

An Experimental Study of Commitment and Observability in Stackelberg Games

John Morgan* Felix J. J. Várdy†

Please send all correspondence to:

CERAS, Ecole Nationale des Ponts et Chaussées

28, rue des Saints-Pères, 75343 Paris, cedex 07

France

September 7, 2001

Abstract

We report on experiments examining the value of commitment in Stackelberg games where the follower chooses whether to pay some cost ε to perfectly observe the leader's action. Várdy (2001) shows that in the unique

*Department of Economics and Woodrow Wilson School, Princeton University, Princeton, NJ 08544.

†CERAS, Ecole Nationale des Ponts et Chaussées, 75343 Paris cedex 07. Corresponding author.

Email: Felix.Vardy@enpc.fr.

pure strategy subgame perfect equilibrium of this game, the value of commitment is lost completely; however, there exists a mixed strategy subgame perfect equilibrium where the value of commitment is fully preserved.

In the data, the value of commitment is almost completely preserved when the cost of looking is small, while it is lost when the cost is large. Nevertheless, for small ε , equilibrium behavior is clearly rejected. Instead, subjects persistently play intuitively appealing non-equilibrium strategies in which the probability of the follower choosing to observe the leader's action is a decreasing function of the observation cost. We find little evidence of convergence to equilibrium play.

JEL Numbers: L13, C91, C72

Keywords: Stackelberg Duopoly, Experiments, Commitment

1 Introduction

It is well-known that a firm can sometimes gain an advantage by committing to an action ahead of its rivals. Indeed, the familiar Stackelberg duopoly model is a prominent example. Here, a leader firm obtains an advantage by committing to produce a large quantity of some homogeneous good. The follower, upon observing the leader's choice, then optimally decides to produce less of the good and the leader thereby gains market share and profit at the expense of its rival. Suppose, however, that to observe the leader's choice, the rival must undertake some investigative activity – perhaps at very low cost. Absent undertaking this activity, the rival simply remains in the dark about the leader's action. How does this option affect

strategic choices in the Stackelberg game? How does it affect the leader's value of commitment?¹

To fix ideas, consider the normal form game, g :

$L \setminus F$	s	c
S	(500, 200)	(300, 100)
C	(600, 300)	(400, 400)

In this game, two firms, designated L(eader) and F(ollower), are competing with one another. Here, the choices of the leader, S and C , correspond to the S(tackelberg) and C(ournot) outputs, respectively. Likewise for the choices of the follower. If firms chose quantities simultaneously, the game is dominance solvable and yields the unique rationalizable outcome (C, c) . In contrast, if L moves first followed by F , and if L 's choice is fully observable, then the unique subgame perfect equilibrium of the game is (S, s) . Thus, the power of commitment yields L an additional 100 points at the expense of 200 points lost on the part of F .

Next, consider a variation of this game that Várdy (2001) refers to as the 'costly leader game.' In this game, L chooses first. Next, F decides whether or not to spend an amount $\varepsilon > 0$ to obtain information about L 's choice. In particular, spending ε lets F learn what L chose precisely; whereas F obtains no information about L 's choice if she does not spend ε . Following this, F then makes a quantity choice as well and payoffs are realized.

¹In this paper, we use the terms 'value of commitment' and 'first-mover advantage' interchangeably. Both terms refer to the extra (equilibrium) payoff the leader gets from moving first, as compared to his (equilibrium) payoff when the players move simultaneously.

If one restricts attention to pure-strategy subgame perfect equilibria, F never chooses to learn about L 's choice, and the outcome of the game is (C, c) . In other words, the value of being a first-mover is completely undermined even for arbitrarily small costs of learning (ε). The intuition is that, since F fully anticipates L 's choice, there is no point in spending anything merely to confirm these expectations. Of course, L anticipates this behavior on F 's part and thus cannot hope to influence F 's choice through making the first move. The game essentially collapses to the simultaneous game and the Cournot outcome results.

In addition to this pure-strategy equilibrium, when costs (ε) are sufficiently low, there is also a subgame perfect equilibrium in mixed strategies where the value of the first-mover advantage to the leader is perfectly preserved. Specifically, for any $\varepsilon < 50$, a subgame perfect equilibrium to this game is where F obtains information about L 's choice exactly half of the time. If F chooses not to learn about F 's choice, she chooses between c and s with equal probability. The leader (L) chooses S with probability $1 - \frac{\varepsilon}{100}$. In this equilibrium, which Várdy (2001) refers to as the 'noisy Stackelberg equilibrium,' L 's payoff is 500 – exactly the same as in the usual Stackelberg game.² This equilibrium suggests that the presence of an option not to observe the leader's choice may not impair the first-mover advantage of the leader at all.

These properties are not unique to this example. Várdy (2001) analyzes a generic class of duopoly Stackelberg games and shows that the properties of pure-

²It is interesting to note that F 's expected payoff is $200 + \varepsilon$, which is higher than that under the usual Stackelberg game.

strategy and noisy Stackelberg equilibria hold quite generally. Thus, one is faced with equilibria arising in the costly leader game that differ dramatically in their implications for the value of being the first mover in this class of games. One approach to resolving these differing predictions is to apply equilibrium refinements to rule out certain of the equilibria in this game. Várdy offers some analysis along these lines, suggesting that it is the pure-strategy equilibrium that survives. By contrast, the approach of this paper is empirical. We conduct controlled laboratory experiments where subjects played the game given above. Our main treatment was to vary the follower's cost of learning the leader's action at five different levels: $\varepsilon = \{1, 15, 30, 45, 60\}$ and observe how this affected outcomes and the empirical frequency of the follower's choosing to observe the leader's choice.

When the cost of observing is low ($\varepsilon = 1, 15$), the Stackelberg outcome occurs more than 87% of the time and the leader retains more than 90% of the value of being a first-mover. In contrast, when costs are high ($\varepsilon = 45, 60$), it is the Cournot outcome that occurs more than 82% of the time, and the leader retains only 15% of the value of being a first-mover. When costs are intermediate ($\varepsilon = 30$), neither Cournot nor Stackelberg constitute more than 50% of the outcomes. Thus, it appears that the cost of observing the leader's choice is an important determinant in equilibrium selection.³

However, especially for small cost, the observed frequency with which the follower observed the leader's choice is inconsistent with equilibrium play. Under

³More precisely, it is an important determinant for all treatments except $\varepsilon = 60$. In this treatment, the Cournot outcome is the unique subgame perfect equilibrium outcome.

the noisy Stackelberg equilibrium, the follower is predicted to observe the leader's choice 50% of the time independent of the cost of observation. When costs are low, the follower chose to 'look' at least 70% of the time; moreover, this frequency increased further, up to 83%, as the cost decreased. Thus, this aspect of the data is inconsistent with the predictions of the noisy Stackelberg equilibrium. Interestingly, the medium cost treatment came closest to the theoretical looking frequency of the noisy Stackelberg equilibrium. Here, subjects chose to look 46% of the time; however, the outcomes and payoffs in this treatment were inconsistent with theoretical predictions. In short, despite finding that the value of commitment is largely preserved when the cost of looking is relatively small, it does not appear that subjects are playing something approximating the noisy Stackelberg equilibrium to support this outcome.

The remainder of the paper proceeds as follows: In section 2, we review the relevant theoretical and empirical literature related to the durability of commitment in duopoly. Section 3 outlines the procedures used in the experiments. Section 4 reports the results of the experiments and compares these to theoretical predictions. In section 5, we examine some alternatives that might help to explain discrepancies between the theory and our results. Finally, section 6 concludes. The instructions used in the experiment are contained in the appendix.

2 Literature Review

In an important paper, Bagwell (1995) points out the fragility of the first-mover advantage in Stackelberg duopoly when the follower receives a noisy signal about the action taken by the leader. That is, in a Stackelberg duopoly model where there is a small chance that the follower receives an inaccurate signal about the leader's choice, Bagwell shows that the set of pure-strategy Nash equilibria in the 'noisy leader game' coincide with the set of pure-strategy equilibria in the simultaneous move Cournot version of the game. Bagwell's point is that the value of commitment is lost entirely when the noise in the signal received by the follower becomes arbitrarily small.

This observation led to a lively debate about the robustness of the first-mover advantage in Stackelberg duopoly. See, for instance, van Damme and Hurkens (vdH, 1997), Oechssler and Schlag (1997), Güth et al. (1998), and Huck and Müller (2000). The first and last of these papers merit particular attention. The paper by van Damme and Hurkens shows that in addition to the pure-strategy equilibria identified by Bagwell, there also exist mixed-strategy equilibria that converge to the Stackelberg outcome. The paper then offers an equilibrium selection procedure that selects these mixed-strategy equilibria over those identified by Bagwell.

Huck and Müller (2000) use laboratory experiments to investigate outcomes in noisy leader games. They find that when the probability of the follower's receiving an incorrect signal is small, outcomes are close to Stackelberg. They argue that the pure strategies identified by Bagwell are not behaviorally relevant. On the

relevance of the noisy Stackelberg equilibria, however, they were not able to draw any firm conclusions. In part, this had to do with the relatively small size of their data set. More importantly, however, is an intrinsic difficulty in differentiating mixed-strategy play in the noisy leader game. This follows from the fact that the only way to identify strategically sophisticated mixed-strategy play from naive play where subject mostly choose Stackelberg quantities is in the choices of subjects when they receive the unexpected (Cournot) signal. When the probability of getting this signal is very low, in the noisy leader, game noisy Stackelberg equilibria become observationally virtually equivalent to ‘noised up’, naive Stackelberg play. In other words, while Huck and Müller observe that the leader’s first-mover advantage is largely preserved when the noise is small, they cannot easily tell whether this is for all the ‘right’ (equilibrium) reasons, or for all the ‘wrong’ (non-equilibrium) reasons. This makes it difficult for Huck and Müller to establish conclusively that subjects are in fact playing the mixed-strategy equilibrium identified by van Damme and Hurkens in the laboratory. They write: “However, this observation [i.e., acceptance of the null hypothesis of mixed-strategy equilibrium play is not really valid, as there are still too few observations. Hence, it is too early to draw a final conclusion regarding the claim of van Damme and Hurkens, and more testing has to be done.”

Huck and Müller (2000) represent the nearest antecedent to the present paper; thus it is useful to distinguish contributions of our study as compared to theirs. First, the theoretical model underlying the experiments differs in the underlying economic reasons for the potential fragility of commitment. Our paper focuses on

how the presence of an option not to learn about the leader's choice can undermine commitment; whereas the existing literature is concerned with how noise in the communication technology can undermine commitment.⁴ As we pointed out in the introduction, there are a variety of circumstances where firms have the option of undertaking costly investigative activities to learn more about the choices of their rivals. Thus, it seems economically relevant to understand this effect empirically. Second, the costly leader theory readily offers a way to distinguish equilibrium play from non-equilibrium play. In particular, the frequency with which followers are predicted to observe the leader's choice in a noisy Stackelberg equilibrium is constant at 50% for all $\varepsilon < 50$. This has directly testable implications for equilibrium play.

3 Experiment

In this section, we describe the design of the experiment and offer some justification for key design choices. We begin with a formal analysis of equilibria arising in the game we studied.

⁴There is a technical difference between the solution concepts underlying the analysis of costly leader games versus noisy leader games. In costly leader games, subgame perfect equilibrium remains the appropriate solution concept – as in the original Stackelberg duopoly. In noisy leader games, Nash equilibrium is the appropriate solution concept.

3.1 The Game

Recall the normal form game g :

$L \setminus F$	s	c
S	(500, 200)	(300, 100)
C	(600, 300)	(400, 400)

We focus attention on a modified version of the standard Stackelberg leader game Γ , in which player F gets to observe L 's action before choosing s or c , if and only if he spends an amount ε on information gathering. In this costly leader game, Γ^ε , the decision whether to expend ε and observe L 's action is denoted by $o = y, n$. Here, we write y if F decides to observe L 's action, and n if he decides not to observe L 's action. We denote F 's pure-strategy of not looking and playing s by (n, s) . The strategy of not looking and playing c is denoted by (n, c) , while the strategy of looking and best responding to his observation is denoted by (y, b) . Finally, F 's strategies 'look and always play s ', 'look and always play c ', and 'look and play c upon observing S and s upon observing C ' are denoted by (y, s) , (y, c) and (y, cs) , respectively. Note that none of these last three strategies may be part of any subgame perfect equilibrium.

Player L 's pure strategies in Γ^ε are the same as in g and Γ , i.e., S and C .

The probability with which F looks at L 's action is denoted by p^ε .⁵ The value of p^ε is therefore equal to the sum of the probabilities that F plays (y, b) , (y, s) , (y, c) and (y, cs) , where only the first strategy is subgame perfect.

⁵As, in equilibrium, p^ε turns out not to depend on ε , we sometimes suppress the superscript ε .

In Γ^ε , we define the outcome of a strategy profile to be the probability distribution that this profile induces on $\{S, C\} \times \{s, c\}$. Hence, an outcome is not concerned with whether F looks at L 's action or not. We use this more restricted definition to preserve comparability of outcomes of Γ^ε with outcomes of g , Γ , and the noisy leader game.

Finally, the payoffs in the costly leader game Γ^ε are as follows. For each pure outcome Ss, Cc, Sc and Cs , the players' payoffs in Γ^ε are in principle the same as in g . For player F , however, the payoff in Γ^ε also depends on whether he looks at L 's action or not. If F looks, the cost ε of looking is subtracted from his payoff.

With these payoffs, we may use the results obtained by Várdy to characterize the set of subgame perfect equilibria (SPE) of Γ^ε :

- Noisy Stackelberg equilibrium: For all $\varepsilon \in [0, 50]$, there exists a SPE characterized by $\Pr_L(S) = 1 - \frac{\varepsilon}{100}$, $p = \frac{1}{2}$, $\Pr_F(y, b) = \frac{1}{2}$ and $\Pr_F(n, s) = \frac{1}{2}$.
- Noisy Cournot equilibrium: For all $\varepsilon \in [0, 50]$, there exists a SPE characterized by $\Pr_L(C) = 1 - \frac{\varepsilon}{100}$, $p = \frac{1}{2}$, $\Pr_F(y, b) = \frac{1}{2}$ and $\Pr_F(n, c) = \frac{1}{2}$.
- Pure-strategy Cournot equilibrium: In the unique pure-strategy SPE of this game, L always plays C and F always plays (n, c) , for all $\varepsilon > 0$.
- Continuum of equilibria: In the special case where $\varepsilon = 50$, there exists a continuum of SPE characterized by $\Pr_L(S) = \frac{1}{2}$, $p = \frac{1}{2}$, $\Pr_F(y, b) = \frac{1}{2}$, and $\Pr_F(n, s) \in \left[0, \frac{1}{2}\right]$.

In our laboratory implementation of this game, we let ε take on five different

values: 1, 15, 30, 45, and 60. Hence, the continuum of equilibria does not play a role in the experiment, as ε is chosen unequal to 50 at all times. To get an intuition for the magnitude of these observation costs, compare them with the 100 point loss the follower risks by not looking. Notice that when $\varepsilon = 60$, only the pure-strategy Cournot equilibrium arises; otherwise, there are three equilibria – the noisy Stackelberg, the noisy Cournot, and the pure-strategy Cournot equilibrium.

The main reason for choosing this set of parameters is that it allows us to explore the durability of commitment near both ends of the parameter space where both the noisy Stackelberg and Cournot outcomes coexist. The treatment where $\varepsilon = 60$ offers a useful benchmark in that the game has a unique subgame perfect equilibrium consisting of the pure-strategy Cournot outcome.

3.2 Experimental Procedures

The experiment consisted of three sessions conducted at Princeton University in January and February of 2001. Subjects participating in the experiment were recruited from the undergraduate population at Princeton University, through a number of E-mail lists. Ten subjects participated in each session and no subject appeared in more than one session. At the beginning of each session, subjects were seated in the same room, at separate computer terminals and given a set of instructions, which were also read aloud by the experimenter. Communication between the subjects was strictly forbidden and, to the best of our knowledge, did not occur.

In each session, the costly leader game was played 100 times in succession. At

the beginning of each period, subjects were randomly paired with one another, randomly assigned the roles of leader and follower, and the cost of observing the leader's choice was randomly generated. This cost (ε) was equally likely to be any of the five values: 1, 15, 30, 45, or 60. The value of ε was displayed to all players throughout the period.

Player A (the leader) then chose between actions labeled 'U'(p) and 'D'(own) corresponding to 'S' and 'C', respectively, in game g . Following this, player B indicated whether or not she wished to pay ε to learn player A's choice. B indicated this by choosing between buttons labeled 'Yes' and 'No.' If B chose 'Yes', then A's choice was displayed on her screen. If B chose 'No,' she received no additional information. Following this, player B chose between actions labeled 'L'(eft) and 'R'(ight), corresponding to 's' and 'c' in game g . At the end of the round, a results screen displayed all of the choices made by players A and B as well as each player's payoffs in that period. The players did not learn the results of matchings other than their own. Figure 1 presents screenshots of the input screens that subjects observed when making their choices. The instructions used in the experiment are contained in the appendix.



Figure 1: Screen shots

Each session lasted about an hour. The subjects' earnings were calculated from the total points they earned during the experiment. For every 1000 points they earned 50 cents, with subject earnings averaging \$18.30 in the experiment. All subjects were paid in cash and in private.

In arriving at these procedures, we considered issues regarding learning, convergence, fairness, feedback, repeated interaction and collusion.

The large number of periods (100) was chosen to give the subjects ample opportunity to experiment and learn. In this way, we wanted to ensure that convergence

to equilibrium play had a fair chance of occurring. The random assignment of the leader and follower roles was chosen to force the subjects to play the game and think it through both from the leader and the follower perspective. This was done to 1) speed up the learning process, 2) make it less likely that fairness issues would play a role as all subjects could expect to be leader and follower about the same number of times, 3) allow us to have every subject play against all other subjects in the group.

In combination with the random and anonymous matching, the fact that all subjects played against all other subjects reduced the probability of playing against a specific co-player to 11%. Of course, more subjects per session would have further reduced the chance that repeated interaction could play a major role. The physical capacity of our laboratory, however, did not allow for more than 10 people. Nevertheless, even with our relatively small number of subjects per session, it is hard to see how one could effectively punish or reward co-players conditional on their previous behavior. In the end, you could be playing anybody out of nine people in any given round. Moreover, in each round, a player only observes the outcome of his own pairing. This makes implicit collusion through punishing or rewarding co-players very difficult. Finally, in none of the post-experiment questionnaires filled out by the subjects there was any reference to specific players, repeated interaction or contingent strategies. Instead, words like ‘the others’ and ‘the group’ suggest that the subjects viewed the interaction as being with an amorphous population in which they could not distinguish between individuals.

4 Results

4.1 Overview

Though the game is relatively simple, we anticipated that, initially, subjects would experiment in their approach to the costly leader game before settling on their strategies for each treatment. As such, the equilibrium predictions of the costly leader game are more likely to be valid for later periods rather than for the initial ‘learning’ phase of the game. Our impression was that by round 40 subject play became fairly stationary. To verify this, we regressed the empirical frequency of followers choosing to observe the leader’s action on the number of the round. For all treatments, we fail to reject the null hypothesis that the round coefficient for rounds 41-100 is zero. (5% confidence level.) Similarly, regressing the empirical frequency with which the leaders choose S on the number of the round yields an identical conclusion.

We begin by presenting summary statistics of play in each of the three sessions. First we make some initial comments, later we offer formal tests. Table 1 shows the percentage of Stackelberg actions taken by leaders for each treatment occurring in

a given session. The tilde symbol ($\tilde{\cdot}$) denotes empirical frequencies.

$(\tilde{S}, \text{ in } \%)$	Session I	Session II	Session III	Overall
$\varepsilon = 1$	89	95	95	93
$\varepsilon = 15$	96	93	80	90
$\varepsilon = 30$	43	43	37	41
$\varepsilon = 45$	15	10	0	8
$\varepsilon = 60$	13	3	3	6

Table 1: Percentage of Stackelberg play by leaders

Notice that the proportion of Stackelberg play decreases as the follower's cost of becoming informed rises. The percentage of Stackelberg play can be roughly divided into three regimes: 1) Low cost ($\varepsilon = 1, 15$) where Stackelberg is the predominant play, 2) High cost ($\varepsilon = 45, 60$) where Cournot is the predominant play, and 3) Medium cost ($\varepsilon = 30$) where neither Stackelberg nor Cournot is played overwhelmingly often.

Table 2 shows the percentage of the time that followers chose observe the leader's action at a cost of ε . It is denoted by \tilde{p} .

$(\tilde{p}, \text{ in } \%)$	Session I	Session II	Session III	Overall
$\varepsilon = 1$	80	85	84	83
$\varepsilon = 15$	75	60	75	70
$\varepsilon = 30$	48	49	42	46
$\varepsilon = 45$	17	12	10	13
$\varepsilon = 60$	8	0	10	6

Table 2: Percentage of looking by followers

Much like Table 1, Table 2 shows that the percentage of games that followers choose to observe the leader's action decreases when the cost of observation increases. According to equilibrium, however, the looking frequency is independent of the observation cost, and either equal to 50% or to 0%. Instead, it decreases rather smoothly, all the way from 83% for $\varepsilon = 1$, to 13% for $\varepsilon = 45$ (and 6% for $\varepsilon = 60$). Hence, for the low cost treatments, the empirical frequency with which followers choose to observe is far in excess of the theoretical predictions (50%) of the noisy Stackelberg equilibrium (or indeed any equilibrium of the costly leader game). For the high cost treatments, the looking frequency is a bit too high for the pure-strategy Cournot equilibrium (0%) and much too low for the noisy Cournot or Stackelberg equilibria (50%).

What happens in the event that a follower chooses to observe the leader's action? In all but four out of 434 cases, the follower best responded to his observation. Thus the data point overwhelmingly in favor of followers best responding following the decision to observe the leader's action.

Table 3 shows the percentage of Stackelberg choices made by followers in in-

stances where they did not observe the leader's choice.⁶

$(\tilde{s} \mid n, \text{ in } \%)$	Session I	Session II	Session III	Overall
$\varepsilon = 1$	88	100	100	95
$\varepsilon = 15$	86	95	86	90
$\varepsilon = 30$	47	67	45	52
$\varepsilon = 45$	20	12	6	13
$\varepsilon = 60$	11	8	6	8

Table 3: Percentage of Stackelberg play by followers after not looking

For treatments with high costs of observation, we observe s being played much less frequently than for treatments with low costs. Under noisy Stackelberg equilibrium followers are predicted to always choose s when not looking, while under (noisy) Cournot equilibrium followers are predicted to always choose c .

4.2 High Cost Treatments

As the preceding tables showed, behavior in high cost treatments ($\varepsilon = 45, 60$) was markedly different from that in low cost treatments ($\varepsilon = 1, 15$). In the low cost treatments, the Stackelberg outcome Ss occurred 92% and 87% of the time, compared to 82% and 87% of Cournot outcomes under high costs. We now take a more detailed look at behavior relative to equilibrium predictions for these treatments.

⁶The weights of the different sessions in determining the overall percentage are determined endogenously by the number of times followers decided to look in a given session. Therefore, the overall percentage of $\tilde{s} \mid n$ is not always equal to the average of the averages of the three sessions.

Unless stated otherwise, the reported p-values are based on χ^2 tests. (See, e.g., Hogg and Craig, 1995, pp.293-301.) We use a significance level of 5%.

$\varepsilon = 60$. This is the only treatment in which there is a unique subgame perfect equilibrium; namely, the Cournot equilibrium. Indeed, there are no beliefs about the strategy of the leader for which it is a best response for the follower to choose to observe the leader's action. That is, observing is a dominated strategy. Iterative deletion of dominated strategies leads to the conclusion that $(C, (n, c))$ is the unique rationalizable strategy profile. The players' equilibrium payoff is 400.

Compared to this benchmark, followers chose to observe the leader's action too often. Namely, 6% of the time, instead of 0%. And they played s too often conditional on not looking (8% of the time, instead of 0%). Leaders chose the Stackelberg action S about 6% of the time under this treatment (instead of 0%), and earned an average payoff of 411 (instead of 400). At all confidence levels, all of these statistics are trivially rejected as having been generated by the pure strategy Cournot equilibrium. The reason is that the null-hypothesis of Cournot equilibrium excludes all randomness.

When we look at the last four rounds of the experiment in which $\varepsilon = 60$ (this corresponds to a total of 60 observations), we see that behavior moves somewhat closer to equilibrium towards the end of the experiment. The values \tilde{p} , $\tilde{s} | n$, and \tilde{S} all go down to 5%, while the leader's average payoff becomes 405.

$\varepsilon = 45$. In this treatment, as in the lower cost treatments, there is a subgame perfect equilibrium where the first-mover advantage is fully retained, as well as equilibria in which it is lost entirely. Subject behavior in this treatment is more consistent with Cournot than with the noisy Stackelberg equilibrium. In both the pure and noisy Cournot equilibria, the followers are predicted to never play s , conditional on not observing L 's choice. The empirical frequency $\tilde{s} | n$ is 13%. In the pure-strategy Cournot equilibrium, leaders are also predicted to never play S . The empirical frequency is 8%. Rejection of the equilibrium hypotheses is again trivial. Finally, in the noisy Cournot equilibrium, L is predicted to play S 45% of the time. This is significantly higher than the observed 8%, with a p-value of .000.

The two Cournot equilibria differ dramatically as to the frequency with which followers choose to observe the leader's action. In the noisy Cournot equilibrium, it is predicted to occur 50% of the time. (Note: the same is true for the noisy Stackelberg equilibrium.) Under the pure Cournot equilibrium the predicted frequency is 0%. The observed 13% is inconsistent with both. Rejection of pure Cournot is trivial, while the p-value of the hypothesis that the observed 13% has been generated by a 50-50 process is .000. While it is possible that the observed frequency is consistent with some fraction playing noisy Cournot and the rest playing pure Cournot, as we pointed out above, the other choices of the followers are not consistent with this hypothesis.

For $\varepsilon = 45$, average play in the last four rounds is very similar to that in the full data set from round 41 onwards. Despite the much smaller number of observations,

all equilibrium tests at the 5% significance level lead to the same conclusions as before. Naturally, the p-values are larger than before.

To sum up, the observed data in the high cost treatment are more consistent with Cournot than with the noisy Stackelberg behavior, and the leader's value of commitment is largely lost. Nevertheless, the data can be rejected as coming from either of the Cournot equilibria. In the $\varepsilon = 60$ treatment, only Cournot outcomes are consistent with equilibrium, and subjects coordinated on Cc outcomes 87% of the time. In the $\varepsilon = 45$ treatment, the value of commitment need not be lost in theory. In practice, however, outcomes are again mostly Cournot (82%).

4.3 Low Cost Treatments

In these treatments ($\varepsilon = 1, 15$), Stackelberg behavior predominates. In this subsection we compare the data to the predictions of the noisy Stackelberg equilibrium.

$\varepsilon = 1$. In this treatment, the leader's value of commitment is largely retained. The average payoff of a leader is 496, compared to 500 in the standard Stackelberg outcome. The noisy Stackelberg equilibrium calls for the leader to choose S 99% of the time. In fact, S was chosen 93% of the time. This difference is statistically significant with a p-value of .000. Conditional on not observing the leader's action, the noisy Stackelberg equilibrium calls for the followers to always play s . In fact, they chose this action 95% of the time.

What is most strikingly inconsistent between the theory and the data is the frequency with which followers chose to observe the leader's action. In any equi-

librium, this frequency is predicted to be no more than 50%. Yet, followers chose to observe 83% of the time in the data. This difference is statistically significant, with a p-value of .000. We find no evidence for convergence to equilibrium play.

$\varepsilon = 15$. The leader's value of commitment is only slightly eroded in this treatment (average payoff of 491 points) and still much closer to the prediction of the noisy Stackelberg equilibrium (500 points) than to any other equilibrium. In this treatment, leaders chose the S strategy 90% of the time versus a theoretical prediction of 85%. This difference is not statistically significant with a p-value of .091. Followers continue to observe the leader's action more frequently (70%) than is called for in equilibrium (50%) and the difference is statistically significant with a p-value of .000. Conditional on being uninformed, they choose s less frequently (90%) than the theory predicts (100%), which leads to a trivial rejection of the noisy Stackelberg hypothesis.

Play does not converge to equilibrium. When we restrict attention to the last four rounds with $\varepsilon = 15$, all the equilibrium tests lead to the same acceptance/rejection results as mentioned above.

To sum up, in the low cost treatments, the value of commitment is largely retained, yet the data do not support the view that the noisy Stackelberg equilibrium is generating this behavior. In particular, the followers observe the leader's action far too frequently. Moreover, this frequency clearly depends on the cost of looking while the noisy Stackelberg equilibrium demands that it be independent of it and

equal to 50%.⁷

4.4 Medium Cost Treatment

In this treatment ($\varepsilon = 30$), the followers' empirical looking frequency (46%) is close to the 50% predictions of the noisy Cournot and noisy Stackelberg equilibria. The difference is not statistically significant with a p-value of .284. In other dimensions, however, the data are inconsistent with the theory. The empirical frequency \tilde{S} , which is equal to 41%, is rejected as coming from any equilibrium, with a p-value of .000. Moreover, in the noisy Stackelberg equilibrium, $\tilde{s} \mid n$ would have been 100%, and in the Cournot and noisy Cournot equilibria 0%. In the data, however, $\tilde{s} \mid n$ is 52%. That is, conditional on not looking, half of the follower population tries to coordinate on the Stackelberg outcome, while the other half coordinates on the Cournot outcome. With 455 points, also the leader's payoff is right in between the two.

5 Alternative Explanations

For low observation costs ($\varepsilon = 1, 15$), we have seen that Stackelberg outcomes predominate and that the leader's value of commitment is largely preserved. Despite this fact, it does not appear that subjects are playing something approximating the noisy Stackelberg equilibrium to support this outcome. What is most strikingly

⁷The null hypothesis that the looking frequency for $\varepsilon = 1$ is equal to the looking frequency for $\varepsilon = 15$ is rejected with a p-value of .00.

inconsistent between the theory and the data is the frequency with which followers chose to observe the leader's action. For low costs, this frequency is significantly larger than predicted and depends strongly on the particular value of the cost. In equilibrium, it would have been independent of the cost.

For high cost treatments ($\varepsilon = 45, 60$), Cournot outcomes predominate and the leader's value of commitment is largely lost. Among the equilibria, pure-strategy Cournot best approximates the data. The (statistically significant) inconsistencies lie in the fact that quantities which should be zero are quite small but not equal to zero.

The data in the medium cost treatment ($\varepsilon = 30$) seem quite far removed from any equilibrium. Neither Stackelberg nor Cournot is played overwhelmingly often, leaders' payoffs are right in between noisy Stackelberg and (noisy) Cournot and, conditional on being uninformed, half of the followers try to coordinate on the Stackelberg and the other half on the Cournot outcome.

In this section, we study some alternative theories and examine whether they can help to better explain the data.

5.1 Quantal Response Equilibrium

We begin by studying the (Logit) quantal response equilibrium (QRE) of our game and look whether this fixed point in 'trembling' strategies can explain the data. The QRE theory was developed by McKelvey and Palfrey (1995). The interested reader may want to consult McKelvey and Palfrey (1995) for details.

As we shall see, the QRE can explain many aspects of the data, but not all. In fact, the QRE does well for large and medium ε . And for the right choice of the error parameter λ , which is an exogenous element of the QRE theory, the predicted probability of looking, p , does become a decreasing function of ε . Also, the observed switch from the prevalence of Stackelberg outcomes for small ε , to the prevalence of Cournot outcomes for large ε is well described by the QRE. Finally, the QRE avoids the subgame perfect equilibria's zero-probability predictions of, e.g., $\tilde{s} \mid n$, thus enabling meaningful statistical analysis. However, for small ε , the QRE fails to predict the looking frequencies correctly. Specifically, the QRE looking probabilities remain too small, no matter how we choose the error parameter λ .

In the experiment, the followers played best responses conditional on observing the leader's action more than 99% of the time. This means that, with respect to true equilibrium play, the deviation in that dimension was at least an order of magnitude smaller than the deviations in the looking behavior and the other actions of L and F . Therefore, we apply the QRE theory to the reduced normal form of Γ^ε , denoted by h' , in which the follower always best responds conditional on observing the leader's action. (Applying the theory to original normal form (h^ε) does not change the results in any essential way, but does complicate the exposition and the calculations.)

The game h' looks as follows:

$L \setminus F$	(n, s)	(n, c)	(y, b)
S	$(5, 2)$	$(3, 1)$	$(5, 2 - \varepsilon)$
C	$(6, 3)$	$(4, 4)$	$(4, 4 - \varepsilon)$

Denoting $\Pr_L(S)$ by α , $\Pr_F(y, b)$ by q_1 and $\Pr_F(n, c)$ by q_2 , a QRE of h' is then characterized by a solution $(\alpha^\varepsilon(\lambda), q_1^\varepsilon(\lambda), q_2^\varepsilon(\lambda))$ to the system:

$$\begin{aligned}\alpha &= \frac{e^{\lambda(5-2q_2)}}{e^{\lambda(5-2q_2)} + e^{\lambda(6-2q_1-2q_2)}} \\ q_1 &= \frac{e^{\lambda(4-\varepsilon-2\alpha)}}{e^{\lambda(4-\varepsilon-2\alpha)} + e^{\lambda(4-3\alpha)} + e^{\lambda(3-\alpha)}} \\ q_2 &= \frac{e^{\lambda(4-3\alpha)}}{e^{\lambda(4-\varepsilon-2\alpha)} + e^{\lambda(4-3\alpha)} + e^{\lambda(3-\alpha)}}\end{aligned}$$

This system does not have explicit solutions in terms of $(\alpha^\varepsilon(\lambda), q_1^\varepsilon(\lambda), q_2^\varepsilon(\lambda))$. However, calculating the solutions for each ε numerically, and plotting, in Figure 2, the unique distinguished curves σ_{QRE}^ε as functions of λ , we get:

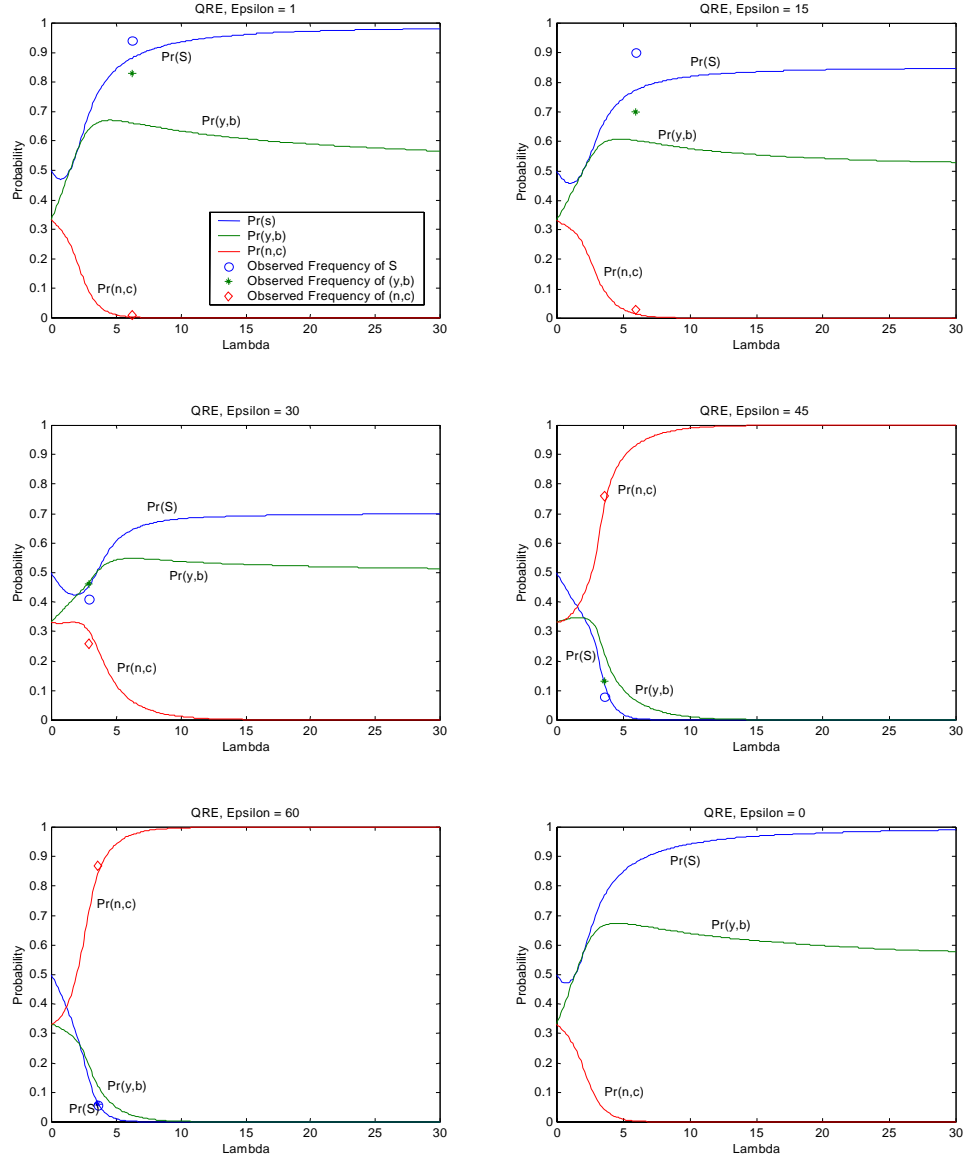


Figure 2: QRE plots

Table 4 reports the maximum likelihood estimates $\hat{\lambda}$ of λ , for the different values of ϵ separately. Also, it tabulates the implied probabilities for the different actions of L and F under the QRE with error parameter $\hat{\lambda}$.⁸ The observed frequencies have been added in the row below the QRE predictions. (They have also been

⁸Because we have applied the QRE to h' instead of h^ϵ , these estimates are based on the reduced data set in which 4 out of 900 observations have been eliminated. These four observations

superimposed on the first five QRE graphs in Figure 2.) All frequencies are denoted in percentages, and have been rounded off to the nearest integer. The maximum likelihood estimates $\hat{\lambda}$ are given in one decimal.

(ML Estimates in %)	$\hat{\lambda}$	S	C	(y, b)	(n, c)	(n, s)
$\varepsilon = 1$	6.2	88	12	66	0	34
		94	6	83	1	16
$\varepsilon = 15$	5.9	77	23	60	2	38
		90	10	70	3	27
$\varepsilon = 30$	2.9	46	54	47	30	23
		41	59	46	26	28
$\varepsilon = 45$	3.6	11	89	22	74	5
		8	92	13	76	11
$\varepsilon = 60$	3.6	6	94	12	84	4
		6	94	6	87	8

Table 4: Maximum likelihood QRE predictions and observed frequencies

A very rough, heuristic measure of how well the maximum likelihood QRE model fits the data is given by comparing the sum of the absolute differences between the predicted probabilities and the observed frequencies. (Note: comparing the Log-Likelihoods over different ε 's does not make sense.) According to this criterion, the correspond to instances is which the follower did not best respond to his observation. In terms of observed frequencies, the only number that changes is the frequency of S when $\varepsilon = 1$. It goes from 93% to 94%. The other entries are the same in both data sets due to rounding.

maximum likelihood QRE predictions are worse for small ε than for large ε . For $\varepsilon = 1, 15, 30, 45, 60$, the sum of the absolute differences between the predicted and the observed frequencies is 38, 48, 20, 23, 13, respectively.

However, if we want to assess how the QRE model does as an integral explanation of our experimental observations, we must calculate $\hat{\lambda}$ for all ε simultaneously. In that case, the log-likelihoods of the observations for the various ε 's must be weighted by the reciprocal of the number of observations for that particular ε . Doing so, we get

$$\hat{\lambda} = 3.8$$

The QRE probability predictions for the players' actions, induced by $\lambda = 3.8$, are shown in Table 5.

(<i>ML</i> Estimates in %)	<i>S</i>	<i>C</i>	(<i>y, b</i>)	(<i>n, c</i>)	(<i>n, s</i>)
$\varepsilon = 1$	78	22	67	3	30
	94	6	83	1	16
$\varepsilon = 15$	69	31	60	8	32
	90	10	70	3	27
$\varepsilon = 30$	53	47	52	21	27
	41	59	46	26	28
$\varepsilon = 45$	9	91	19	77	3
	8	92	13	76	11
$\varepsilon = 60$	5	95	10	87	3
	6	94	6	87	8

Table 5: QRE predictions for $\hat{\lambda} = 3.8$ and observed frequencies

Also in this case the fit is better for large ε than for small ε . For $\varepsilon = 1, 15, 30, 45, 60$, the sum of the absolute differences is 64, 62, 36, 17, 11. More formally, the hypothesis that the empirical looking frequencies for $\varepsilon = 1$ (83%) has been generated by the maximum likelihood QRE (67%) is rejected with a p-value of .000. The analogous hypothesis for $\varepsilon = 15$ is rejected with a p-value of .009. In fact, for $\varepsilon = 1$, the QRE hypothesis can be rejected for every λ , with a p-value of .000. For $\varepsilon \geq 30$, the maximum likelihood QRE predictions do better. For $\varepsilon = 30, 60$, the QRE hypothesis with respect to the looking frequencies cannot be rejected. The p-values are .094 and .144, respectively. For $\varepsilon = 45$, the p-value is .040. This means that, at the 5% level, the QRE hypothesis is just rejected, but not very convincingly. The leader's behavior points in the same direction of bad fit for small cost and better fit for large cost. The empirical frequency \tilde{S} can be rejected as coming from the maximum likelihood QRE for $\varepsilon = 1, 15$ and 30, with p-values equal to .000, .000, .001, respectively. But for $\varepsilon = 45, 60$, the hypothesis cannot be rejected, with p-values equal to .639 and .615.

Finally, let us go back to the original interpretation of λ . The parameter λ measures the level of error associated to the players' best-response functions. When $\lambda = 0$, the players' actions are completely random, while for $\lambda = \infty$ there is no error in their actions whatsoever. If one hypothesizes that the error rate of individuals decreases with experience, one would expect $\hat{\lambda}$ to increase from early rounds to later rounds. In fact, an equivalent hypothesis could be formulated for the (weighted)

log-likelihoods. Table 6 shows that the data are consistent with both hypotheses.

(<i>ML</i> Estimates)	1 – 20	21 – 40	41 – 60	61 – 80	81 – 100
$\hat{\lambda}$	2.5	3.3	3.5	3.8	4.1
<i>Log L</i>	-7.5	-6.4	-6.2	-5.7	-5.4

Table 6: Maximum likelihood QRE estimates by round

Table 6 suggests that most learning is going on in the first twenty rounds. The jump in both $\hat{\lambda}$ and the (weighted) log-likelihood from rounds 1 – 20 to rounds 21 – 40 is greater than the total jump from rounds 21 – 40 to rounds 81 – 100. After the first twenty rounds, the ‘error reduction’ becomes much less, but remains positive.

To sum up, our tables and graphs suggest that for low cost treatments ($\varepsilon = 1, 15$), the QRE can explain the underlying pattern of the data to some extent, but does not get the empirical frequencies correct. For example, the probability of looking, p , which in h' is equal to $\Pr(y, b)$, does indeed become a decreasing function of ε . However, the dependence of p on ε is not strong enough to explain the looking frequencies as observed in the data. Also, for all λ , p stays below 68% no matter how small we choose ε . That is, even when we choose $\varepsilon = 0$. (See last graph of Figure 2).

For $\varepsilon = 45, 60$ and, to a lesser extent, 30, the QRE predictions perform well. QRE correctly predicts the switch from a predominance of Stackelberg outcomes to Cournot outcomes. Moreover, the maximum likelihood predictions of actions and looking probabilities are pretty close to the observed frequencies.

5.2 Risk Aversion

For low cost, we have seen that the QRE cannot fully explain the followers' looking behavior. So far, however, we have not controlled for risk preferences of the subjects. That is, all equilibrium predictions implicitly relied on the assumption of risk neutrality. We now investigate whether risk aversion can explain the high frequency of looking for small ε .

In order to introduce risk aversion into the model, we transform all the payoffs of the game used in the experiment by a concave, increasing function $m : \mathbb{R} \rightarrow \mathbb{R}$.

The payoff table then looks like:

$L \setminus F$	s	c
S	$m(500), m(200)$	$m(300), m(100)$
C	$m(600), m(300)$	$m(400), m(400)$

Note that this transformation has not changed the ordering of the payoffs. The underlying structure of the game has therefore remained the same.

Calculating the equilibrium probability p_m of looking, in the noisy Stackelberg equilibrium of the transformed game, we get:

$$p_m = \frac{m(600) - m(500)}{m(600) - m(400)}$$

which can be written as

$$p_m = \frac{m(600) - m(500)}{m(600) - m(500) + m(500) - m(400)}$$

Now, because m is concave:

$$m(500) - m(400) > m(600) - m(500)$$

It follows that

$$\begin{aligned}
 p_m &= \frac{m(600) - m(500)}{m(600) - m(500) + m(500) - m(400)} \\
 &< \frac{m(600) - m(500)}{m(600) - m(500) + m(600) - m(500)} \\
 &= \frac{1}{2}
 \end{aligned}$$

(For the noisy Cournot equilibrium we get $p_m > \frac{1}{2}$.) Therefore, risk aversion pushes the probability of looking in the wrong direction. The intuition is that in the noisy Stackelberg equilibrium L gets 500 for sure if he plays S . If he plays C , his payoff is uncertain. Hence, if L becomes risk averse, F must start looking less than before, in order for L to remain indifferent between S and C . For the noisy Stackelberg equilibrium, which, from an outcome perspective, seemed to apply for $\varepsilon = 1, 15$, risk aversion makes p smaller than $\frac{1}{2}$, while the data show that $\tilde{p} > \frac{1}{2}$. Therefore, risk aversion does not help in rationalizing excessive looking for small ε .

5.3 Decision-theoretic Play

Finally, what if subjects play this game as if it were a decision problem? What we mean by this is that a follower formulates some belief about the probability of L playing S or C . Then, he looks at his cost ε and calculates whether, for his beliefs, observing L 's action is worthwhile or not. Because L 's payoff does not depend on ε , from a naive, decision theoretic view point, F 's beliefs about L 's strategy will not depend on ε either. (Obviously, what is missing from this way of reasoning, is a strategic perspective. Followers should realize that player L can go through the same chain of thought, assess whether F will look or not, and subsequently adapt

his frequency of playing S and C . In turn, this should change F 's beliefs and so forth.)

To see whether such a more or less naive decision theoretic perspective could explain part of the puzzle why, for small observation costs, the probability of looking is so high, we now calculate for which beliefs $\check{\alpha}$ of F about $\text{Pr}_L(S)$ looking is a best reply, and how this depends on ε .

If F plays the pure strategy (n, s) while he believes that L plays S with probability $\check{\alpha}$, F 's expected payoff $\Pi_{(n,s)}^F(\check{\alpha})$ is:

$$\begin{aligned}\Pi_{(n,s)}^F(\check{\alpha}) &= 2\check{\alpha} + 3(1 - \check{\alpha}) \\ &= 2 - \check{\alpha}\end{aligned}$$

If F plays (n, c) , his expected payoff is:

$$\begin{aligned}\Pi_{(n,c)}^F(\check{\alpha}) &= \check{\alpha} + 4(1 - \check{\alpha}) \\ &= 4 - 3\check{\alpha}\end{aligned}$$

Finally, if F plays (y, b) :

$$\begin{aligned}\Pi_{(y,b)}^F(\check{\alpha}) &= 2\check{\alpha} + 4(1 - \check{\alpha}) - \varepsilon \\ &= 4 - 2\check{\alpha} - \varepsilon\end{aligned}$$

For $\varepsilon = 30$, we have plotted F 's payoffs of playing (n, s) , (n, c) and (y, b) as functions of $\check{\alpha}$. From Figure 3 it can be seen that (y, b) is a best for response for F for all $\check{\alpha} \in [.30, .70]$.

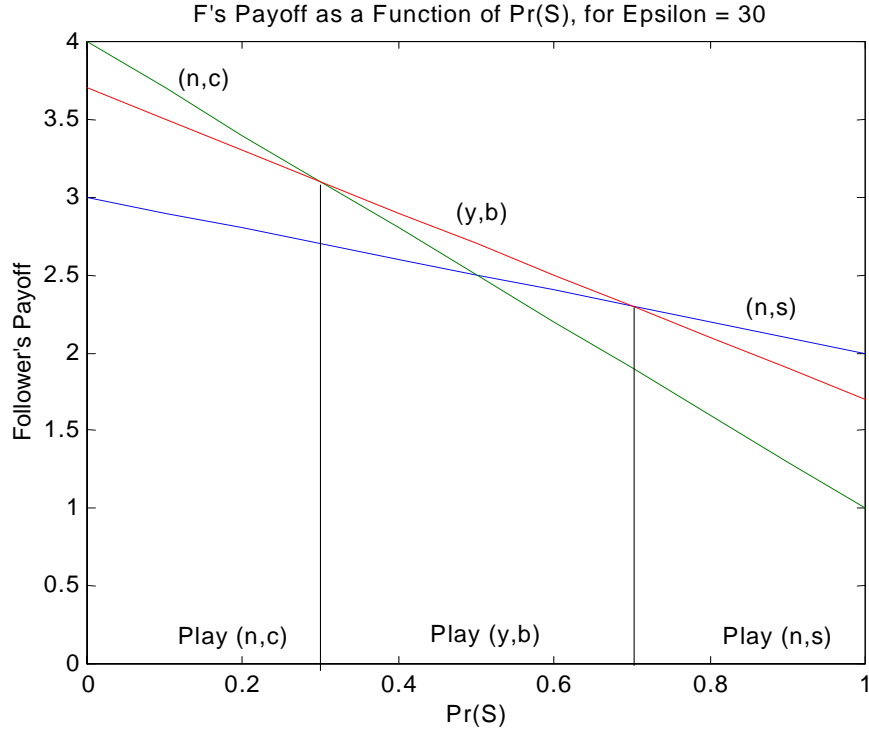


Figure 3: Best response map

More generally, (y, b) is a best reply when $\check{\alpha} \in [\varepsilon, 1 - \varepsilon]$. Among other things this implies that looking is irrational when $\varepsilon = 60$.

The belief interval $[\varepsilon, 1 - \varepsilon]$ for which (y, b) is a best reply is strictly decreasing in ε . Thus, if as postulated, F 's belief about $\Pr_L(S)$ does not depend much on ε , then the likelihood that $\check{\alpha}$ is an element of $[\varepsilon, 1 - \varepsilon]$ increases when ε decreases. Hence, the probability that the follower looks at the leader's action also increases. In fact, when $\varepsilon \downarrow 0$ then the predicted looking probability goes to 1.

Thus, a decision theoretic approach on the part of the follower that ignores the strategic effect of a change in ε could be an explanation for the strong dependence of the observed frequency of looking, \tilde{p} , on the cost ε , and for the high frequency

of looking when ε is very small.

We conclude this section by checking for which values of ε a follower should in fact have looked at the leader's action, given our data. From Tables 1 and Figure 3 we can conclude that F should have looked for $\varepsilon = 1$ and $\varepsilon = 30$. For $\varepsilon = 15$, he should have played s without looking, while for $\varepsilon = 45$ and 60 , his optimal strategy was to play c without looking. Finally, player L should have played S for $\varepsilon = 1, 15$, and C for $\varepsilon = 30, 45$, and 60 .

6 Conclusion

Our results suggest that, at least for small observation costs, the leader's value of commitment is robust. Once the follower's cost of observing the leader's action grows large, however, the theoretical possibility of preservation of the first-mover advantage is not borne out in the data. In a nutshell, variations in observation cost are a key determinant in whether commitment is valuable or not.

While the first-mover advantage is largely preserved in low cost treatments, it does not appear to be supported by equilibrium play. In particular, we can confidently reject the hypothesis that the empirical frequencies are being generated by the noisy Stackelberg equilibrium. Moreover, as the costs of becoming informed changes, the followers' empirical looking frequency changes systematically, a finding at odds with equilibrium theory.

The data is better fit in a model where quantal response equilibrium is used as solution concept. Under QRE, follower play is correctly predicted to depend on the

cost of being informed. However, certain aspects of the data cannot be explained by QRE. Specifically, we confidently reject the hypothesis that, for low costs of becoming informed, the empirical frequency of followers choosing to observe the leader's action derives from the maximum likelihood QRE. Indeed, for very low cost, this hypothesis is confidently rejected for every possible error parameter of that model.

Thus, we are left mostly with negative conclusions. Clearly, selecting equilibria on the basis of pure strategies does not lead to a good description of behavior. That said, the mixed-strategy noisy Stackelberg equilibrium is not a good description of behavior either. Adding other factors, such as risk aversion, actually moves the theory further from the data. QRE addresses some of these problems but cannot account for the high frequency with which the followers choose to observe the leader's action when the observation cost is small cost. It is this aspect of the data that seems to present the greatest difficulty for the existing theory.

In our view, the results reported here should suggest a change in the focus of the debate on the value of commitment. Specifically, it appears to us that the emphasis on equilibrium selection between Cournot and noisy Stackelberg equilibrium does not truly go to the heart of the matter. For the costly leader game, a satisfactory theory must deal with the systematic changes in looking behavior occurring as the cost parameter changes, and with the high frequencies with which the followers choose to observe the leader's choice when the cost of doing so is low. This remains a task for future research.

A Instructions

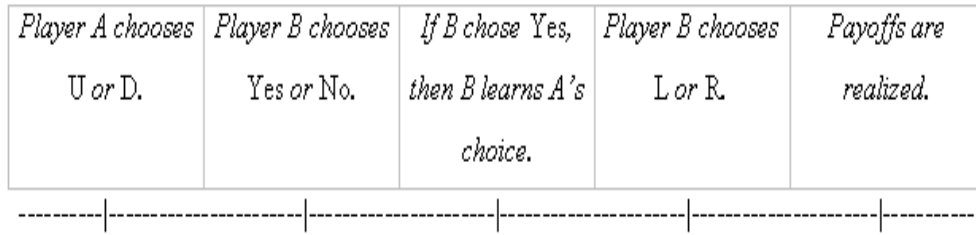
Thank you for participating in this experiment on the economics of decision making. If you follow the instructions carefully and make good decisions you can earn a considerable amount of money. At the end of the experiment you will be paid in cash and in private. The experiment will take about one hour and 15 minutes.

There are 10 people participating in this session. They have all been recruited in the same way that you have and are reading the same instructions that you are for the first time. Please refrain from talking to the other participants during the experiment.

You are about to play the same game 100 times in succession. In each round you will be paired with a co-player. These pairings are random. Neither you nor your co-player will know who you are paired with. And since the pairings are random your co-player will change from one round to the next.

In each round, you will be randomly assigned the role of Player A or the role of Player B. Whatever role you are assigned, your co-player is assigned the opposite role. In any given round, it is equally likely that you are a Player A or a Player B.

The following Time Line summarizes the play of the game:



First, player A chooses U(p) or D(own). Then, player B decides whether he wants to observe player A's choice, i.e., he chooses Yes or No. If B chose Yes, then Player B is informed of A's choice. Finally, player B chooses either L(eft) or R(ight).

The cost to Player B of choosing Yes is listed at the top of the screen. It will change from round to round. If Player B chooses to observe Player A's choice, Player A's choice will be marked by an arrow.

The payoffs in this game are summarized in the following matrix. The first number in a cell represents Player A's payoff, while the second number represents Player B's payoff.

		Player B's Choice	
		L	R
Player A's Choice	U	500, 200	300, 100
	D	600, 300	400, 400

This matrix should be read as follows. Suppose Player A has chosen U(p) and player B has chosen L(eft). Then, player A gets a payoff of 500, while Player B

gets a payoff of 200. The other matrix entries have a similar interpretation. Thus, if Player A chooses D(own) and Player B chooses R(ight), Player A receives 400 and Player B receives 400. If Player A chooses U(p) and Player B chooses R(ight), Player A receives 300 and Player B receives 100. Finally, if Player A chooses D(own) and Player B chooses L(eft), Player A receives 600 and Player B receives 300.

In addition, Player B's payoff depends on whether or not he chose to observe player A's choice. If player B observed A's choice, his payoff is reduced by the cost listed at the top of the screen.

Player B's cost of observing A's choice can take on any of five possible values: 1, 15, 30, 45, or 60 points. The value in a particular round is determined at random. All values are equally likely to occur in any given round.

Your cash earnings are calculated from the total points you earned during experiment. For every 1000 points you earn 50 Cents. For example, if you score a total of 20.000 points, you earn \$10.

References

- [1] Bagwell, K. (1995). ‘Commitment and Observability in Games,’ *Games and Economic Behavior* 8, 271–280.
- [2] van Damme, E., and Hurkens, S. (1997). ‘Games with Imperfectly Observable Commitment,’ *Games and Economic Behavior*, 21, 282–308.
- [3] Güth, W., Kirchsteiger, G., and Ritzberger, K. (1998). ‘Imperfectly Observable Commitments in n -Player Games,’ *Games and Economic Behavior* 23, 54–74.
- [4] Hogg, R.V., and Craig, A.T. (1995). *Introduction to Mathematical Statistics*. Upper Saddle River, NJ: Prentice-Hall.
- [5] Huck, S., and Müller, W. (2000). ‘Perfect versus Imperfect Observability—An Experimental Test of Bagwell’s Result,’ *Games and Economic Behavior* 31, 174–190.
- [6] McKelvey, R., and Palfrey, T. (1995). ‘Quantal Response Equilibria for Normal Form Games,’ *Games and Economic Behavior* 10, 6–38.
- [7] Oechssler, J. and Schlag, K. (1997). ‘Loss of Commitment: An Evolutionary Analysis of Bagwell’s Example,’ Discussion Paper no. B-410, Sonderforschungsbereich 303, University Bonn.
- [8] Schelling, T. (1960). *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

- [9] Selten, R. (1975). 'Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games,' *International Journal of Game Theory* 4, 25–55.
- [10] Von Stackelberg, H. (1934). *Marktform und Gleichgewicht*, Vienna/Berlin: Springer-Verlag.
- [11] Várdy, F. (2001). 'Stackelberg Leader Games with Observation Costs,' in: Ph.D. dissertation, Princeton University.
- [12] Weibull, J.W. (1997). *Evolutionary Game Theory*. Cambridge, MA: MIT Press.