{uninfo}}$)			-0.068	(-0.888)	-0.068	(-1.674)
α_{14} ($\Delta 50_{i,t}^{A,\text{uninfo}}$)			-0.160	(-1.596)	-0.160	(-1.838)
α_{15} ($\Delta 75_{i,t}^{A,\text{uninfo}}$)			-0.051	(-0.309)	-0.051	(-0.210)
α_{16} ($\Delta 100_{i,t}^{A,\text{uninfo}}$)			-0.075	(-0.828)	-0.075	(-0.753)
α_{17} ($L25_{i,t}^{\text{uninfo}}$)			-0.057	(-0.524)	-0.057	(-1.989)
α_{18} ($L50_{i,t}^{\text{uninfo}}$)			0.039	(0.461)	0.039	(0.560)
α_{19} ($L75_{i,t}^{\text{uninfo}}$)			0.199	(1.119)	0.199	(0.916)
ρ_1 ($\Delta_{i,t-1}^{\text{info}}$)	0.125^{***}	(2.850)	0.097^{**}	(2.155)	0.097	(0.965)
ρ_2 ($\Delta_{i,t-1}^{\text{uninfo}}$)	0.333^{***}	(7.168)	0.334^{***}	(7.115)	0.334^{***}	(5.496)
adj. R^2	0.355		0.369		0.369	
Number of Observations	841		841		841	

Note: Parameter estimates for subject dummies not shown. t-values in parentheses

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

Table 7: Results of the Δ regressions. t-values in parentheses

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 when we break down the data by treatments as the coefficients $\alpha_5 - \alpha_8$ 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 0.358, implying that for each leader's quantity above 20 units, informed followers overreacted by an average of 0.358 units. This is consistent with the hypothesis that informed followers viewed leaders who were trying to exploit their first-mover advantage as "unfair" and "punished" them by overreacting. Regressions D2-D3 show that the overreaction of informed followers to leaders' quantities above 20 units was present in all treatments as the coefficients $\alpha_9 - \alpha_{12}$ are all highly significant. Moreover, with the exception of treatment NOISE50, informed followers overreacted more on average in treatments that had higher noise levels (0.35 units in treatment NOISE0, 0.598 units in treatment NOISE25, 0.232 units in treatment NOISE50, and 0.869 units in treatment NOISE75).

As for uninformed followers, it appears from Regression D1 that they did not react to past deviations of leaders from either the equilibrium or the Cournot quantities, as the estimated coefficient α_3 is very small and barely significant, while the estimated coefficient α_4 is not significant. Regressions D2-D3 show a similar picture as the coefficients $\alpha_{13} - \alpha_{19}$ are all insignificant as well.

Regressions D1-D3 also show that the behavior of informed, and especially uninformed, followers was persistent since ρ_1 and ρ_2 are both significant and positive with $\rho_2 > \rho_1$. If follower i was informed (uninformed) in round t and chose a positive Δ in round $t-1$, then, other things being equal, the same follower chose in round t a Δ that was $0.097 - 0.125$ units (0.33 units) higher than the Δ chosen in round t by follower k who chose $\Delta = 0$ in round $t-1$. As we 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.[24]

Observation 4 *The followers' tendency to over- and underreact can be summarized as follows:*

(i) Irrespective of whether followers were informed or uninformed, their modal behavior was to

[24]We also included a 2-period lagged values of $\Delta_{i,t}^{\text{info}}$ and $\Delta_{i,t}^{\text{uninfo}}$ in the regression but their coefficients were highly insignificant. In addition, 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.

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 (54% of all 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.

(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. With the exception of treatment* NOISE50, *this effect was stronger in treatments with higher levels of noise. Informed followers did not react to deviations of leaders from their equilibrium quantities.*

(iv) Uninformed followers did not react to past deviations of leaders from either their equilibrium or the Cournot quantities.

(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 particularly high if a follower was uninformed in round* t*.*

3.5 Why were leaders soft?

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

The first reason might be that leaders 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 gave up his first-mover advantage, we should expect that leaders will be more inclined to cut their quantities to promote collusion in treatments with less noise. Yet, contrary to this logic, the leaders' quantities were higher in treatments with less noise. Moreover, a large body of experiments on finitely repeated games shows that subjects tend to collude in early rounds and behave more strategically in the final rounds (e.g., Selten and Stoecker, 1986). In our experiments however, the output of leaders declined as sessions progressed rather than increased (the coefficient of T_t^3 in the leaders' regressions is negative and significant). Therefore it appears that the soft behavior of leaders was not motivated by collusion considerations.

		Mean reactions of informed and uninformed followers							
		NOISE0	NOISE25		NOISE50		NOISE75		NOISE100
q^A	$BR^B(q^A)$	$q^{B,\text{info}}$	$q^{B,\text{uninfo}}$	$q^{B,\text{info}}$	$q^{B,\text{uninfo}}$	$q^{B,\text{info}}$	$q^{B,\text{uninfo}}$	$q^{B,\text{info}}$	$q^{B,\text{uninfo}}$
13	24	24	–	–	–	–	–	–	23
14	23	19.2	–	–	–	–	–	–	20
15	23	22.33	20	21.5	19	23	19.67	–	19.67
16	22	–	–	–	21.5	–	25.6	21	21.8
17	22	–	–	22	20	–	21.67	24.67	19.78
18	21	21	–	21	–	–	23	20.5	19.77
19	21	21	21	21	22.5	22	20	20.25	19.83
20	20	20.7	18.47	20.4	19.5	20.31	22.02	23.33	20.48
Cases with $q^B < 20$		0/19	-	0/8	-	0/3	-	0/11	-

Table 8: Followers' average response to below–Cournot leader quantities

Second, it could be that leaders were reluctant to fully exploit their first-mover advantage because they do not like inequality. However, it is reasonable to expect that this concern for inequality will greatly diminish (or even disappear completely) 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 that followers will exploit them if they attempt to play collusively. Table 8 shows for each treatment how followers responded when leaders chose quantities below the Cournot quantity of 20 units. These quantities can be interpreted as attempts by leaders to induce collusive outcomes that give both players higher profits than they can get at the Cournot outcome.

As Table 8 shows, in 40 out of 41 times in which leaders chose quantities strictly below 20 units, informed followers responded with quantities that exceeded those selected by the leaders and therefore ended up getting a larger profit than the leaders. Based on this observation, it might be thought that as sessions progress, leaders will begin to exploit their first-mover advantage more often. Yet, the coefficient γ_2 in the leaders' regressions was negative and significant, implying that exactly the opposite happened: as a session progressed, leaders chose smaller and smaller quantities and exploited their first-mover advantage to a lesser degree.

This brings us to the 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 Footnotes 5 and 9 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 roughly 45% of all cases (401 out of 900 cases) and underreacted (i.e., chose $\Delta < 0$) in only 17.5% of the cases (158 out of 900 cases) and that informed followers tended to overreact more the farther away was the leader's quantity 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.

Our experiments yield several important observations. 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 susceptible to follower's deviations from best-response. 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.[25]

Second, informed followers are willing to sacrifice small amounts in order to hurt leaders who try to exploit their first-mover advantage. In particularly, the willingness of informed followers

[25] Although uninformed followers may end up underreacting to the leader's quantity (and thereby boosting the leader's payoff), risk aversion implies that leaders should be more concerned with potential overreactions by uninformed followers.

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 the leader's quantities exceed the symmetric Cournot output. Instead, they seem to simply try to play a best-response against their prediction on 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.[26] 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, Huck, and Ockenfels, 1996).

[26] 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."

References

[1] Bagwell, K. (1995), "Commitment and observability in games," *Games and Economic Behavior* 8, 271-280.

[2] Camerer, C., Knez, M. and Weber, R. A. (1996), "Timing and virtual observability in ultimatum bargaining and 'weak link' coordination games," *Social Science Working Paper No. 970*, California Institute of Technology.

[3] Chakravorti, B. and Y. Spiegel (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," *Bellcore Economics Discussion Paper No.* 103.

[4] Cooper, R., DeJong, D. V., Forsythe, R. and Ross, T. W. (1993), "Forward induction in the Battle–of–the–Sexes games," *American Economic Review*, 83, 1303-1316.

[5] Fershtman, C. and Gneezy, U. (2001), "Strategic delegation: An experiment," *Rand Journal of Economics*, Vol. 32, 352-368.

[6] Fershtman, C. and Kalai, E. (1997), "Unobserved delegation," *International Economic Review*, Vol. 38, 763-774

[7] Fershtman, C. and Judd, K. and Kalai, E. (1991), "Observable contracts: Strategic delegation and cooperation," *International Economic Review*, Vol. 32, 551-559.

[8] Fischbacher, U. (1999), "Z-Tree, Zurich Toolbox for Readymade Economic Experiments," *Working paper 21*, Institute for Empirical Research in Economics, University of Zurich.

[9] Güth, W., Huck, S., and Rapoport, A (1998), "The limitations of the positional order effect," *Journal of Economic Behavior and Organization*, 34, 313-325.

[10] Güth, W., Huck, S., and Ockenfels, P. (1996), "Two-level ultimatum bargaining with incomplete information," *Economic Journal*, 106, 593-604.

[11] Harrison G. and McCabe K. (1996), "Expectations and fairness in a fimple bargaining experiment," *International Journal of Game Theory*, 25, 303-327.

[12] Huck, S., and W. Müller, (2000), "Perfect versus imperfect observability: An experimental test of Bagwell's result," *Games and Economic Behavior* 31, 174-190.

[13] Huck, S., W. Müller, and H.-T. Normann, (2001): "Stackelberg beats Cournot - On collusion and efficiency in experimental markets," *Economic Journal* 111 (474), 749-765.

[14] Katz, M. (1991), "Game-playing agents: Unobservable contracts as precommitments," *Rand Journal of Economics*, Vol. 22, 367-388.

[15] Königstein, M. (2000), "Measuring treatment effects in experimental cross–sectional time series," In *Equity, Efficiency and Evolutionary Stability in Bargaining Games with Joint Production* (M. Königstein ed.), c. 2. Berlin/Heidelberg/New York: Springer.

[16] Kübler, D. and W. Müller, (2000), "Simultaneous and sequential price competition on heterogeneous duopoly markets: experimental evidence," Working Paper, Humboldt University.

[17] Laffont, J. J. and Tirole, J. (1993), "*A Theory of Incentives in Procurement and Regulation*," Cambridge, MA: The MIT Press.

[18] Maggi, G. (1999), "The value of commitment with imperfect observability and private information," *RAND Journal of Economics*, 30, 555-574.

[19] Müller, W. (2001), "The quality of the signal matters: A note on imperfect observability and the timing of moves," *Journal of Economic Behavior and Organization* 45(1), 99-106.

[20] Oechssler, J. and K. Schlag, (2000), "Does noise undermine the first-mover advantage? An evolutionary analysis of Bagwell's example," *International Game Theory Review* 2, 83-96.

[21] Rapoport, A. (1997), "Order of play in strategically equivalent games in extensive form," *International Journal of Game Theory* 26, 113 - 136.

[22] Rogers, W. H. (1993), "Regression standard errors in clustered samples," *Stata Technical Bulletin* 13, 19–23.

[23] Rubinstein, A. (1989), "The electronic mail game: strategic behavior under 'almost common knowledge," *American Economic Review*, 79, 385-391.

[24] Schelling, T. C. (1960), "*The Strategy of Conflict*," Cambridge, MA: Harvard University Press.

[25] Schotter, A. Zheng, W. and B. Snyder, (2000), "Bargaining through agents: An experimental study of delegation and commitment," *Games and Economic Behavior*, Vol. 30, 248-292.

[26] Selten, R., and R. Stoecker, (1986), "End behavior in sequences of finite prisoner's dilemma supergames: A learning theory approach," *Journal of Economic Behavior and Organization* 7, 47-70.

[27] Smith, V. L. (1989), "Theory, experiment and economics," *Journal of Economic Perspectives* 3, No. 1, 151-169.

[28] Stackelberg, H. von (1934): *Marktform und Gleichgewicht*, Vienna and Berlin: Springer.

[29] STATA Corp. (1999): *Reference manual*, release 6, STATA Press Texas.

[30] Suits, D. (1984), "Dummy variables: Mechanics v. interpretation," *Review of Economics and Statistics*, vol. 66, 177-180.

[31] van Damme, E., and S. Hurkens, (1997), "Games with imperfectly observable commitment," *Games and Economic Behavior* 21, 282-308.

[32] Weimann, J., C-L Yang, and C. Vogt, (2000), "An experiment on sequential rent-seeking," *Journal of Economic Behavior and Organization* 41, 405-426.

[33] White, H. (1980), "A heteroskedasticity-consistent covariance matrix estimator and a direst test for heteroskedasticity," *Econometrica* 48, 817–30.

5 Appendix

5.1 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.

[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.

5.2 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 players.

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	210	196
15	480	465	450	435	420	405	390	375	360	345	330	315	300	285	270	255	240	225	210	195
16	496	480	464	448	432	416	400	384	368	352	336	320	304	288	272	256	240	224	208	192
17	510	493	476	459	442	425	408	391	374	357	340	323	306	289	372	255	238	221	204	187
18	522	504	486	468	450	432	414	396	378	360	342	324	306	288	270	252	234	216	198	180
19	532	513	494	475	456	437	418	399	380	361	342	323	304	285	266	247	228	209	190	171
20	540	520	500	480	460	440	420	400	380	360	340	320	300	280	260	240	220	200	180	160
21	546	525	504	483	462	441	420	399	378	357	336	315	294	273	252	231	210	189	168	147
22	550	528	506	484	462	440	418	396	374	352	330	308	286	264	242	220	198	176	154	132
23	552	529	506	483	460	437	414	391	368	345	322	299	276	253	260	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