Tilburg University

Oligopoly limit-pricing in the lab

Müller, W.; Spiegel, Y.; Yehezkel, Y.

Published in:
Games and Economic Behavior

Publication date:
2009

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Müller, W., Spiegel, Y., & Yehezkel, Y. (2009). Oligopoly limit-pricing in the lab. *Games and Economic Behavior*, 66(1), 373-393.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Oligopoly limit-pricing in the lab

Wieland Müller^{a,b,c}, Yossi Spiegel^{d,*}, Yaron Yehezkel^d

^a Department of Economics, Tilburg University, PO Box 90153, 5000LE Tilburg, The Netherlands

^b CentER, Tilburg, The Netherlands

^c TILEC, Tilburg, The Netherlands

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

Received 12 September 2006

Available online 3 July 2008

Abstract

We examine the behavior of senders and receivers in the context of oligopoly limit pricing experiments in which high prices chosen by two privately informed incumbents may signal to a potential entrant that the industry-wide costs are high and that entry is unprofitable. The results provide strong support for the theoretical prediction that the incumbents can credibly deter unprofitable entry without having to distort their prices away from their full information levels. Yet, in a large number of cases, asymmetric information induces incumbents to raise prices when costs are low. The results also show that the entrants' behavior is by and large "bi-polar:" entrants tend to enter when the incumbents' prices are "low" but tend to stay out when the incumbents' prices are "high."

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C72; C91

Keywords: Oligopoly limit pricing; Multi-sender signaling; Full information equilibrium; Unprejudiced beliefs; Experiments; Strategy method

1. Introduction

A common feature of signaling games is that senders need to distort their actions away from their full information levels in order to credibly signal their private information to receivers. In an early application of signaling models to Industrial Organization, Milgrom and Roberts (1982) used this feature to explain the rationale behind limit pricing behavior: an incumbent will cut his price before entry occurs in order to signal his own low cost to a potential entrant and hence deter the latter's entry into the market. This prediction has received experimental support by Cooper et al. (1997a, 1997b).

However, many real-life applications involve more than one privately informed player. For example, in oligopolistic markets, privately informed firms may wish to use their prices or output levels as a way to signal to a potential entrant that entry is unprofitable. Likewise, oligopolistic firms may wish to use their investments in (uninformative)

* Corresponding author.

E-mail addresses: w.mueller@uvt.nl (W. Müller), spiegel@post.tau.ac.il (Y. Spiegel), yehezkel@post.tau.ac.il (Y. Yehezkel).

URLs: <http://center.uvt.nl/staff/muller/> (W. Müller), <http://www.tau.ac.il/~spiegel> (Y. Spiegel), <http://www.tau.ac.il/~yehezkel/> (Y. Yehezkel).

advertising as a way to signal to consumers that the quality of their products is high. This begs the obvious question of whether results obtained in single-sender games carry over to settings with multiple senders.

Harrington (1987) shows that the answer to this question is “yes.” He studies a Cournot oligopoly limit pricing model in which two incumbents with the same private information choose their output levels, and based on the resulting price, a potential entrant decides whether or not to enter. Harrington shows that there exists an equilibrium in which the incumbents deter entry by distorting their output levels downward.¹ An important feature in Harrington’s model is that the entrant observes only an aggregate signal of the incumbents’ choices (the market price), but not their individual choices (the output level of each incumbent). Bagwell and Ramey (1991) show that this feature is critical: when the entrant observes the incumbents’ individual prices, the incumbents may be able to credibly deter unprofitable entry without having to distort their prices away from their full information levels. Moreover, they show that this “full information equilibrium” is the only separating equilibrium that survives a belief-based refinement, called unprejudiced beliefs (UPB), which is based on the minimal number of deviations needed to generate a particular out-of-equilibrium price combination.²

Bagwell and Ramey’s (1991) results are striking and indicate that there may be a fundamental difference between multi-sender and single-sender signaling games. Yet, beside the “full information equilibrium,” the oligopoly limit pricing game that Bagwell and Ramey consider also admits a continuum of separating and pooling equilibria that involve both upward as well as downward price distortions. Given this multiplicity of equilibria, it is natural to wonder whether multiple senders can nonetheless credibly signal their private information while playing their “full information” strategies. And, if they cannot, then which type of behavior emerges when receivers face multiple senders who share the same piece of private information? Moreover, how do receivers react when they receive conflicting signals from identically informed senders?

To address these questions, we conducted a series of experiments based on a variant of the Bagwell and Ramey (1991) oligopoly limit pricing model. Subjects were randomly matched in groups of three: two incumbents and one entrant. The two incumbents privately observed whether the common marginal cost of all three firms is high or low and then independently chose their prices. The entrant observed the incumbents’ prices and then had to decide whether or not to enter. Entry was profitable only when the marginal cost was low. We ran both full information treatments in which the entrant was informed about the realized marginal cost and asymmetric information treatments in which the entrant was uninformed about the realized marginal cost and had to infer it from the incumbents’ prices.

The main findings in our experiments are as follows. First, asymmetric information does not have a significant effect on the subjects’ behavior. In particular, the subjects’ behavior both under full as well as under asymmetric information is by and large consistent with the full information equilibrium. This means that when information is asymmetric, incumbents are able to credibly deter entry when cost is high without having to distort their prices away from the full information prices. These results provide strong support for Bagwell and Ramey’s (1991) theoretical results and are in stark contrast to earlier experimental results from single-sender signaling games (see Camerer, 2003, Chapter 8).

Second, while Bagwell and Ramey (1991) derive the full information equilibrium by imposing the UPB refinement, we find that the entrants’ behavior in our experiments is inconsistent with the UPB refinement. Nonetheless, the entrants’ behavior is consistent with the main idea behind the UPB refinement which is that small deviations from an equilibrium should have a small effect on the entrants’ beliefs. In particular, the entrants’ behavior in our experiments seems to follow a “bi-polar” pattern: entrants tend to enter when the incumbents’ prices are “close” to the full information prices when costs are low, but tend to stay out when the incumbents’ prices are “close” to the full information prices when costs are high. Moreover, entry rates fall (increase) with the Euclidean distance from the full information

¹ In Harrington (1987), which extends the monopoly setting of Harrington (1986), the incumbents and the entrant have the same marginal cost and entry is unprofitable when this cost is high. Consequently, the incumbents deter entry by producing a small quantity which signals high marginal cost to the entrant. This is in contrast with Milgrom and Roberts (1982), where the incumbent’s and the entrant’s costs are uncorrelated; hence, in their model, low prices, which signal that the incumbent’s marginal cost is low, deter entry. Still, in all three models, prices are distorted away from their full information levels in order to credibly signal that entry is unprofitable.

² The Bagwell and Ramey model has been extended by Schultz (1999) who studies the case where one incumbent wishes to deter entry while the other wishes to encourage it and by Martin (1995) who studies the case where the firms’ costs are imperfectly correlated and the incumbents do not know each others’ costs. Additional multi-sender signaling games have been studied by Matthews and Fertig (1990), de Bijl (1997), Hertzendorf and Overgaard (2002), and Fluet and Garella (2002) in the context of advertising, and by Schultz (1996) and Martinelli and Matsui (2002) in the context of electoral competition.

prices under low (high) costs and is affected in general by the combination of the incumbents' prices rather than by only one of these prices (e.g., the minimal or the maximal price).

The remainder of the paper is organized as follows. Section 2 describes our experimental design and establishes our main research questions. The experimental implementation is described in Section 3. In Sections 4 and 5 we discuss our experimental results. Concluding remarks are in Section 6. Appendix A contains the written instructions that were given to the subjects.

2. Experimental design and main research questions

Our experiments are based on the Bagwell and Ramey (1991) model of oligopoly limit pricing. There are two incumbents and one potential entrant. All three firms have the same marginal cost, which is either low or high with equal probabilities. The strategic interaction between the three firms evolves over two periods. At the beginning of period 1, the two incumbents privately observe whether marginal cost is high (state H) or low (state L), and based on this information, they simultaneously choose their period-1 prices. The potential entrant observes the incumbents' period-1 prices, updates his belief about whether marginal cost is high or low, and decides whether or not to enter the market. If the entrant decides to enter, then all three firms compete in period 2 by choosing prices. If the entrant stays out, only the two incumbents compete in period 2. Entry is profitable only when cost is low.

The crucial feature of this game is that there are two senders (the two incumbents) who have the same private information which they signal to the receiver (the entrant) through their simultaneous price choices. As a result, the receiver observes two independent signals about the same unknown variable rather than just one signal. One important implication of this is that in state L , a monopolist incumbent can fool the entrant into believing that entry is unprofitable by mimicking the state H prices. However, when there are two incumbents, both of them need to deviate in state L in order to mimic the state H —a unilateral deviation from the state L prices is insufficient for generating the (pair of) state H prices. This feature creates a wide variety of interesting coordination problems that distinguish our experimental design from previous signaling games with a single sender that have been experimentally tested in the lab.

Although the Bagwell and Ramey (1991) model, as well as other limit pricing model like Milgrom and Roberts (1982) and Harrington (1987), has two periods, the interesting signaling interaction occurs only in period 1; period 2 simply involves a straightforward oligopoly game. In order to focus on the equilibrium predictions of the underlying oligopoly signaling game without needlessly complicating our experimental design, we therefore follow Cooper et al. (1997a, 1997b) and collapse the original two-period model into a single-period game by adding the period-2 equilibrium payoffs to players' period-1 earnings. Specifically, the two incumbents were only asked to choose their period-1 prices from the set $\{30, 40, 50, 60, 70, 80\}$, and the entrant was only asked to make an entry decision after observing the pair of incumbents' prices. We then computed the resulting Nash equilibrium payoffs in period-2 (either a 3-firm oligopoly game in case of entry or a duopoly game in case of absent entry) and added them to the duopoly period-1 profits of the two incumbents. The entrant's payoff was either equal to his Nash equilibrium payoff in case of entry or to $\bar{\pi}$ in case he decided to stay out.³ The value of $\bar{\pi}$ was equal to 10 in some treatments or to 14 in others.

Fig. 1 presents the actual payoff tables that were used in our experiments.⁴

Fig. 1 shows that the incumbents have 4 relevant payoff tables, depending on whether the marginal cost is low (State L) or high (State H) and on whether entry does or does not take place. As mentioned earlier, the two incumbents were privately informed about the relevant cost state. Hence, in choosing their strategies, the incumbents had to compare either the two upper tables (in state L) or the two lower tables (in state H). Moreover, notice that the payoff of each incumbent increases with the price of the rival incumbent, and is first increasing and then decreasing with the incumbent's own price.⁵ These features may also fit other types of multi-sender signaling games. Hence, although the

³ To make the entrant's payoff from entry more equal across the two cost states, we assumed that the entrant needs to pay an entry cost of 1 in state L and 1.5 in state H .

⁴ The payoff tables are derived from a variant of the Bagwell and Ramey (1991) model. For details, see the Technical Appendix, available at <http://www.tau.ac.il/~spiegel>. In order to make the instructions more "neutral," the incumbents were asked to choose one of the numbers in the set $\{1, 2, 3, 4, 5, 6\}$, where 1 corresponds to a price of 30, 2 corresponds to a price of 40, and so on. In addition, incumbents were referred to as "A-participants," the entrant as "B-participant," states H and L as "state 1" and "state 2," respectively, and the entrants' decisions were referred to as "X" (entry) and "Y" (no entry).

⁵ It is also easy to see from Fig. 1 that the joint payoff of the two incumbents is maximized at (60, 60) in state L and at (70, 70) in state H .

		State L / Entry						State L / No entry					
		Incumbent 2's Price						Incumbent 2's Price					
		30	40	50	60	70	80	30	40	50	60	70	80
Incumbent 1's Price	30	42 42	45 43	48 38	52 27	55 18	58 18	53 53	56 55	60 50	63 38	66 30	70 30
	40	43 45	48 48	53 45	58 35	63 18	68 18	55 56	60 60	65 56	70 46	75 30	80 30
	50	38 48	45 53	50 50	58 43	65 28	72 18	50 60	56 65	61 61	70 55	76 40	83 30
	60	27 52	35 58	43 58	52 52	60 38	68 18	38 63	46 70	55 70	63 63	71 50	80 30
	70	18 55	18 63	28 65	38 60	48 48	58 30	30 66	30 75	40 76	50 71	60 60	70 41
	80	18 58	18 68	18 72	18 68	30 58	42 42	30 70	30 80	30 83	30 80	41 70	53 53

		State H / Entry						State H / No entry					
		Incumbent 2' Price						Incumbent 2's Price					
		30	40	50	60	70	80	30	40	50	60	70	80
Incumbent 1's Price	30	-3 -3	-4 8	-6 13	-8 11	-9 8	-11 8	1 1	0 13	-1 18	-3 16	-5 13	-6 13
	40	8 -4	8 8	8 15	8 15	8 8	8 8	13 0	13 13	13 20	13 20	13 13	13 13
	50	13 -6	15 8	16 16	18 18	20 13	21 8	18 -1	20 13	21 21	23 23	25 18	26 13
	60	11 -8	15 8	18 18	21 21	25 18	28 8	16 -3	20 13	23 23	26 26	30 23	33 13
	70	8 -9	8 8	13 20	18 25	23 23	28 15	13 -5	13 13	18 25	23 30	28 28	33 20
	80	8 -11	8 8	8 21	8 28	15 28	21 21	13 -6	13 13	13 26	13 33	20 33	26 26

Payoff table for the Entrant

		State L	State H
Your Choice	Enter	18	6
	Stay Out	$\bar{\pi}$	$\bar{\pi}$

Fig. 1. Payoff tables for the incumbents and the entrant. As usual, the upper left number in each cell of the incumbents' tables is the row players' payoff and the lower right number is the column player's payoff.

payoffs tables were derived from an oligopoly limit pricing game, we believe that our experiments can shed light on a broad class of multi-sender signaling games, including games in which competing firms signal their product qualities to consumers, or electoral competition games in which competing candidates signal some policy relevant state of nature to voters.

As for the entrant, Fig. 1 shows that following entry, the entrant's payoff is higher in state *L* than in state *H*. As mentioned earlier, the entrant's payoff when staying out, $\bar{\pi}$, was equal either to 10 or to 14, independent of the cost state.

2.1. The equilibria in our experimental design

Under full information, the entrant should enter in state L but stay out in state H . Since this decision is independent of the incumbents' prices, the two incumbents will simply play the Nash equilibrium prices which are (40, 40) in state L and (60, 60) in state H .

Under asymmetric information, our experimental design admits both separating and pooling equilibria. We consider separating equilibria first. In these equilibria, the two incumbents play different prices in state L and in state H . Consequently, the entrant can infer the state from the incumbent's prices and hence enters only in state L . Since entry takes place in state L , two incumbents can do no better than play the full information equilibrium, (40, 40). On the other hand, in state H , at least one incumbent sets a price different than 40 (this is enough to ensure separation). Let $(\hat{p}_i, \hat{p}_j)$ be the incumbents' prices in state H in a separating equilibrium, where $\hat{p}_i$ or $\hat{p}_j$ or both are different than 40. The equilibrium is supported by the entrant's belief that any deviation from $(\hat{p}_i, \hat{p}_j)$ is associated with state L ; hence any deviation from $(\hat{p}_i, \hat{p}_j)$ will trigger entry. A necessary and sufficient condition for $(\hat{p}_i, \hat{p}_j)$ to be a separating equilibrium choice in state H is that $\pi_i(\hat{p}_i, \hat{p}_j, Out) \geq \pi_i(BR_i(\hat{p}_j), \hat{p}_j, In)$, where $\pi_i(\hat{p}_i, \hat{p}_j, Out)$ is the equilibrium payoff of incumbent i given that in equilibrium the entrant stays out and $\pi_i(BR_i(\hat{p}_j), \hat{p}_j, In)$ is incumbent i 's payoff when he deviates from $\hat{p}_i$ by playing a best response against $\hat{p}_j$ and given that the deviation triggers entry. In other words, the necessary and sufficient condition requires that no incumbent can profitably deviate if the deviation triggers entry.

To illustrate, suppose that the state is H and the incumbents' prices are (80, 80). Given that there is no entry in equilibrium, the payoff of each incumbent is 26. However, if one incumbent deviates unilaterally to 60 or to 70, then even if entry is induced, the incumbent's payoff increases to 28 (if entry is not induced, the incumbent's payoff increases to 33). Clearly, then (80, 80) cannot be an equilibrium. On the other hand, (70, 70) can be chosen in a separating equilibrium in state H . To see why, suppose that the entrant believes that the state is L unless the two incumbents play (70, 70). If the incumbents indeed play (70, 70), then given that the entrant stays out, each incumbent gets a payoff of 28. However, once the incumbent deviates from 70, entry takes place and the incumbent's payoff is at most 25. Hence, when supported by appropriate out-of-equilibrium beliefs, (70, 70) is immune to unilateral deviations and hence can be chosen in a separating equilibrium. Using this logic, one can verify that in a separating equilibrium, the incumbents prices in state H must belong to the set $\{(50, 50), (50, 60), (50, 70), (60, 60), (60, 70), (70, 70)\}$. It is worth noting that with the exception of (60, 60), prices in this set are not mutual best responses. These prices are nonetheless played in separating equilibria because they are supported by the entrant's belief that any deviation from the equilibrium is associated with state L ; hence any deviation will trigger entry.

We now turn to pooling equilibria. In these equilibria, each incumbent plays the same price in states L and H (the two incumbents though need not choose the same prices). Consequently, after observing the incumbents' prices, the entrant maintains his prior belief that the two states are equally likely and therefore expects that his payoff from entry will be $18 \times 0.5 + 6 \times 0.5 = 12$. Assuming that entrants are not too risk-averse or too risk-lovers, entry will take place if the entrant's payoff from staying out, $\bar{\pi}$, is 10 but not when $\bar{\pi} = 14$.⁶ Since entry takes place when $\bar{\pi} = 10$, the incumbents might as well play the full-information equilibrium. This equilibrium however is not pooling since the incumbents' play (40, 40) in state L and (60, 60) in state H ; hence there are no pooling equilibria when $\bar{\pi} = 10$.

Next suppose that $\bar{\pi} = 14$, so in equilibrium the entrant stays out. Suppose that the incumbents' prices are $(\hat{p}_i, \hat{p}_j)$ in both states. Then, similarly to separating equilibria, a necessary and sufficient condition for a pooling equilibrium is that $\pi_i(\hat{p}_i, \hat{p}_j, Out) \geq \pi_i(BR_i(\hat{p}_j), \hat{p}_j, In)$. That is, neither incumbent can increase his payoff by playing a best-response against his rival's price when the deviation triggers entry, no matter what the state is. For instance, suppose that the incumbents' prices are (70, 70). Given that there is no entry in a pooling equilibrium, the payoff of each incumbent in state L is 60. If an incumbent deviates to 50, his payoff increases to 65 even though the deviation triggers entry. Clearly then, (70, 70) cannot be a pooling equilibrium. On the other hand, (60, 60) can be a pooling equilibrium: if the entrant holds an (out-of-equilibrium) belief that the state is L unless the incumbents' prices are (60, 60), then any deviation away from (60, 60) will trigger entry. Hence a unilateral deviation away from 60 will lead to a payoff of at most 58 which is below the equilibrium payoff of 63. Using this logic, one can verify when $\bar{\pi} = 14$, the incumbents' prices in a pooling equilibrium must belong to the set $\{(50, 50), (50, 60), (60, 60)\}$ in both states.

⁶ Clearly, a risk averse entrant may decide to stay out even when $\bar{\pi} = 10$, while a risk-loving entrant may decide to enter even if $\bar{\pi} = 14$.

Table 1
Equilibrium price pairs in our experiments

$\bar{\pi} = 10$							$\bar{\pi} = 14$						
Price	30	40	50	60	70	80	Price	30	40	50	60	70	80
30	–	–	–	–	–	–	30	–	–	–	–	–	–
40		L	–	–	–	–	40		L	–	–	–	–
50			H	H	H	–	50			P, H	P, H	H	–
60				H	H	–	60				P, H	H	–
70					H	–	70					H	–
80						–	80						–

Note: H—incumbents' separating prices in state H, L—incumbents' separating prices in state L, P—incumbents' pooling prices. The full information equilibrium is in bold font.

Table 1 summarizes the various equilibrium prices in our experiments. We use “H” to indicate separating equilibrium prices in state H , “L” to indicate separating equilibrium prices in state L , and “P” to indicate pooling equilibrium prices. For instance, when $\bar{\pi} = 14$, (50, 50) is part of both a separating equilibrium in which (50, 50) is chosen in state H and part of a pooling equilibrium in which (50, 50) is chosen in both states. Since the incumbents are symmetric, we only present distinct price pairs (i.e., (50, 60) and (60, 50) are considered to be the same price pair). Entry takes place only following the price pair (40, 40). Bold face letters indicate the full information equilibrium prices.

Given the multiplicity of separating equilibria, Bagwell and Ramey (1991) propose a belief-based refinement which they call “Unprejudiced Beliefs” (UPB). The refinement works as follows: suppose that the entrant believes that the incumbents play (40, 40) in state L and $(\hat{p}_1, \hat{p}_2)$ in state H . What should the entrant believe if instead he observes the price pair $(\hat{p}_1, p_2)$, where $p_2 \neq \hat{p}_2$ and $p_2 \neq 40$? One possibility is that the state is H and incumbent 2 unilaterally deviated from $\hat{p}_2$. A second possibility is that the state is L and both incumbents deviated from (40, 40). Bagwell and Ramey (1991) argue that the first possibility is more reasonable because it involves a single deviation rather than two as in the second possibility. Accordingly, the UPB refinement requires that the entrant will maintain his beliefs following a unilateral deviation from a putative equilibrium. But then, if a unilateral deviation does not alter the entrant's beliefs, the incumbents' prices must be mutual best responses, otherwise each incumbent can profitably deviate in state H (the deviation does not trigger entry because it does not alter the entrant's beliefs that entry is unprofitable). This implies in turn that the UPB refinement eliminates all separating equilibria except the full information equilibrium. This result is striking because it implies that with two incumbents, all separating equilibria which involve distortions away from the full information prices can be eliminated.⁷ For example, (70, 70) can no longer be an equilibrium in state H because at least one incumbent will deviate from 70 to 60. Under the UPB refinement, the deviation will not alter the entrant's belief that the state is H and will therefore not trigger entry. As a result, the deviating incumbent will increase his payoff from 28 to 30. Note that the UPB refinement has no bite in the case of pooling equilibria, because each incumbent chooses the same prices in both states, so any non-equilibrium price pair is the same number of deviations away from the price choices in state L and the price choices in state H .

2.2. Research questions

Having fully characterized the equilibria in our experimental design, we are ready to state the main research questions that we are going to address with our experimental data. The first research question concerns the case where $\bar{\pi} = 10$. As Table 1 shows, in this case we only have separating equilibria. The question that we ask in this case is whether the subjects in our experiments learn to play one or more of the separating equilibria in Table 1?

The second research question concerns the case where $\bar{\pi} = 14$. As Table 1 shows, here we have both separating and pooling equilibria. Therefore the question that we ask in this case is whether the subjects learn to play a separating or a pooling equilibrium, and if so which equilibrium they learn to play?

⁷ In fact, Yehezkel (2007) shows that the full information equilibrium is the only separating equilibrium that survives the UPB refinement in any multi-sender signaling game in which the senders know each other's types.

Related to the first two research questions is our third research question: is the subjects' behavior under asymmetric information similar to their behavior under full information, as Bagwell and Ramey predict, or does asymmetric information affect the subjects' behavior in a systematic manner?

As we shall see, our experiments provide a strong support for Bagwell and Ramey's (1991) prediction that the full information equilibrium emerges even under asymmetric information. Bagwell and Ramey obtain this prediction by using the UPB refinement. Our fourth research question then is whether the entrants' behavior is indeed consistent with the UPB refinement? And, to the extent that its is not, what are the main determinants of the entrants' behavior?

3. Experimental implementation

We ran 14 sessions of the oligopoly limit-pricing game, using the software tool kit *z-Tree* (Fischbacher, 2007). All sessions were held at Tilburg University in September–October 2003 and in May 2005. Each session included 12 different subjects who were randomly recruited students from various departments, mainly from Economics and Business Administration.

Upon arrival at the lab, subjects were assigned a computer screen and received written instructions in English (see Appendix A) which they read in private. After reading the instructions, subjects were allowed to ask clarifying questions (which were answered in private) and were asked to fill in a short questionnaire to ensure that they understand the instructions. Sessions consisted of 48 rounds. At the beginning of each round, we randomly created four groups of three subjects: two who played role A (incumbents) and one who played role B (entrant). Following Cooper et al. (1997a, 1997b), subjects switched roles. In our design, role switching took place every 8 rounds. As a result, each subject played role A for exactly 32 rounds and role B for exactly 16 rounds.⁸ The purpose of role switching was to enhance subjects' learning and enable them to better understand the decision problem of subjects in the other player role and therefore the overall game. The random matching was meant to minimize potential repeated game effects (our experiments are based on a “one shot” oligopoly signaling game that does not feature repeated game effects).⁹

We implemented 6 different treatments which are summarized in Table 2. The treatments differed with respect to the information structure (full or asymmetric information), the entrants' payoff from staying out, $\bar{\pi}$ (pooling equilibria exist only when $\bar{\pi} = 14$), and the experimental method (sequential play or strategy method).

In the four sequential play (SP) treatments, the incumbents were informed about the cost state and were asked to simultaneously choose prices. Then, before making their entry decision, entrants observed the prices of the incumbents with whom they were matched, and in the two full information treatments, were also informed about the state. At the end of each round, subjects were informed about the incumbents' prices, the entrant's decision, the realized state, their own profit in the last round, and their own cumulative profit. The SP sessions lasted about 90 minutes.

In the two strategy method (SM) treatments, incumbents and entrants were simultaneously asked to specify their complete strategies without observing the cost state: incumbents were asked to specify a price for each of the two possible states, while entrants were asked to specify an entry decision for each of the possible 21 distinct incumbents' price pairs. Then the subjects were informed about the realized cost state, the incumbents' prices for the realized

Table 2
Overview of the experimental design

Treatment	# of sessions	# of subjects	$\bar{\pi}$	Equilibrium prediction	Experimental method
FULL10	2	24	10	unique SPE	sequential play
ASYM10	3	36	10	separating equilibria	sequential play
FULL14	2	24	14	unique SPE	sequential play
ASYM14	3	36	14	pooling & separating equilibria	sequential play
ASYM10-SM	2	24	10	separating equilibria	strategy method
ASYM14-SM	2	24	14	pooling & separating equilibria	strategy method

⁸ With 48 rounds, we had 6 blocks of 8 rounds. Roles were fixed within each block. The sequence of roles over the 6 blocks were AABAAB for a third of the subjects, ABAABA for another third, and BAABAA for the last third.

⁹ On average, the same two incumbents played against each other in 2.9 rounds out of the 32 rounds in which each subject played the role of an incumbent.

state (but not the prices for the other state), the entry decision given the incumbents' actual prices (but not the entry decisions for all other 20 possible price pairs), the subjects' own profit, and their own cumulative profit. To ensure that subjects have a clear idea of the timing in the original asymmetric information treatments, we started each SM session with six additional rounds of sequential play as in Treatments ASYM10 and ASYM14.¹⁰ Moreover, to facilitate the entrants' task and speed up things, the most recent strategy of each entrant was displayed on the entrant's screen as a default at the beginning of every round; entrants were then free to change this default.¹¹ The SM sessions lasted about 120 minutes.

Subjects were paid 1 Euro for every 70 points in the SP sessions, and 1 Euro for every 50 points in the longer and more complicated SM sessions.¹² The average earnings in the SM sessions were about 35 Euros compared to 21 Euros in the SP sessions.

4. Results from the sequential play (SP) treatments

In this section we consider the results from the SP treatments. The results from the SM treatments are discussed in the next section. We will first collect a number of facts about subjects' behavior that will eventually enable us to answer the first three research questions posed in Section 2.2. We begin with the incumbents' behavior. Fig. 2 provides information about the incumbents' price choices. Each row corresponds to a single treatment. The left diagram in each row shows the evolution of average prices in states L and H , while the middle and right diagrams show histograms of price choices.¹³

The histograms in Fig. 2 reveal that incumbents mostly chose the full information prices (40 in state L and 60 in state H) both under full as well as under asymmetric information. These choices were more frequent in rounds 25–48 and were more frequent in treatments with $\bar{\pi} = 10$ than in the corresponding treatments with $\bar{\pi} = 14$. Thus, the incumbents' behavior provides strong support for Bagwell and Ramey's (1991) prediction that the oligopoly limit pricing model gives rise to a full information equilibrium behavior, especially in treatments with $\bar{\pi} = 10$ which do not admit pooling equilibria. Moreover, the left diagrams in Fig. 2 show that in all treatments, the incumbents' average prices in state H were remarkably close to the full information price of 60 right from the start, whereas in state L , they exhibited a moderate, but statistically significant, downward time trend towards the full information price of 40.¹⁴ Together with the fact that the standard deviations of the incumbents' prices were smaller in rounds 25–48 than in rounds 1–24, this suggests that as sessions have progressed, incumbents gradually "learned" to play the full information equilibrium.

While Fig. 2 provides strong support for the full information equilibrium, the histograms in this figure also reveal a substantial fraction of upward deviations in state L from the full information price of 40 to 50, and, to a lesser extent, to 60. As Table 1 shows, such price choices in state L are inconsistent with separating equilibrium behavior. The frequency of the upward deviations in state L was larger in treatments with $\bar{\pi} = 14$ (which admit pooling equilibria) than in the corresponding treatments with $\bar{\pi} = 10$ (which do not admit pooling equilibria) and was also larger in the asymmetric information treatments than in the corresponding full information treatments. To examine whether these cross-treatment differences are significant, we estimate the following random effects ordered probit equations:

$$p_{ist} = \beta \times TREATM + \alpha_i + \varepsilon_{it},$$

where p_{ist} is incumbent i 's price in round t when the state is $s = L, H$, $TREATM$ is a treatment dummy, α_i is a fixed subject-specific component, and ε_{it} is a subject-specific error term that may vary across observations from the same subject. For example, when testing for the differences between Treatments FULL10 and ASYM10 in state L ,

¹⁰ More precisely, subjects were informed that the experiment would consist of two "phases" and that they would be informed about the second phase only after the completion of the first phase. To implement role-switching in the first phase, each subject played 4 rounds as an incumbent and 2 rounds as an entrant. For more detail, see <http://www.tau.ac.il/~spiegel>.

¹¹ A possible drawback of this feature is that the entrants' may get lazy and fail to update their screens. However, if this is true then asking the entrants to fill in their entire entry strategies from scratch in every round is equally or even more problematic since the entrants may get lazy and fail to think hard about their entry strategies.

¹² To prevent bankruptcy problems, we gave each subject 100 points at the start of the experiment. It turned out, however, that incumbents got a negative payoff in only 9 out of 5376 price choices that were made in our experiments.

¹³ See also the additional data on the session level available at <http://www.tau.ac.il/~spiegel/papers/MSY-technicalappendix.pdf>.

¹⁴ The Pearson correlation between the round numbers and the average prices in state L was significant at the 1% level in all treatments.

Fig. 2. Incumbents' behavior in the sequential play (SP) treatments.

Table 3

p-values of (two tailed) pairwise cross-treatment differences of incumbents' average price choices (data from all rounds)

State <i>L</i>	FULL10	ASYM10	FULL14	State <i>H</i>	FULL10	ASYM10	FULL14
ASYM10	0.534	–	–	ASYM10	0.842	–	–
FULL14	0.000	0.101	–	FULL14	0.716	0.504	–
ASYM14	0.017	0.010	0.708	ASYM14	0.494	0.461	0.385

we included in the regression only price choices made in these treatments in state *L* and set *TREATM* equal to 1 if an observation came from Treatment FULL10 and equal to 0 if it came from Treatment ASYM10.¹⁵ The two-tailed *p*-values for the null hypothesis that $\beta = 0$ are presented in Table 3.

Table 3 has two important implications: first, it confirms that there are no significant cross-treatment differences in the incumbents' prices in state *H*. Together with Fig. 2, this implies that the incumbents' prices in state *H* were close to the full information equilibrium price 60 across all treatments. By contrast, Fig. 2 and Table 3 imply that the incumbents' prices in state *L* were significantly higher in Treatment FULL14 than in Treatment FULL10 and in Treatment ASYM14 than in Treatment ASYM10. Hence, holding the information structure fixed, an increase in the entrants' payoff from staying out, $\bar{\pi}$, induced the incumbents to raise their prices in state *L*.

Second, Table 3 shows that the incumbents' prices in Treatments FULL10 and FULL14 were not significantly different than their prices in Treatments ASYM10 and ASYM14, respectively. This implies in turn that holding $\bar{\pi}$ fixed, asymmetric information had no significant effect on the incumbents' behavior relative to their behavior under full information. Once again, this provides strong support for the main implication of the Bagwell and Ramey model of oligopoly limit pricing.

Next, we turn to the entrants' behavior. Table 4 provides a first look at the entrants' behavior by showing the average entry rates in the various treatments.

Table 4 shows that entrants tended to enter in state *L* but stay out in state *H*. The fact that this is true under both full and asymmetric information suggests that, by and large, incumbents' managed to credibly signal the state to the entrants. On the other hand, the table also shows a large fraction of unprofitable entry decisions (entry in state *H* and no entry in state *L*). The frequency of these decisions was smaller in rounds 25–48 than in rounds 1–24, suggesting that similarly to the incumbents, entrants also “learned” to play the full information equilibrium as sessions have progressed. Interestingly, unprofitable entry decisions occurred not only under asymmetric information, but also under full information, albeit less often.¹⁶ As might be expected, unprofitable entry in state *H* was generally more common when $\bar{\pi} = 10$, while unprofitable no entry in state *L* was more common when $\bar{\pi} = 14$. Taken together, the two types of unprofitable entry decisions imply that entry rates were higher when $\bar{\pi} = 10$.

So far, we examined the incumbents' and entrants' behavior in isolation. In order to answer our first three research questions, we now study their joint behavior. Fig. 3 shows the frequency of price pairs played by the incumbents in rounds 25–48 and the corresponding entry rates. The top number in each cell shows the overall entry rate at the relevant price pair (regardless of which state was realized), while the two numbers in parentheses directly below it represent the frequency with which the relevant price pair was played in state *L* (left number) and state *H* (right number). For instance, in Treatment FULL10, the frequency of (40, 40) was 72% in state *L*, 0% in state *H*, and the entry rate following (40, 40) was 97%.

A few interesting observations emerge from Fig. 3. First, the incumbents had a strong tendency to play the full information equilibrium prices, especially when $\bar{\pi} = 10$ and especially in state *H*: (40, 40) was the modal price pair

¹⁵ For the estimations we use the “reoprob” procedure written by Frechette (2001) for the software package STATA. Unfortunately, this procedure cannot be used if a price combination is chosen in one treatment but not in a comparison treatment. We therefore used data from all rounds (rather than focus only on rounds 25–48) in order to minimize the number of price combinations that were never chosen in any given treatment. In Treatment ASYM10 however, the price 30 was never chosen in state *H* and was chosen only once in Treatments FULL10 and FULL14 and only four times in Treatment ASYM14. Hence, when estimating the random effects ordered probit equations in state *H* we grouped together the price choices 30 and 40.

¹⁶ Some unprofitable no-entry decisions in state *L* were in fact intentional: several subjects reported in the post-experimental questionnaire that sometimes they chose to stay out in state *L* because this choice increased the incumbents' payoffs by more than it lowered their own payoff. The post-experimental questionnaire also indicates that some subjects decided to always stay out because this strategy gave them a sufficiently high payoff. Since the payoff from staying out was 14 in treatment FULL14 but only 10 in treatment FULL10, it is not surprising that more entrants adopted this strategy in treatment FULL14.

Table 4
Average entry rates in the various treatments

	Treatment			
	FULL10	ASYM10	FULL14	ASYM14
Rounds 1–48				
State <i>L</i>	93%	90%	75%	62%
State <i>H</i>	5%	18%	4%	11%
Rounds 25–48				
State <i>L</i>	94%	91%	84%	68%
State <i>H</i>	0%	7%	2%	7%

in state *L* and (60, 60) was by far the most commonly played price pair in state *H*. Moreover, the entry rates were very high following (40, 40) (over 94% in Treatments FULL10, ASYM10, and ASYM10, and 68% in Treatment ASYM14), but very low following (60, 60) (below 7% in all treatments). These observations provide additional support for Bagwell and Ramey's (1991) prediction that the oligopoly limit pricing model gives rise to full information equilibrium behavior.

Second, when the incumbents did not play the full information equilibrium prices, their deviations from these prices were small: in state *L*, the second most commonly played price pair was (40, 50), with (40, 60) being third, while in state *H* the second most commonly played price pair was (60, 70), with (60, 50) being third. In all of these cases, only one incumbent deviated from the full information price pair. In fact, the only treatment in which both incumbents deviated consistently from the full information equilibrium prices was Treatment ASYM14, in which the incumbents played (50, 50) and (50, 60) in state *L* about 15% of the time each. In state *H* by contrast, the incumbents never consistently played a price pair that involved deviations by both incumbents from (60, 60). Interestingly, Table 1 shows that the price pairs (40, 50) and (40, 60) are out-of-equilibrium prices in state *L*, while (60, 70) and (60, 50) can be played in a separating equilibrium in state *H*.

Third, the only price choice which is consistent with a pooling equilibrium play is (50, 60); its frequency in Treatment ASYM14 was 15% in state *L* and 12% in state *H*. While this price choice did not deter entry altogether (as it should in a pooling equilibrium), the entry rate following (50, 60) was merely 34%.

Fourth, Fig. 3 shows that the entry rates were high following “low” price pairs (i.e., those in the upper part of each table), but low following “high” price pairs (i.e., those in the lower part of each table). We shall explore this pattern in more detail in the next subsection.

The next three results summarize the discussion so far and provide answers to the first three research questions stated in Section 2.2

Result 1. When the entrant's payoff from staying out was low ($\bar{\pi} = 10$), the subjects appear to have learned to play the full-information separating equilibrium.

Result 2. When the entrant's payoff from staying out was high ($\bar{\pi} = 14$), the subjects tended to play the full-information separating equilibrium, although in state *L* there is a substantial fraction of upward deviations from the full information price of 40. These deviations are inconsistent with either separating or pooling equilibrium choices.

Result 3. (i) Holding the entrant's payoff from staying out fixed, asymmetric information had no significant effect on the incumbents' behavior.

(ii) When the entrant's payoff from staying out was low ($\bar{\pi} = 10$), asymmetric information had no significant effect on the entrants' behavior.

(iii) When the entrant's payoff from staying out was high ($\bar{\pi} = 14$), asymmetric information had no significant effect on the entrants' behavior in state *H* but discouraged entry in state *L*.

We end this section with two more comments about the subjects' behavior in the SP treatments. First, the relatively fast convergence to the full information equilibrium (see Fig. 2) stands in contrast to the typical pattern in experiments on single-sender signaling games in which the equilibrium play emerges only gradually (Camerer, 2003, Chapter 8). For example, in the monopoly limit pricing experiments of Cooper et al. (1997a, 1997b), the convergence

Fig. 3. The frequencies of various price pairs and the corresponding entry rates. The top number in each cell is the relevant entry rate and the two numbers in parenthesis directly below it are the frequencies with which the relevant price pair was played in state L (the left number) and in state H (the right number). There are 192 observations on outcomes in the full information treatments and 288 observations on outcomes in the asymmetric information treatments.

to equilibrium evolves along the following characteristic adjustment process: initially, incumbents ignore the threat of entry and choose their myopic maxima. As a result, entrants are able to infer the cost state and hence stay out when entry is unprofitable. Incumbents realize this and then try to deter entry even when entry is profitable. If the game admits pooling equilibria, play settles into an efficient pooling equilibrium. Otherwise, the pooling attempts are upset by increased entry rates which induce the incumbents to separate when entry is unprofitable. It is plausible that the incumbents in our experiments also played myopically in early rounds before realizing that they might be able to deter entry in state *L* by mimicking the state *H* prices. But since there are two incumbents, mimicking the state *H* prices in state *L* requires the incumbents to coordinate their actions. Given that coordination is hard to achieve under the random matching scheme employed in our experiments, incumbents mostly continued to play their full information strategies throughout.¹⁷

¹⁷ Answers given in the post-experimental questionnaire indicate that some subjects understood that in state *L* incumbents have an incentive to mimic the state *H* prices but this strategy can succeed only if the two incumbents cooperate. For instance, subject 5 in the second session of ASYM10 writes that as an incumbent, s/he “[...] tried to choose 4 [price 60] in which case the B participant [entrant] would be confused about the state but the other A participant [incumbent] did not cooperate.” (Explanation in italics added.)

Second, as mentioned earlier, there was a substantial fraction of upward deviations in state L from the full information equilibrium price 40, especially in treatments with $\bar{\pi} = 14$. These deviations can be interpreted as attempts by the incumbents to collude in state L by playing (50, 50) or (60, 60) instead of playing the Nash equilibrium (40, 40).¹⁸ At least under asymmetric information, these collusive attempts may have the additional effect of signaling to the entrant that the state is H and thereby deter entry.¹⁹

5. Results from the strategy method (SM) treatments

The previous section shows that by and large, asymmetric information had no significant effect on subjects' behavior. Bagwell and Ramey (1991) derive this result by imposing the UPB refinement which eliminates all separating equilibria which involve price distortions. The question then is whether the entrants' behavior in our experiments was consistent with this refinement, and if not, what were the main determinants of the entrants' behavior?

To address this question, we conducted the strategy method (SM) treatments, in which the incumbents were asked to specify a price for each of the two possible states, while the entrants were asked to specify an entry decision for each of the possible 21 distinct incumbents' price pairs. These treatments therefore allow us to observe the entire strategy of the incumbents and the entrants. This feature is particularly useful given that in the SP treatments, we only have a limited number of observations on the entrants' responses following price pairs that were not played often (or not played at all).

Fig. 4 provides a first look at the incumbents' behavior in the SM treatments and shows that it was similar to their behavior in the SP treatments.

Fig. 4. Incumbents' behavior in the strategy method (SM) treatments.

¹⁸ Answers given in the post-experimental questionnaire indicate that some incumbents were indeed trying to collude, especially in state L .

¹⁹ Indeed, as we shall see below, entry rates following the price pairs (50, 50) and (60, 60) were substantially lower than those following (40, 40).

Table 5

The relative frequency of strategy choices by incumbents in the strategy-method treatments

ASYM10-SM: Rounds 1–24								ASYM10-SM: Rounds 24–48							
P_H	30	40	50	60	70	80	Total	P_H	30	40	50	60	70	80	Total
P_L								P_L							
30	0	0	0	0.8	0	0.5	1.3	30	0.5	0	0	0.3	0	0	0.8
40	0	0.8	5.7	47.9	2.1	0	56.5	40	0	0.5	1.8	60.7	3.4	0	66.4
50	0	0.3	4.9	22.9	1	0.8	29.9	50	0	0.3	2.3	21.9	0.3	0.3	25
60	0	0.3	0.5	6.8	1.8	0	8.3	60	0	0.5	0.5	3.6	0.3	0	4.9
70	0	0	0	1	0.5	0	1.6	70	0	0.3	0.3	0	1.6	0.3	2.3
80	0	0	0.8	0.3	0.5	0.8	2.3	80	0	0	0	0	0	0.5	0.5
Total	0	1.3	12	79.7	4.9	2.1	100	Total	0.5	1.6	4.9	86.5	5.6	1.1	100

ASYM14-SM: Rounds 1–24								ASYM14-SM: Rounds 25–48							
P_H	30	40	50	60	70	80	Total	P_H	30	40	50	60	70	80	Total
P_L								P_L							
30	0.5	0	0.3	0.8	0.5	0.3	2.3	30	0	0	0.3	0.5	0	0	0.8
40	0	0	0.5	52.3	2.9	0.3	56	40	0	0	0	62.5	0.3	0.3	63
50	0	0	0.5	21.1	1.3	0	22.9	50	0	0	0.3	18.5	1	0.3	20.1
60	0	0	1	10.7	4.9	0	16.7	60	0	0	1.3	9.1	4.7	0	15.1
70	0	0	0.5	0.3	0.3	0	1	70	0	0	0.3	0	0.5	0	0.8
80	0	0	0	0.5	0	0.5	1	80	0	0	0	0.3	0	0	0.3
Total	0.5	0	2.9	85.7	9.9	1	100	Total	0	0	2.1	90.9	6.5	0.5	100

Note: Each table summarizes a total of 384 incumbent strategy choices.

Table 5 provides a closer look at the incumbents' behavior by presenting the frequencies of the incumbents' strategies in each half of the SM treatments (the entries in each matrix sum up to 100).²⁰ For instance, in rounds 1–24 of Treatment ASYM10-SM, the full information strategy “40 in state L and 60 in stage H ” was chosen in 47.9% of the cases.

Table 5 shows that the incumbents' modal strategy choice in the SM treatments is the full information strategy (40 in state L and 60 in state H), and its frequency was larger in rounds 25–48 than in rounds 1–24. This suggests that as sessions progressed, the incumbents learned to play the full information strategies, with over 60% of the incumbents playing this strategy in rounds 25–48. The second most common strategy was 50 in state L and 60 in state H . The frequency of this strategy varied from 18.5% in Treatment ASYM14-SM to 22.9% in Treatment ASYM10-SM. Interestingly, this strategy is not an equilibrium strategy. The only other frequently played strategy was the pooling equilibrium strategy in which subjects played 60 in both states. Not surprisingly, this strategy was played more frequently in Treatment ASYM14-SM (which admits pooling equilibria) than in Treatment ASYM10-SM (which does not admit pooling equilibria).

Next we turn to the entrants' behavior. Table 6 presents the entry rates associated with each distinct price pair in rounds 25–48 of Treatments ASYM10-SM and ASYM14-SM (we will discuss the reason for writing some entry rates in bold font shortly).²¹ The table shows that by and large, entrants chose to enter when the incumbents' prices were (40, 40), but chose to stay out when the incumbents' prices were (60, 60).

We now ask whether the entrants' response to other price pairs is consistent with the UPB refinement. To this end, recall that the UPB refinement implies that unilateral deviations from (40, 40) in state L and (60, 60) in state H

²⁰ Notice that Table 5 differs from Fig. 3 in that it presents the incumbents' strategies (the prices that each incumbent chose for the two possible states), while Fig. 3 shows the price pairs that two incumbents played at a given state (in the SP treatments the incumbents specified their choices only for the state that was realized but not for the other state).

²¹ We excluded from the data 11 subjects who did not respond to the incumbent's prices: 2 of these subjects always entered (both participated in session 2 of treatment ASYM10-SM) while 9 always stayed out (one participated in session 2 of treatment ASYM10-SM, three participated in session 1 of treatment ASYM14-SM, and additional five participated in session 2 of treatment ASYM14-SM). Including these 11 subjects in the analysis would have only added noise and made it harder to understand the behavior of entrants who did respond to the incumbents' prices.

Table 6

Average entry rates in rounds 25–48 of treatments ASYM10-SM and ASYM14-SM (some data excluded)

ASYM10-SM: Session 1							ASYM10-SM: Session 2						
Price	30	40	50	60	70	80	Price	30	40	50	60	70	80
30	85	85	85	85	85	77	30	65	65	56	67	56	67
40	–	94	94	85	69	60	40	–	50	71	68	56	60
50	–	–	77	34	10	27	50	–	–	50	46	32	35
60	–	–	–	11	12	10	60	–	–	–	26	40	39
70	–	–	–	–	19	19	70	–	–	–	–	19	32
80	–	–	–	–	–	10	80	–	–	–	–	–	33

ASYM14-SM: Session 1							ASYM14-SM: Session 2						
Price	30	40	50	60	70	80	Price	30	40	50	60	70	80
30	89	89	75	44	56	65	30	96	86	88	84	79	82
40	–	99	86	65	56	67	40	–	91	82	82	71	77
50	–	–	44	43	22	22	50	–	–	46	29	29	34
60	–	–	–	0	11	21	60	–	–	–	12	11	20
70	–	–	–	–	11	22	70	–	–	–	–	18	18
80	–	–	–	–	–	33	80	–	–	–	–	–	18

should not affect the entrant's beliefs and hence his entry decision. Consequently, the entry rates following the price pairs $(x, 40)$ and $(x, 60)$, where $x \in \{30, 50, 70, 80\}$, should be the same as those following $(40, 40)$ and $(60, 60)$, respectively.²² (The refinement does not pin down the entrant's belief following $(40, 60)$ since this pair is one deviation away from both $(40, 40)$ and $(60, 60)$.) To test this hypothesis, we run the following random effects probit regressions:

$$\text{Prob}[Entry_{it}(p_1, p_2)] = F(\gamma_{40}D_{40,40} + v_i + \varepsilon_{it}),$$

and

$$\text{Prob}[Entry_{it}(p_1, p_2)] = F(\gamma_{60}D_{60,60} + v_i + \varepsilon_{it}),$$

where $Entry_{it}(p_1, p_2)$ is the entry decision of subject i in round t at a given price pair (p_1, p_2) , $D_{40,40}$ and $D_{60,60}$, respectively, are dummy variables which are equal to 1 if the entry decision was made following the price pair $(40, 40)$ and $(60, 60)$ and are equal to 0 otherwise, v_i is a fixed subject-specific component, and ε_{it} is a subject-specific error term that may vary across observations from the same subject. For example, to test whether the entry rates following the price pairs $(40, 40)$ and $(x, 40)$ are significantly different, we included in the regression only the entry rates following these two price pairs. Hence, the variables γ_{40} and γ_{60} measure the effect that unilateral deviations from $(40, 40)$ and $(60, 60)$ had on the entry rates. In each regression, we only used data from rounds 25–48 (in which the subjects are already experienced) and excluded subjects who did not react to the incumbents' prices.

The results (which we do not report in detail) show that with few exceptions, the entry rates in Table 6 following $(x, 40)$ and $(x, 60)$, where $x \in \{30, 50, 70, 80\}$, are significantly different than those following $(40, 40)$ and $(60, 60)$, respectively.²³ This implies that the entrants' behavior was inconsistent with the UPB refinement: by and large, unilateral deviations from $(40, 40)$ discouraged entry, while unilateral deviations from $(60, 60)$ encouraged entry.

While the entrants' behavior was inconsistent with the UPB refinement, the entry rates in Table 6 suggest that entrants' behavior was nonetheless consistent with the general idea behind the UPB refinement, in the sense that "small" deviations from equilibrium had a "small" effect on the entrants' behavior. In particular, the entrants tended to enter following price pairs that were "close" to $(40, 40)$ but tended to stay out following prices that were close to $(60, 60)$.

To examine this "bi-polar" pattern in more detail, we perform a hierarchical agglomerative cluster analysis of the 21 average entry rates in each of the four SM sessions. In general, a hierarchical cluster analysis is a statistical method for

²² It should be noted however that since there were unprofitable entry decisions even under full information (failure to enter in state L and entry in state H), the correspondence between the entrants' behavior and their beliefs is only imperfect.

²³ We were unable to check for the significance of unilateral deviations from $(60, 60)$ in session 1 of treatment ASYM14-SM, using probit regressions because there was no entry at all following $(60, 60)$.

identifying relatively homogeneous clusters of observations based on their characteristics. An agglomerative analysis starts with each observation as a separate cluster and then merges the two closest clusters into a single cluster. This process is repeated sequentially, thereby reducing the number of clusters at each step until only one cluster is left. To determine the distance between every two possible clusters, we use Ward's method in which the successive clustering steps are chosen to minimize at each step the variance within clusters (see Kaufman and Rousseeuw, 1990).²⁴ In general, Ward's method is regarded as very efficient, but tends to create clusters of small size.

The cluster analysis reveals that each of the four SM sessions featured two very distinct clusters: one cluster with “high” entry rates which are marked in Table 6 by bold font (cluster *HIGH*) and a second cluster with “low” entry rates (cluster *LOW*). This “bi-polar” entry behavior was more pronounced in Treatment ASYM10-SM which does not admit pooling equilibria. Mann-Whitney U tests reveal that the difference between the average entry rates across the two clusters is highly significant ($p < 0.001$) in all four sessions. It should also be noted that price pairs in which the minimal price was either 30 or 40 were always part of cluster *HIGH*, while, apart from (80, 80), price pairs in which both prices were at least 50 and their sum was at least 120 were always part of cluster *LOW*.

Having identified two clusters in the entrants' data, we now analyze the entrants' behavior within each cluster. To this end, we run random-effects probit regressions of the form:

$$\text{Prob}[Entry_{it}(p_1, p_2)] = F(\alpha_0 + \alpha_1 DIS40 + \alpha_2 DIS60 + \alpha_3 Min + \alpha_4 Max + \alpha_5 CLUSTER + \varepsilon_{it}),$$

where $DIS40 \equiv \sqrt{(p_1 - 40)^2 + (p_2 - 40)^2}$ and $DIS60 \equiv \sqrt{(p_1 - 60)^2 + (p_2 - 60)^2}$ are the Euclidean distances of (p_1, p_2) from (40, 40) and (60, 60); $Min \equiv \min\{p_1, p_2\}$ and $Max \equiv \max\{p_1, p_2\}$ are the minimal and maximal components of (p_1, p_2) ; $CLUSTER$ is a dummy variable which is equal to 1 if (p_1, p_2) belongs to cluster *HIGH* (i.e., the bold-faced entry rates in Table 6), and is equal to 0 if (p_1, p_2) belongs to cluster *LOW*; and $Entry_{it}(p_1, p_2)$ and ε_{it} are defined as above. As before, we restrict attention to rounds 25–48 and exclude from the data subjects who did not respond to the incumbent's prices. The regression results are shown in Tables 7 and 8. Regression 2 in each table includes the $CLUSTER$ dummy while regression 1 does not include it. The coefficients in the tables are the marginal effects of the relevant explanatory variables on the probability of entry, evaluated at the mean of the respective explanatory variable. In other words, each coefficient in Tables 7 and 8 represents the partial derivative of $\text{Prob}[Entry_{it}(p_1, p_2)]$ with respect to the relevant explanatory variable, when the derivative is evaluated at the mean value of the variable. The marginal effect of the dummy variable $CLUSTER$ is the change in $\text{Prob}[Entry_{it}(p_1, p_2)]$ due to a switch from cluster *LOW* to cluster *HIGH*.

Table 7
Results of entry rates regressions for treatment ASYM14-SM (marginal effects)

	Asym10-SM, session 1		Asym10-SM session 2	
	Regression 1	Regression 2	Regression 1	Regression 2
$\alpha_1(DIS40)$	−0.0383*** (0.0047)	−0.0228*** (0.0048)	−0.0053 (0.0040)	0.0003 (0.0042)
$\alpha_2(DIS60)$	0.0259*** (0.0028)	0.0132*** (0.0030)	0.0050** (0.0025)	0.0024 (0.0026)
$\alpha_3(Min)$	−0.0003 (0.0024)	0.0033 (0.0024)	−0.0086*** (0.0023)	−0.0056** (0.0024)
$\alpha_3(Max)$	0.0186*** (0.0035)	0.0098*** (0.0035)	0.0033 (0.0032)	0.0004 (0.0033)
$\alpha_5(CLUSTER)$		−0.4639*** (0.0484)		−0.2240*** (0.0580)

Note: Standard errors in parentheses.

** $p < 0.05$,

*** $p < 0.01$.

²⁴ More precisely, the distance between every two clusters, X and Y , is given by $D(X, Y) = ESS(XY) - ESS(X) - ESS(Y)$, where XY is the union of X and Y , and $ESS(\cdot)$ is the error sum of squares. Given a cluster $Z = \{z_1, \dots, z_N\}$ with N observations, $ESS(Z) = \sum_i |z_i - \frac{\sum_i z_i}{N}|^2$, where $|\cdot|$ is Euclidean distance between observation z_i and the mean observation in cluster Z .

Table 8
Results of entry rates regressions for treatment ASYM14-SM (marginal effects)

	Asym14-SM, session 1		Asym14-SM session 2	
	Regression 1	Regression 2	Regression 1	Regression 2
$\alpha_1(DIS40)$	−0.0646*** (0.0069)	−0.0522*** (0.0067)	−0.0162*** (0.0055)	−0.0097* (0.0056)
$\alpha_2(DIS60)$	0.0564*** (0.0048)	0.0546*** (0.0045)	0.0218*** (0.0032)	0.0130*** (0.0036)
$\alpha_3(Min)$	0.0189*** (0.0034)	0.0284*** (0.0036)	−0.0089*** (0.0028)	−0.0038 (0.0030)
$\alpha_3(Max)$	0.0358*** (0.0050)	0.0301*** (0.0050)	0.0069 (0.0045)	0.0023 (0.0045)
$\alpha_5(CLUSTER)$		−0.5754*** (0.0636)		−0.3455*** (0.0704)

Note: Standard errors in parentheses.

* $p < 0.1$,

*** $p < 0.01$.

Tables 7 and 8 indicate that the entrants' behavior in the strategy method (SM) treatments exhibits a “bi-polar” pattern: relatively high entry rates following “low” incumbent price and relatively low entry rates following “high” incumbent prices. In particular, the coefficient of the *CLUSTER* dummy is highly significant even after controlling for other possible determinants of entry rates. The tables show that a switch from cluster *LOW* to cluster *HIGH* leads to a large drop in the probability of entry which ranges from 22.4 percentage points in session 2 of Treatment ASYM10-SM, to 57.5 percentage points in session 1 of Treatment ASYM14-SM.

Moreover, the tables show that with the exception of session 2 in Treatment ASYM10-SM, entry rates increased with the Euclidean distance from the price pair (60, 60) but decreased with the Euclidean distance from the price pair (40, 40), and were also affected by the minimum and the maximum price chosen by the incumbents. It should be noted that in general, the entrants' behavior was affected by the combination of incumbents' prices rather than by only one of these prices (say the minimum or the maximum price). The fact that *both* prices matter means that there is an important difference between single sender signaling games in which the receiver observes a single signal and multi-sender signaling games in which the receiver observes more than one signal about the same unknown parameter.

We summarize the main findings on the entrants' behavior as follows:

Result 4. The entrants' behavior is inconsistent with the formal definition of UPB refinement, but is nonetheless consistent with the general idea behind the UPB refinement, which is that “small” deviations from equilibrium prices should have “small” effects on the receivers' (entrants') behavior. In particular, it appears that the entrants followed a “bi-polar” decision rule and tended to enter following “low” price pairs but tend to stay out following “high” price pairs.

Finally, given that the entrants' behavior is inconsistent with the UPB refinement, one may wonder how come we nonetheless get such strong support for the full information equilibrium? To examine this issue, we have computed the incumbents' empirical expected payoffs for each of the four SM sessions, given the actual average entry rates in rounds 25–48 of that session. For instance, recalling that the payoff of each incumbent in state *L* following (30, 30) is 42 if there is entry and 53 if there is no entry, and noting that the average entry rate in rounds 25–48 of session 1 of ASYM10-SM following (30, 30) was 85%, the expected payoff of each incumbent in state *L* of that session following (30, 30) was $0.85 \times 42 + 0.15 \times 53 = 43.6$. The incumbents' empirical expected payoffs in session 1 of ASYM10-SM are presented in Table 9.

Using these payoff matrices, it is easy to check that the unique Nash equilibrium in the reduced-form game between the two incumbents is (40, 40) in state *L* and (60, 60) in state *H*. Similar computations shows that the unique Nash equilibrium in the incumbents' reduced game coincides with the full-information equilibrium in all SM sessions. This fact might explain the incumbents' tendency to play full information equilibrium under asymmetric information.

Table 9

Incumbents' expected payoff matrices given entrants' actual behavior in rounds 25–48 in session 1 of treatment ASYM10-SM

ASYM10-SM, Session 1, state <i>L</i>							ASYM10-SM, Session 1, state <i>H</i>						
Price	30	40	50	60	70	80	Price	30	40	50	60	70	80
30	43.6	46.7	49.8	53.7	56.7	60.8	30	−2.4	−3.4	−5.3	−7.3	−8.4	−9.9
40	44.8	48.7	53.7	59.8	66.7	72.8	40	8.8	8.3	8.3	8.8	9.6	10.0
50	39.8	45.7	52.5	65.9	74.9	80.0	50	13.8	15.3	17.2	21.3	24.5	24.7
60	28.7	36.7	50.9	61.8	69.7	78.8	60	11.8	15.8	21.3	25.5	29.4	32.5
70	19.8	21.7	38.8	48.6	57.7	67.7	70	8.8	9.6	17.5	22.4	27.1	32.1
80	20.8	22.8	26.8	28.8	38.9	51.9	80	9.2	10.0	11.7	12.5	19.1	25.5

Note: Only the row player's payoffs are shown.

6. Conclusion

We have examined the strategic behavior of senders and receivers in the context of oligopoly limit pricing experiments in which two incumbents are privately informed about whether entry is profitable or not. Using the intuition from monopoly limit pricing games, one might expect that the incumbents would have to distort their prices away from their full information levels in order to credibly signal that entry is unprofitable. However, as Bagwell and Ramey (1991) first showed, this need not be the case when the entrant can observe the prices of both incumbents. Intuitively, starting from the full information prices, incumbents might wish to fool the entrant when entry is profitable into believing that it is unprofitable. However, this requires both incumbents to mimic the prices that would have been chosen had entry been unprofitable. But since the incumbents cannot coordinate their prices, the entrant will not be fooled and hence entry will take place only when it is profitable. As a result, the incumbents will simply play their full information strategies.

While the full information equilibrium is very appealing and while Bagwell and Ramey (1991) prove that it is the only separating equilibrium that survives when the entrants' beliefs are unprejudiced, the oligopoly limit pricing model also admits additional equilibria which involve price distortions. The question then is whether the full information equilibrium emerges and whether unprofitable entry is credibly deterred without price distortions by the incumbents. The latter question is important because separating equilibria which involve upward price distortions are anticompetitive in the short run and may fail to lead to actual entry in the long run. It is therefore natural to wonder if these anticompetitive equilibria emerge or not.

Our experimental results provide strong support for the full information equilibrium: incumbents' prices quickly converge to the full information equilibrium levels both under full and under asymmetric information and entrants learn to correctly interpret the incumbents' prices and (tend to) enter when it is profitable to do so but stay out otherwise. This behavior is particularly pronounced in treatments that do not admit pooling equilibria. Our results also show that the entrants seem to follow a "bi-polar" decision rule and tend to enter following "low" price pairs but tend to stay out following "high" price pairs. Moreover, their entry rates fall (increase) as prices get further away from the full information prices when entry is profitable (unprofitable).

Our results leave several interesting questions about multi sender games unanswered. First, in our experiments the incumbents' interests were perfectly aligned as both incumbents were interested in deterring entry. However, one can think about real-life situations in which multiple informed players may wish to signal conflicting messages to the same uninformed player. For example, if the entrant's product is a substitute for the product of one incumbent but a complement for the product of the other incumbent, then the first incumbent will be interested in deterring entry while the second incumbent will be interested in promoting it (see Schultz, 1999).²⁵ The question then is what kind of equilibrium emerges when the senders "compete" with each other on the receiver's beliefs. In particular, out of equilibrium actions by the senders may have a new meaning in this context because the receiver may interpret a deviation from a putative equilibrium as an attempt to discredit the signal sent by the rival sender.

²⁵ Other scenarios in which the incumbents will have conflicting interests include the case where one incumbent operates in the upstream market while the other operates in the downstream market and the entrant considers entry into one of these markets, and the case where each incumbent operates in a separate market and wishes to deter entry into his own market (but does not care if there is entry into the other market). Note that in the latter case, the incumbents' costs could be positively correlated, negatively correlated, or uncorrelated.

Second, the two incumbents in our experiments shared the same information. This setting is a natural starting point for experimental studies of oligopoly limit pricing games because it differs from the monopoly limit price game only in the number of senders but not in the information structure. However, in some applications, (e.g., Martin, 1995), the senders' information may only be imperfectly correlated. If the incumbents' costs are imperfectly correlated, then in a separating equilibrium, any price combination may emerge with positive probability on the equilibrium path. This suggests in turn that a price distortion would be needed to support a separating equilibrium.²⁶

Third, a crucial feature of our experimental setting is that the entrant can observe the individual price of each incumbent. Without this feature, the incumbents would not face a coordination of the kind we explored in our experiments. The obvious question is whether absent this coordination problem, e.g., when the entrant observes a signal only about the *aggregate* behavior of the incumbents as in Harrington (1987), the oligopoly limit pricing game would involve price distortions or not.

In light of these open questions, we view our results as a first attempt to examine the behavior of individuals in the context of multi-sender signaling games. We believe that further experimental studies are needed in order to enhance our understanding of how individuals behave in such situations.

Acknowledgments

We thank Kyle Bagwell, David Cooper, Dirk Engelmann, Nikos Georgantzis, Hans Theo Normann, Dean Williamson, two anonymous referees, an associate editor, and seminar participants at Ben Gurion University, Brunel University, Helsinki Center of Economic Research, The 2006 International Industrial Organization Conference in Boston, Séminaire Parisien d'Economie Industrielle, Tel Aviv University, University College London, University Paris 1 Pantheon-Sorbonne, and the Wissenschaftszentrum Berlin (WZB) for helpful comments. The first author acknowledges financial help from the Netherlands Organisation for Scientific Research (NWO) through a VIDI grant.

Appendix A. Instructions (treatment FULL10/ASYM10)

Please read these instructions closely!

Please do not talk to your neighbors and remain quiet during the entire experiment.

If you have a question, please raise your hand. We will come up to you to answer it.

A.1. Introduction

In this experiment you can earn money by interacting with other participants.

Your earnings will be measured in "Points." The number of points that you will earn will depend on the decisions that you and the other participants make.

At the beginning of the experiment, every participant will receive 100 Points as an initial endowment.

Your total number of points at the end of the experiment will be equal to the sum of the points you have earned in each round plus your initial endowment.

For every 70 Points you will be paid 1 Euro in cash.

Your identity will remain anonymous to us as well as to the other participants.

A.2. Description of the experiment

The experiment will consist of 48 rounds. The events in each round are as follows:

At the beginning of each round, you will be first randomly assigned into one of 4 groups of 3 participants each. In each group, two participants will be assigned to act in role A and one participant will be assigned to act in role B.

After the 4 groups were formed, a random draw for each group will determine the "state": with probability $\frac{1}{2}$, the state will be "state 1" and with probability $\frac{1}{2}$ it will be "state 2."

²⁶ Martin (1995) shows that if the incumbents do not know each other's costs, then apart from entry deterrence effects there are also cost-type revelation effects as each incumbent may also try to affect the beliefs of the rival incumbent about its own costs.

The realized state (state 1 or state 2) will be announced to the two A-participants in each group. Then, each A-participant will be asked to choose one of the numbers 1, 2, 3, 4, 5, or 6.

[The next sentence only in Treatment ASYM10] After the two A-participants have chosen their numbers, the B-participant will be informed about the two numbers chosen by the two A-participants in his/her group, but not about which state was realized.

[The next sentence only in Treatment FULL10] After the two A-participants have chosen their numbers, the B-participant will be informed both about the state (whether it is state 1 or state 2) and about the two numbers chosen by the two A-participants in his/her group.

Then, the B-participant will have to choose between option “X” and option “Y.”

The payoff in each round is described below.

A.3. Payoffs

The payoff of each A-participant in each round will depend on the state, on the numbers chosen by the two A-participants, and on the option chosen by the B-participant.

The payoff of each B-participant in each round will depend on his/her own decision and on the state, but not on the numbers chosen by the two A-participants in his/her group.

All necessary information about the precise payoffs is included in the 5 tables that appear at the end of these instructions. [See Fig. 2.]

All A-participants have the same payoff tables, and all B-participants have the same payoff tables.

A.4. The payoffs of the A-participants

The 4 tables that specify the payoffs of the A-participants result from the realized state and the decision of the B-participant with whom the two A participants are grouped in the relevant round.

Each of the 4 tables corresponds to one possible combination of the realized state and the option chosen by the B-participant. One table corresponds to the case where the realized state is 1 and participant B chose option X, the second table corresponds to the case where the realized state is 1 and participant B chose option Y, the third table corresponds to the case where the realized state is 2 and participant B chose option X, and the fourth table corresponds to the case where the realized state is 2 and participant B chose option Y.

The rows in the A-participants’ payoff tables correspond to the participant’s own chosen number (each row corresponds to one of the 6 possible choices that the participant can make) and the columns correspond to the choice made by the other A-participant (each column corresponds to one of the 6 possible choices that the other A-participant can make). The upper left number in each cell (in blue) corresponds to the participants’ own payoff and the bottom right number in each cell (in red) corresponds to the other A-participant’s payoff.

Note that since the B-participant in each group chooses between options X and Y after the two A-participants in their group have chosen their numbers, the A-participants will not know, when making their choice, which payoff table is relevant for them. (Given the announced state there could be two relevant payoff tables depending on the B-participant’s choice.)

A.5. The payoffs of the B-participants

The rows in participant B’s payoff table correspond to the option that the B-participant chooses (X or Y) and the columns correspond to the realized states (state 1 or state 2). Therefore, each of the 4 cells specifies the B-participant’s payoff for a specific combination of the B-participant’s own choice and the realized state.

A.6. Role assignment and information during the experiment

The experiment will consist of 48 rounds.

Your role in the experiment will alternate between role A and role B. The roles are fixed for 8 consecutive rounds. After 8 rounds, new roles are assigned to all participants and these roles remain fixed for another 8 rounds.

Each participant will act exactly 32 rounds in role A and exactly 16 rounds in role B.

Your computer screen (see the top line) indicates in every round which role you have in that round.

Please remember that in every round, 4 groups of 3 participants are randomly selected from the pool of all participants in the room. We will make sure that each of the 4 groups will always consist of two A-participants and one B-participant.

At the end of each round, you will be given the following information about what happened in your own group during the round: what was the realized state, what were the numbers chosen by the two A-participants, what was the option chosen by the B-participant, and what was your own payoff.

References

- Bagwell, K., Ramey, G., 1991. Oligopoly limit pricing. *RAND J. Econ.* 22, 155–172.
- Camerer, C., 2003. *Behavioral Game Theory: Experiments on Strategic Interaction*. Princeton Univ. Press.
- Cooper, D.J., Gravin, S., Kagel, J.H., 1997a. Signaling and adaptive learning in an entry limit pricing game. *RAND J. Econ.* 28, 662–683.
- Cooper, D.J., Gravin, S., Kagel, J.H., 1997b. Adaptive learning vs. equilibrium refinements in an entry limit pricing game. *Econ. J.* 107, 553–575.
- de Bijl, P., 1997. Entry deterrence and signaling in markets for search goods. *Int. J. Ind. Organ.* 16, 1–19.
- Fischbacher, U., 2007. Zurich toolbox for readymade economic experiments. *Exper. Econ.* 10, 171–178.
- Fluet, C., Garella, P., 2002. Advertising and prices as signals of quality in a regime of price rivalry. *Int. J. Ind. Organ.* 20, 907–930.
- Frechette, G.R., 2001. Random-effects ordered probit. *Stata Technical Bull.* 59, 23–27.
- Harrington, J., 1986. Limit pricing when the potential entrant is uncertain of its cost function. *Econometrica* 54, 429–437.
- Harrington, J., 1987. Oligopolistic Entry Deterrence Under Incomplete Information. *RAND J. Econ.* 18, 211–231.
- Hertzenndorf, M., Overgaard, P., 2002. Prices as signals of quality in duopoly: Type-dependent unit costs. Preliminary and incomplete draft. Updated version can be downloaded from <http://www.econ.au.dk/vip.htm/povergaard/pbohome/pbohome.html>.
- Kaufman, L., Rousseeuw, P.J., 1990. *Finding Groups in Data: An Introduction to Cluster Analysis*. Wiley, New York.
- Martin, S., 1995. Oligopoly limit pricing: Strategic substitutes, strategic complements. *Int. J. Ind. Organ.* 13, 41–65.
- Martinelli, C., Matsui, A., 2002. Policy reversals: Electoral competition with privately informed parties. *J. Public Econ. Theory* 4, 39–61.
- Matthews, S., Fertig, D., 1990. Advertising signals of product quality. Discussion paper No. 881. CMSEMS, Northwestern University.
- Milgrom, P., Roberts, J., 1982. Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica* 50, 443–459.
- Schultz, C., 1996. Polarization and inefficient policies. *Rev. Econ. Stud.* 63, 331–344.
- Schultz, C., 1999. Limit pricing when incumbents have conflicting interests. *Int. J. Ind. Organ.* 17, 801–825.
- Yehezkel, Y., 2007. On the robustness of the full-information separating equilibrium in multi-sender games. Mimeo. Available at <http://www.tau.ac.il/~yehezkel/Multi-sender.pdf>.