Transparency in broadband markets: An experimental study of how information disclosure affects quality*

Bastian Henze and Florian Schuett†
TILEC & CentER, Tilburg University

September 2010

PRELIMINARY – DO NOT CITE OR CIRCULATE

Abstract

We propose an experimental design to investigate the role of transparency in a duopoly market with vertical differentiation. Firms choose both the quality and the price of their product, while consumers differing in their taste for quality choose from which firm to buy. We compare three different treatments in which we vary the share of consumers who are informed about quality. Specifically, we have a full-information treatment in which all consumers are informed, a no-information treatment in which none of them are informed, and a subset treatment in which half of them is informed and the other half uninformed. We find that, contrary to theoretical predictions, firms do not differentiate quality under full information. Rather, both tend to offer services of similar, and high, quality, entailing more intense price competition than predicted by theory. Under no information, we observe a "lemons" outcome where quality is low. At the same time, firms manage to maintain prices substantially above marginal cost. In the subset treatment, quality is close to the full-information level. This suggests that a signaling equilibrium tends to materialize when some consumers are informed and others are not. Our results have implications for consumer protection policies aimed at increasing transparency, such as recent EU legislation mandating that ISPs disclose information on their network management practices.

Keywords: Vertical differentiation, experimental economics, quality signaling, informational externalities

* Part of this research project was funded by the Dutch Ministry of Economic Affairs (Ministerie van Economische Zaken), whose support we gratefully acknowledge. We thank Eric van Damme, Pierre Larouche, Wieland Müller, Jan Potters, Bert Willems, and particularly Jasper Sluijs, as well as seminar participants at TILEC and CERRE for helpful comments and suggestions. The usual disclaimer applies.

† Corresponding author. Postal address: Tilburg University, Department of Economics, PO Box 90153, 5000 LE Tilburg, Netherlands. Email: `f.schuett@uvt.nl`.

1 Introduction

Consumer protection policies often involve transparency requirements. Typically, these requirements mandate that sellers disclose certain information that is supposed to help consumers assess the quality of the products on offer. When consumers cannot observe quality, economic theory predicts that firms will supply inefficiently low quality (this is a variant of the famous lemons problem (Akerlof, 1970)). At the same time, theory does not make sharp predictions as to the desirability of mandatory disclosure. When quality is observable, firms are predicted to engage in vertical differentiation in order to relax price competition (Shaked and Sutton, 1982). While average quality would increase, so would prices; the overall effect on consumer surplus is ambiguous. Moreover, because according to the theory producer surplus is higher when quality is observable, firms should have an incentive to voluntarily disclose the information themselves. Transparency requirements should not be necessary.

In many markets producers fail to disclose such information, however, leaving consumers uncertain about the level of quality they can expect. An example is the market for broadband internet access, where producers typically advertise only the maximum bandwidth, rather than other, more informative measures of quality.[1] As part of the updated European Regulatory Framework for Electronic Communications,[2] European lawmakers have included transparency requirements for Internet Service Providers to mitigate such concerns. The idea behind this transparency policy is that ISPs are required to be transparent on how actively they manage their network, so that consumers can come to a more informed purchasing decision when deciding on a broadband Internet subscription.[3]

In this paper, we assess the effect of transparency by means of a laboratory experiment. Our experimental market consists of two sellers offering a product to four buyers. Sellers simultaneously choose a level of quality for their product. Then, they observe each other's quality and simultaneously post a price. Finally, buyers decide whether and from which seller to buy. We study four treatments that differ in the amount of information that buyers possess

[1] In this context, quality refers to the actual performance of a broadband connection beyond advertised maximum down- and upload speed, including information on parameters such as latency, capacity, peak hours, network management and overall reliability, as well as measures of average and minimum (guaranteed) bandwidth. Even though consumers may learn about the quality of their connection by experience, switching costs prevent them from sampling a wide range of providers before making their definite choice.

[2] See Directive 2009/140/EC of the European Parliament and of the Council of 25 November 2009 amending Directives 2002/21/EC on a common regulatory framework for electronic communications networks and services, 2002/19/EC on access to, and interconnection of, electronic communications networks and services, and 2002/20/EC on the authorization of electronic communications networks and services, OJ 2009 L337/37; and Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009 amending Directive 2002/22/EC on universal service and users' rights relating to electronic communications networks, Directive 2002/58/EC concerning the processing of personal data and the protection of privacy in the electronic communications sector and Regulation (EC) No 2006/2004 on consumer protection cooperation, OJ 2009 L337/11.

[3] This legislation is part of a broader debate about the issue of network neutrality. For a survey of the economic literature, see Schuett (2010).

about quality. In the first treatment, none of the buyers observe the quality on offer (we call this treatment "no info"). In the second treatment, all of the buyers perfectly observe quality ("full info"). In the third treatment, half of the buyers perfectly observes quality, while the other half does not ("subset"). In the fourth treatment, all of the buyers observe an imperfect signal about quality ("signal").

The evidence from our experiment suggests that transparency is a more effective tool to raise welfare and consumer surplus than theory would lead one to expect. This insight is based on three main findings. First, in the full-info treatment competition is fiercer than predicted. Producers fail to vertically differentiate their products; instead, both tend to offer the highest level of quality and compete vigorously on price. Second, while we do observe a lemons outcome in the no-info treatment, sellers manage to sustain prices substantially above marginal cost, allowing them to "rip off" buyers particularly in the early rounds of the experiment. Together, these two findings make transparency potentially more valuable to buyers and less valuable to sellers than the theory predicts. Third, we find that these effects of transparency survive when we make more realistic assumptions about how information disclosure translates into buyers' ability to evaluate quality. Not all buyers may have access to or understand the information disclosed. In our intermediate treatments, where buyers are less than perfectly informed, offered quality is significantly greater than under no information, and often close to the full-information level. There is some evidence that this result can be attributed to sellers signaling quality through prices. Our results imply that the case for transparency requirements may be stronger than previously thought. They may also explain why in many real-world markets, such as broadband, we do not observe voluntary disclosure.

The remainder of this paper is organized as follows. In Section 2 we describe the experimental design and outline theoretical predictions and procedures. In Section 3, we present and discuss our results. Section 4 concludes.

2 The experiment

2.1 The environment

The environment we study is a finitely repeated version of the following three-stage game. There are two firms, indexed by $j \in \{A, B\}$, and n consumers, indexed by $i \in \{1, \dots, n\}$. At stage 1, each firm j chooses the quality q_j of its product from the set $Q = \{1, 2, \dots, 10\}$. At stage 2, firms observe each other's quality and post a price p_j. At stage 3, consumers observe prices as well as varying amounts of information about quality depending on the treatment and make a purchasing decision.

Consumer i has a taste for quality θ_i drawn at the beginning of each period from a uniform distribution on the support $[\underline{\theta}, \overline{\theta}]$. A consumer's taste for quality is her private information; firms merely know the distribution from which it is drawn. Consumers buy at most one unit and have utility

$$u_i = \begin{cases} v + \theta_i q_j - p_j & \text{if } i \text{ buys from firm } j \\ 0 & \text{if } i \text{ doesn't buy} \end{cases}$$

That is, utility includes a fixed component v that is independent of quality. Utility increases linearly with quality according to the consumer's marginal valuation for quality given by θ_i, and decreases with price. Firm j's per-unit cost of providing quality q_j is cq_j. That is, high quality is more costly to provide than low quality, and the higher cost is a variable cost incurred for every unit that is sold. There is no fixed cost.

In the experiment, subjects repeat the above game for a total of 32 periods, including two trial periods which do not count towards their earnings. (In what follows, we will refer to the duration of the experiment as 30 periods.) Subjects are randomly assigned to be either a seller or a buyer. In each session, there are three groups of sellers and three groups of buyers. The groups remain the same for the entire experiment. At the beginning of every period, each pair of sellers is randomly matched to a group of buyers. None of the subjects know to which group they are matched.

There are four different treatments in which buyers have varying degrees of information on the quality of the product supplied by the sellers, and which are otherwise identical: *no info*, *full info*, *subset*, and *signal*. In the *no info* treatment, buyers only observe prices and have no information on the quality supplied. In the *full info* treatment, buyers perfectly observe the quality supplied by each seller. In the *subset* treatment, half the buyers perfectly observe quality while the other half do not. Whether a buyer is informed or uninformed is independent of her taste for quality θ. In the *signal* treatment, all buyers observe an imperfect signal s_j about the quality of seller j's product, $j = A, B$. The signal satisfies $s_j = q_j + \varepsilon_j$ where ε_j is drawn from a uniform discrete distribution on $\{-r, \ldots, 0, \ldots, r\}$, for some positive integer r, with truncation so as to avoid values of s_j outside the range of possible quality levels Q. The signals are the same for all buyers in a group and unobservable to the sellers.

The parameters we used for the experiment are $c = 1$, $\underline{\theta} = 1$, $\overline{\theta} = 4$, $v = 20$, $n = 4$, and $r = 5$. These parameter values imply that the efficient outcome is for all consumers to obtain the highest possible quality because the marginal cost of quality (c) is (weakly) lower than the marginal valuation for quality of the lowest type ($\underline{\theta}$). Note that prices affect efficiency only to the extent that they discourage some consumer from buying. As long as all consumers buy, prices have no effect on efficiency; they merely affect the distribution of surplus. The parameters $\underline{\theta}$ and $\overline{\theta}$ were chosen so as to obtain integer values for the equilibrium prices in the full-information treatment. The parameter v was chosen so as to obtain a relatively even distribution of surplus between buyers and sellers in the equilibrium of the full-information treatment; it also ensures that the market is covered.

We calibrated the range of the error term, r, in such a way that a consumer in the *signal* treatment who tries to determine which seller offers higher quality solely based on the signals he receives would make the same amount of mistakes as an average consumer in the *subset* treatment, given the distribution of qualities offered that we observe in the other three treatments.[4] While a fully informed consumer perfectly observes the quality, uninformed and imperfectly informed consumers make mistakes. An uninformed consumer has a 50% chance of correctly identifying the higher-quality product. Thus, an average consumer in the *subset* treatment (made up of one half of informed and one half of uninformed consumers) has a 75% chance of correctly identifying the higher-quality product. We weight the mistakes this average consumer is bound to make by their magnitude to account for the fact that mistakenly choosing a product of quality $q = 1$ instead of one of quality $q = 10$ is worse than mistakenly choosing one of quality $q = 5$ instead of one of quality $q = 6$. Based on the distribution of qualities offered, we can compute a measure of the total weighted mistakes of a hypothetical average consumer in the *subset* and *signal* treatments. We calibrate the range of the error term in such a way that the total weighted mistakes in both treatments are about the same. The idea of this calibration is to make both treatments roughly comparable in terms of the information held by an average consumer.

2.2 Theoretical predictions: equilibrium in the stage game

In this section, we analyze the one-shot version of the game which is repeated for 30 periods in the experiment, with each group of sellers having a probability of one third of being matched with each of the buyer groups in each period. Because players do not know with which group they are matched and the probability of being matched with a particular group is only one third, we do not expect the dynamic aspect to be of much relevance for the interaction between sellers and buyers. We do expect it to be relevant for the interaction between sellers, however, and will comment on the dynamic aspects where they are likely to be relevant.

2.2.1 Full information

Qualitatively, the equilibrium is as in Shaked and Sutton's (1982) model of vertical differentiation. We derive here the equilibrium for the specific formulation described above. In the *full info* treatment, all consumers observe the qualities of the products on offer. Suppose that the market is covered (as mentioned above, this will be the case in equilibrium for parameter values we have chosen). We solve the game backward from the price setting stage and assume

[4] The sessions for the *no info*, *full info*, and *subset* treatments were conducted first.

without loss of generality that $q_A > q_B$.[5] A consumer with taste for quality θ obtains utility $v + \theta q_A - p_A$ from buying firm A's product, and utility $v + \theta q_B - p_B$ from buying firm B's product. The value of θ at which a consumer is indifferent is θ^*, defined by

$$\theta^* q_A - p_A = \theta^* q_B - p_B \iff \theta^* = \frac{p_A - p_B}{q_A - q_B}. \tag{1}$$

Suppose $\underline{\theta} \leq \theta^* \leq \overline{\theta}$, and let $\Delta\theta \equiv \overline{\theta} - \underline{\theta}$. The probability that any consumer i prefers to buy from A is $1 - (\theta^* - \underline{\theta})/\Delta\theta$, while the probability that i prefers to buy from B is $(\theta^* - \underline{\theta})/\Delta\theta$. Since the θs are independently distributed, the expected demand for the firms is

$$D_A(p_A, p_B, q_A, q_B) = \frac{n}{\Delta\theta}\left(\overline{\theta} - \frac{p_A - p_B}{q_A - q_B}\right), \tag{2}$$

$$D_B(p_B, p_A, q_B, q_A) = \frac{n}{\Delta\theta}\left(\frac{p_A - p_B}{q_A - q_B} - \underline{\theta}\right). \tag{3}$$

From this we can derive the firms' best-response functions in the pricing game. Firm j's problem is

$$\max_{p_j} (p_j - cq_j) D_j(p_j, p_{-j}, q_j, q_{-j}),$$

from which we obtain the best responses

$$p_A = \frac{p_B + cq_A + \overline{\theta}(q_A - q_B)}{2}, \tag{4}$$

$$p_B = \frac{p_A + cq_B - \underline{\theta}(q_A - q_B)}{2}. \tag{5}$$

Solving for the Nash equilibrium prices, we find

$$p_A^* = \frac{q_A(2\overline{\theta} - \underline{\theta} + 2c) - q_B(2\overline{\theta} - \underline{\theta} - c)}{3}, \tag{6}$$

$$p_B^* = \frac{q_A(\overline{\theta} - 2\underline{\theta} + c) - q_B(\overline{\theta} - 2\underline{\theta} - 2c)}{3}. \tag{7}$$

For the specific parameter values we use in the experiment, we have $p_A^* = 3q_A - 2q_B$ and $p_B^* = q_A$.

Plugging the equilibrium prices into the profit function, and simplifying, we obtain

$$\pi_A(q_A, q_B) = \frac{n}{\Delta\theta}(q_A - q_B)\left(\frac{2\overline{\theta} - \underline{\theta} - c}{3}\right)^2, \tag{8}$$

$$\pi_B(q_B, q_A) = \frac{n}{\Delta\theta}(q_A - q_B)\left(\frac{\overline{\theta} - 2\underline{\theta} + c}{3}\right)^2. \tag{9}$$

As can be seen from these expressions, the high-quality firm's profit is increasing in its own quality whatever its rival's quality, and the low-quality firm's profit is decreasing in its own

[5] When $q_A = q_B$, all consumers buy from the firm with the lower price. In case prices are the same, assume that consumers randomize.

quality whatever its rival's quality. It follows that the game has two subgame perfect pure-strategy Nash equilibria (PSNE): one where firm A offers the highest possible quality while firm B offers the lowest, and another where firm B offers the highest and firm A the lowest. For our specific parameter values, the equilibria are $(q_A = 10, q_B = 1, p_A = 28, p_B = 10)$, and $(q_A = 1, q_B = 10, p_A = 10, p_B = 28)$. There is also a mixed-strategy Nash equilibrium (MSNE) in which both firms randomize between the highest and lowest quality. In this equilibrium, each firm chooses $q_j = 1$ with probability 1/5 and $q_j = 10$ with probability 4/5.

A buyer's expected surplus is given by

$$\begin{aligned} E(u_i) &= v + \Pr(\theta \leq \theta^*) \left[q_B E(\theta|\theta \leq \theta^*) - p_B \right] + \Pr(\theta > \theta^*) \left[q_A E(\theta|\theta > \theta^*) - p_A \right] \\ &= v + \frac{\theta^* - \underline{\theta}}{\Delta\theta} \left[q_B \frac{\theta^* + \underline{\theta}}{2} - p_B \right] + \frac{\overline{\theta} - \theta^*}{\Delta\theta} \left[q_A \frac{\theta^* + \overline{\theta}}{2} - p_A \right]. \end{aligned}$$

2.2.2 No information

The equilibrium concept we employ for the case where consumers cannot observe quality is perfect Bayesian equilibrium (PBE). That is, all players' strategies must be optimal given beliefs, and beliefs must be derived from equilibrium strategies using Bayes' rule whenever possible. Consider the following profile of strategies and beliefs. Both firms choose the lowest quality $(q_A = q_B = 1)$ and price at marginal cost $(p_A = p_B = c \cdot 1)$. Consumers believe the quality of the products on offer is the lowest possible (1). Consumers' out-of-equilibrium believes are that any firm choosing a different price also has the lowest quality. Given these beliefs, no firm has an incentive to unilaterally deviate, and beliefs are derived from equilibrium strategies. This clearly is a PBE of the game.

We cannot entirely exclude that there might be other equilibria in which prices signal quality. Even though there cannot be an equilibrium in which consumers' beliefs about q_j depend only on p_j (not p_{-j}) (why would a firm ever choose a high quality then?), it is possible that there is an equilibrium in which beliefs about qualities depend on the *pair of prices* (p_A, p_B). Because firms observe their competitor's quality, they can condition their own price on the qualities chosen at the first stage. Therefore, prices may jointly convey information about qualities. For the case of two levels of quality, Fluet and Garella (2002) show that there exist separating equilibria when the quality differential is sufficiently large. In Appendix A.1 we show that their result extends to our particular setup.

The theoretical literature has not addressed the question of whether there can be a signaling equilibrium when firms choose from a set of more than two possible levels of quality, and such an analysis is beyond the scope of this paper. Intuitively, however, it seems that adding more levels of quality makes it difficult to sustain a separating equilibrium. The result from the two-quality case suggests that quality levels need to be sufficiently far apart. With

	no info	*full info*	
		PSNE	MSNE
Quality (supplied)	1	5.5	8.2
$\|\Delta\text{Quality}\|$	0	9	2.88
Price (posted)	1	19	12.5
$\|\Delta\text{Price}\|$	0	18	5.76
Markup	0	13.5	4.32
Sellers' Surplus	0	60	19.2
Buyers' Surplus	86	74	116.8
Total Surplus	86	134	136
% Efficiency[a]	61.4	95.7	97.1

[a] % Efficiency = Total Surplus/140

Table 1: Equilibrium averages predicted by theory

our specification, a quality differential of 5 is needed for separation (see Appendix A.1), but actual quality levels in our experiment are only one unit apart.[6] Therefore, it is unlikely that separation would remain an equilibrium when there are more than two levels of quality.

If we accept the premise that the stage game has a unique equilibrium, then the only equilibrium in the finitely repeated version is to play the unique equilibrium of the stage game in every period. Thus, theory predicts that tacit collusion between sellers to set a price above marginal cost is not an equilibrium.

Table 1 shows the theoretically predicted averages for certain key variables in the *full info* and *no info* treatments. For example, the pure-strategy equilibrium in the *full info* treatment is associated with an average quality of 5.5 (= (10 + 1)/2) and an average price of 19. It is characterized by an expected absolute quality difference of 9, owing to sellers' strategy of maximum quality differentiation. The mixed-strategy equilibrium features higher quality, lower price, and lower differentiation on average as sellers often wind up with equal quality products. The theoretically predicted values for buyers' surplus illustrate the fact that the theory does not make sharp predictions as to the desirability of transparency from the point of view of consumer protection: consumers are better off in the *no info* treatment than in the pure-strategy equilibrium of the *full info* treatment. The values for sellers' surplus show that sellers are unambiguously better off under *full info*.

[6] One may wonder whether there is an incentive for firms to differentiate themselves, and thus choose only qualities that are sufficiently far apart in equilibrium. Suppose firms play only $q_A = 1$ and $q_B = 10$. Then, firm B has an incentive to deviate and choose a lower quality while keeping the same price. Firm A must punish this kind of behavior for maximal differentiation to be an equilibrium. But Firm A has no incentive to do so since its own payoff is unaffected by this deviation.

2.2.3 Subset of informed consumers

Once again, the literature has so far not analyzed the case of multiple quality levels. For the case of two quality levels, Yehezkel (2008) derives conditions under which separation can occur. He shows for a given quality differential that a separating equilibrium is sustainable when the number of informed consumers is sufficiently large. In Appendix A.2, we extend his result to the setup of our experiment by showing that the quality differential required to sustain a separating equilibrium decreases with the number of informed consumers. The intuition is that the presence of more informed consumers makes charging a high price less attractive for a low-quality firm than for a high-quality firm: at equal prices, informed consumers will not buy from the low-quality firm regardless of their taste for quality (that is, even though they might prefer the low-quality product at a lower price). Moreover, the low-quality firm's cost of mimicking the high-quality firm increases with the latter's price, because as the price increases, more and more uninformed consumers will refrain from buying. The Appendix shows that for our particular parameter values, a quality differential of 1 is sufficient to sustain a separating equilibrium. It also shows that the prices needed to sustain a separating equilibrium are greater than the prices under *full info*.

While this is of course no proof that separation would be sustainable when firms can choose between more than two quality levels, it is strongly suggestive of the fact that the presence of informed consumers makes the existence of a separating equilibrium, in which prices signal qualities, more likely. The informed consumers exert a positive informational externality on uninformed consumers: by creating a cost of mimicking high-quality firms, informed consumers discipline low-quality firms, thereby helping make prices informative about qualities. This benefits the uninformed consumers who can infer quality from prices.

2.2.4 Imperfect signal

To the best of our knowledge, there are no theoretical papers dealing with the case of imperfect information about quality. Intuitively, however, we would expect a similar argument as in the subset treatment to apply. That is, an imperfect but informative signal about quality should make it costly for low-quality firms to choose a high price, because consumers are likely to correctly identify which of the two products on offer is of higher quality. Moreover, we have calibrated the imperfect information treatment in such a way that an average consumer is about as informed as in the subset treatment, in the sense defined in Section 2.1 above. Therefore, we would expect that a signaling equilibrium could also occur in the imperfect information treatment, and that the theoretical predictions for both treatments would be similar.

2.3 Procedures

The experiment was conducted in the CentERlab of Tilburg University in the Netherlands. We organized 16 computerized sessions in between March 24, 2010 and May 26, 2010.[7] Each session, including briefing and debriefing, lasted for approximately 75 minutes. Subjects were recruited through e-mail lists of students interested in participating in experiments. 210 subjects in total participated in the experiment, of which 47% were female. The average age of subjects was 22.85 years, and people of 22 different nationalities participated. 92% of the subjects were students in the Faculty of Economics and Business Administration of Tilburg University, and 90% had participated in other economic experiments before. While payment was based on performance, the average earnings for subjects were around EUR 12.

Subjects took the role of either sellers or imperfectly informed buyers. The fully informed buyers were automated, i.e., their role was played by the computer. Fully informed buyers are present in the *full info* and *subset* treatments. We chose this procedure because informed buyers have a rather mechanical task, which is to calculate which of two possible quality-price pairs maximizes their payoff, given their induced taste for quality. We discuss the implications of having automated buyers in Section 3.3. As a consequence of this procedure, the number of subjects per session differed depending on the treatment. There were six subjects per session in the *full info* treatment (of which we ran four sessions in total), twelve in the *subset* treatment (five sessions), and 18 in the *no info* (three sessions) and *signal* (four sessions) treatments.

It was explained to subjects that their earnings would depend on their own choices as well as those of the other participants in the session. Subjects were informed that the exchange rate for converting their payoff from experimental currency units to EUR was 50 to 1, and that their earnings would be paid out anonymously and in private. They were told that they would be randomly assigned to be either a buyer or a seller, that sellers would be in groups of two and buyers in groups of four, and that these groups would remain the same throughout the entire experiment, with random rematching of groups occurring every period. They were instructed that as a seller they would have to choose the price and the "grade" of a product. Following Holt and Sherman (1990), we chose the more neutral labeling of quality as grade to avoid the possible distortion of results that might result from the positive connotations of the word quality. It was explained to the subjects that for sellers, a high grade is more costly to provide than a low grade, while for buyers, a high grade is more valuable than a low grade. They were instructed that as a buyer they would have to choose from which of the two sellers offering products to their group to buy, and that they could also refrain from buying.

[7] The experiment was programmed and run using the software z-Tree (Fischbacher, 2007).

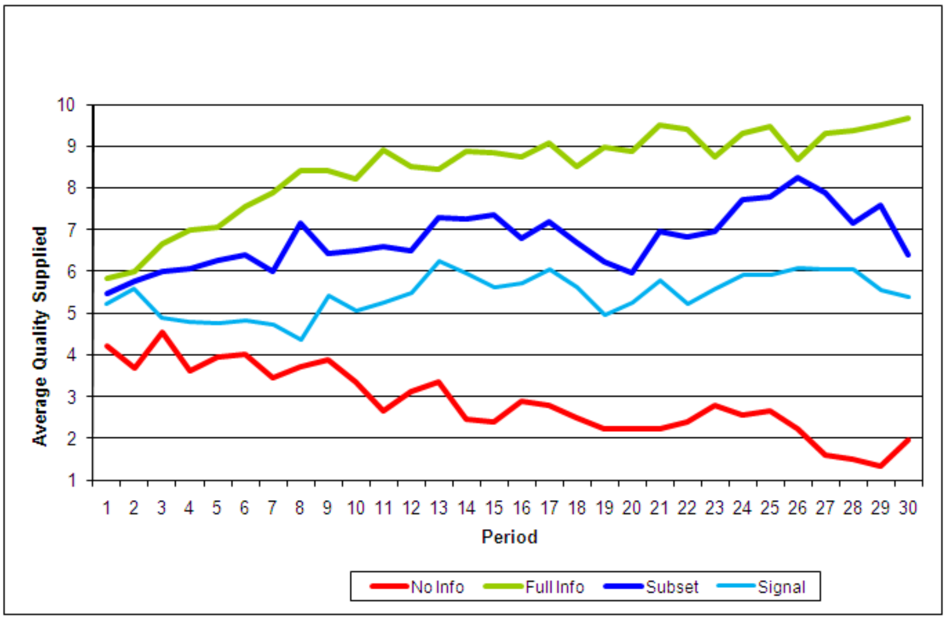

Figure 1: Average quality supplied by treatment and period

The instructions explained in detail how the payoff of each participant would be determined: for sellers, as a function of the price and the number of units sold, minus the cost which depends on the grade chosen; for buyers, as a function of the grade obtained and the price paid, and depending on their individual valuation parameter θ. The cost parameter c was revealed only to the sellers, once the computerized part of the experiment started. Keeping cost information private is standard procedure and is known to improve convergence to competitive equilibrium (Smith, 1994). The instructions are available from the authors upon request.

3 Results

3.1 Description of the data

In this subsection, we describe the basic patterns in the data for each treatment and contrast them with the theoretical predictions. Figure 1 shows the observed averages of quality supplied by treatment and period. All four treatments start out at relatively similar levels of quality. Subsequently, quality supplied trends downward under *no info* and upward in the other three treatments. Quality is highest in the *full info* treatment, followed by the *subset* and *signal* treatments. Figure 2 shows the observed averages of prices posted by treatment and period. Prices in the *full info* treatment are relatively low and stable over time, while prices in the other three treatments are initially somewhat higher and then trend downward

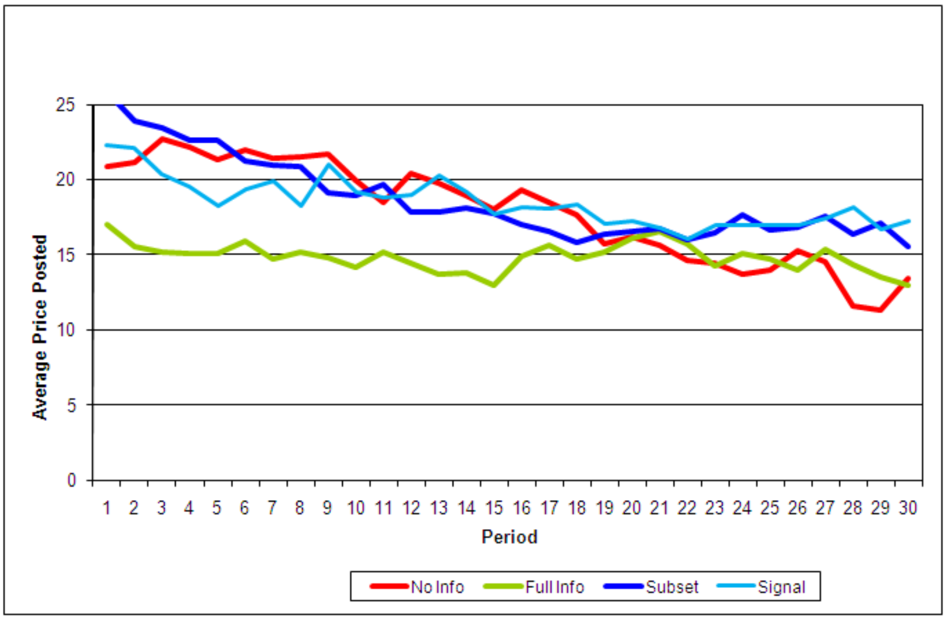

Figure 2: Average price posted by treatment and period

toward the *full info* level.

Tables 2 and 3 confirm these observations. The tables show the observed averages over the entire 30 periods and the final 15 periods, respectively, for the same variables of interest as those for which Table 1 gives theoretical benchmarks, . In the *full info* treatment, average quality exceeds the theoretical benchmarks for both PSNE and MSNE. This is especially visible in the final 15 periods, where average quality is 9.14. The observation that the average quality difference for the final 15 periods is .99 reflects the fact that both sellers choose similar, high levels of quality. They do not vertically differentiate their products as theory would predict.

In the *no info* treatment, average quality is low, albeit slightly above the theoretically predicted level. This provides evidence that the predicted lemons outcome does indeed tend to materialize. Prices, however, are nowhere near the prediction of marginal-cost pricing, and even exceed the *full info* levels. Apparently, buyers do not anticipate the lack of quality in early periods, allowing sellers to get away with high prices. The fact that a downward pressure on prices seems to emerge toward the later periods suggests that buyers do eventually learn to expect low quality from both sellers, however.

The *subset* and *signal* treatments show the highest price levels and also the highest level of quality differentiation. Average quality in these treatments is between the *no info* and *full info* levels. As argued in Section 2.2 above, high prices are consistent with a signaling interpretation: sellers may try to signal high quality through high prices.

	no info	*signal*	*subset*	*full info*
Quality (offered)	2.87	5.44	6.78	8.46
$\lvert\Delta$Quality$\rvert$	1.80	2.33	2.44	1.14
Price (posted)	17.86	18.47	18.66	14.85
$\lvert\Delta$Price$\rvert$	3.48	3.73	3.82	3.43
Markup	14.99	13.03	11.87	6.40
Sellers' Surplus	50.47	45.53	41.33	22.00
Buyers' Surplus	37.15	64.16	80.22	112.17
Total Surplus	87.61	109.68	121.55	134.17
% Efficiency[a]	61.69	78.98	85.93	94.95

[a] % Efficiency = Total Surplus/$(40 + 10\sum_{i=1}^{4}\theta_i)$

Table 2: Observed treatment averages over all 30 periods

	no info	*signal*	*subset*	*full info*
Quality (offered)	2.26	5.67	7.10	9.14
$\lvert\Delta$Quality$\rvert$	1.15	2.46	2.54	.99
Price (posted)	15.05	17.26	16.60	14.86
$\lvert\Delta$Price$\rvert$	2.90	3.48	3.32	3.54
Markup	12.80	11.59	9.51	5.72
Sellers' Surplus	42.66	42.76	35.20	18.44
Buyers' Surplus	43.79	71.80	92.13	121.06
Total Surplus	86.44	114.56	127.34	139.50
% Efficiency[a]	60.60	82.59	89.38	97.82

[a] % Efficiency = Total Surplus/$(40 + 10\sum_{i=1}^{4}\theta_i)$

Table 3: Observed treatment averages over the final 15 periods

3.2 Comparison across treatments

3.2.1 Nonparametric tests

To test for differences in quality across the treatments that seem to be apparent in Figure 1 and Tables 2 and 3, we performed conservative nonparametric rank-sum tests that treat each experimental session as one data point. Tables 4 and 5 show the (one-tailed) p-values from Mann-Whitney rank-sum tests under the null hypothesis of no difference between any two treatments, based on the data from all 30 periods and from the final 15 periods, respectively.

Because there are only three to five data points in each treatment, the tests have modest statistical power to detect differences across treatments. Yet, even with this conservative test, we can reject equality of the level of quality supplied in the *no info* versus the *subset* and *full info* treatments at the 5% level, and versus the *signal* treatment at the 10% level over all 30 periods and at the 5% level over the final 15 periods. Moreover, we can reject equality of the *signal* and *subset* treatments versus the *full info* treatment at the 5% level.[8] Thus, more information appears to consistently improve the level of quality supplied statistically.

Not reported here are the results of comparisons for price posted. None of the differences in price turn out to be statistically significant.

	signal	*subset*	*full info*
no info	0.057	0.018	0.029
signal		0.095	0.014
subset			0.008

Table 4: p-values for one-sided rank sum tests of differences in quality supplied (all 30 periods)

	signal	*subset*	*full info*
no info	0.029	0.018	0.029
signal		0.056	0.014
subset			0.008

Table 5: p-values for one-sided rank sum tests of differences in quality supplied (final 15 periods)

3.2.2 Parametric tests

We also ran parametric tests, which have the advantage of using more of the information contained in the data and therefore have greater power to detect differences. The statistical model used for the parametric analysis is a random effects model with a clustered error structure at the session level. The functional form is

$$y_{gst} = \alpha + X_s'\beta + u_s + \varepsilon_{gst} \tag{10}$$

where

$$X_s'\beta = \beta_1 \text{Signal}_s + \beta_2 \text{Subset}_s + \beta_3 \text{FullInfo}_s.$$

The indices g, s, and t designate group (i.e. the particular match of a seller with a buyer group), session and period, respectively. Note thus that we have three observations of the independent variable y_{gst} per period as there are three matchings of seller and buyer groups in each period. Signal_s, Subset_s, and FullInfo_s are zero-one dummies attaining a value of one in the sessions of the *signal*, *subset* and *full info* treatment respectively and a value of zero otherwise. Note that in the *no info* treatment these dummies are jointly zero. The term u_s represents a random intercept which is assumed to be randomly drawn from a normal

[8] Some of the significance levels may seem counterintuitive. This stems from the different number of sessions for each treatment.

distribution with mean zero for each individual session. This accounts for differences between sessions caused by unobserved heterogeneity of subjects, that is, differences between the subjects which can neither be observed nor controlled for by the experimenter. Finally, ε_{gst} is an error term with an expected value of zero that is assumed to be uncorrelated with both with the random intercept u_s and all independent variables in every single period. To sustain that assumption (which is crucial in order to obtain unbiased estimators) in our framework, it is however essential to cluster ε_{gst} for each individual session. This is necessary because it cannot be ruled out that within each individual session there are possible dynamic session effects that would violate this assumption if not accounted for (Fréchette, 2007).

Tables 6 through 9 report the results of the parametric tests for differences in quality supplied and price posted in the *no info* baseline versus the three other treatments. The results for quality confirm the results from the nonparametric tests: the difference in quality is strongly significant for all three treatments. The differences in prices still do not come out as statistically significant, suggesting that the lack of difference detected by the nonparametric tests is not due to their lack of power.

	Coef.	Std. Err.	z	$P > \|z\|$	95% Conf.	Interval
Constant	2.87	.70	4.09	0.000	1.49	4.25
Signal	2.57	.87	2.95	0.003	.86	4.28
Subset	3.91	.82	4.78	0.000	2.31	5.52
Full Info	5.58	.71	7.92	0.000	4.20	6.97
σ_u	.99					
σ_ε	2.04					
ρ[a]	.19					

[a] Fraction of variance due to unobserved heterogeneity

Table 6: Parametric regression for quality supplied (all 30 periods)

	Coef.	Std. Err.	z	$P > \|z\|$	95% Conf.	Interval
Constant	2.26	.77	2.92	0.003	.74	3.77
Signal	3.42	.89	3.83	0.000	1.67	5.16
Subset	4.84	.86	5.66	0.000	3.16	6.52
Full Info	6.88	.78	8.78	0.000	5.35	8.42
σ_u	.93					
σ_ε	1.99					
ρ[a]	.18					

[a] Fraction of variance due to unobserved heterogeneity

Table 7: Parametric regression for quality supplied (final 15 periods)

	Coef.	Std. Err.	z	$P > \|z\|$	95% Conf.	Interval
Constant	17.86	2.00	8.92	0.000	13.94	21.79
Signal	.61	2.88	0.21	0.833	−5.04	6.25
Subset	.79	2.68	0.30	0.767	−4.46	6.05
Full Info	−3.01	2.14	−1.41	0.159	−7.21	1.18
σ_u	3.83					
σ_ε	5.58					
ρ[a]	.32					

[a] Fraction of variance due to unobserved heterogeneity

Table 8: Parametric regression for price posted (all 30 periods)

	Coef.	Std. Err.	z	$P > \|z\|$	95% Conf.	Interval
Constant	15.05	2.83	5.32	0.000	9.51	20.60
Signal	2.21	3.49	0.63	0.527	−4.64	9.05
Subset	1.55	3.21	0.48	0.629	−4.74	7.85
Full Info	−.19	3.0	−0.06	0.949	−6.07	5.69
σ_u	4.01					
σ_ε	4.62					
ρ[a]	.43					

[a] Fraction of variance due to unobserved heterogeneity

Table 9: Parametric regression for price posted (final 15 periods)

3.3 Discussion

Is information the driving force behind our results, or is it something else? One potential criticism concerns our use of automated buyers. Therefore, by construction, the treatments in which there is more information also have more automated buyers. While in the *no info* and *signal* treatments, all buyers are played by subjects, half of the buyers in the *subset* treatment and all of the buyers in the *full info* treatment are played by the computer. Do our results reflect the fact that unlike human subjects the computer makes no mistakes?

Mistakes by buyers clearly reduce efficiency by not maximizing surplus, given the available qualities on offer. But efficiency is mainly driven by which qualities the sellers choose, and we find a large difference in these between treatments. Is there reason to think that mistakes by buyers would decrease the quality *on offer*? First, let us think about when we expect people to make mistakes. When they are facing offers $(q_A, p_A), (q_B, p_B)$ such that $q_A \geq q_B$ and $p_A \leq p_B$, we don't expect many mistakes. We expect mistakes when $q_A \geq q_B$ and $p_A \geq p_B$. Do such mistakes favor the low-quality seller (thereby providing stronger incentives to offer low quality ex ante, compared to a situation with automated buyers)? There is no obvious reason why they should. Suppose for example $q_A = 1, p_A = 11$ and $q_B = 10, p_B = 31$. Should a buyer with $\theta = 2.89$ buy from A or B? From A, he gets 11.89, from B he gets 17.9, so he should choose B. But there is no reason why mistakes should generally favor A. Buyers with a different θ could get it wrong and buy from B. Do sellers have an incentive to induce, or prevent, mistakes? In our example, the high-quality seller would want to price lower in order to make it clear that his product is more attractive. This reduces the profitability of his product. But the low-quality seller might want to decrease price to make it less clear which product is more attractive (and thus induce mistakes). Moreover, while for this particular θ, it is clear which seller gains from mistakes and which seller loses out, sellers do not know buyers θs. For some θs, the high-quality seller will instead have incentives to induce mistakes, and the low-quality seller to prevent them. Since each theta is equally likely, there does not seem to be any reason to expect mistakes to systematically favor low-quality sellers.

Buyer mistakes also cannot explain why we detect a significant difference in quality between the *no info* and *signal* treatment, since there are no automated buyers in either. On these grounds we conclude that information is the more likely source of the differences, and that automated buyers play, at most, a minor role.

4 Conclusion

We have presented the results of an experiment investigating the role of transparency in a duopoly market with vertical differentiation. Firms choose both the quality and the price of their product, while consumers differing in their taste for quality choose from which firm to buy. We have compared three different treatments in which we vary the share of consumers who are informed about quality. Specifically, we have a full-information treatment in which all consumers are informed, a no-information treatment in which none of them are informed, and a subset treatment in which half of them is informed and the other half uninformed. We have found that, contrary to theoretical predictions, firms do not differentiate quality under full information. Rather, both tend to offer services of similar, and high, quality, entailing more intense price competition than predicted by theory. Under no information, we observe a "lemons" outcome where quality is low. At the same time, firms manage to maintain prices substantially above marginal cost. In the subset treatment, quality is close to the full-information level. This suggests that a signaling equilibrium tends to materialize when some consumers are informed and others are not. Our results have implications for consumer protection policies aimed at increasing transparency, such as recent EU legislation mandating disclosure of information on Internet Service Providers' network management.

Appendix A Signaling equilibria with two levels of quality

A.1 The *no info* treatment

Suppose there are only two possible levels of quality, q_L and q_H. A strategy in this game must include prices for every possible combination of qualities chosen at the quality-choice stage. A symmetric, fully revealing equilibrium in the pricing game requires the following price structure:

- $p_A = p_B = p_{LL}$ when $q_A = q_B = q_L$,
- $p_A = p_{LH}$, $p_B = p_{HL}$ when $q_A = q_L$, $q_B = q_H$,
- $p_A = p_{HL}$, $p_B = p_{LH}$ when $q_A = q_H$, $q_B = q_L$, and
- $p_A = p_B = p_{HH}$ when $q_A = q_B = q_H$,

with $p_{LL} < p_{HH}$ and $p_{LH} < p_{HL}$ and all four prices being different from each other.

Consider the following profile of strategies and beliefs. Consumers believe that $(q_j = q_H, q_{-j} = q_L)$ if and only if $p_j = p_{HL}$ and $p_{-j} \notin \{p_{LL}, p_{HH}\}$; they believe that $q_A = q_B = q_H$ if and only if $p_A = p_B = p_{HH}$; and they believe that $q_A = q_B = q_L$ otherwise. Prices satisfy the following conditions:

$$\begin{aligned} p_{LL} &= cq_L, & (11)\\ p_{HH} &= cq_H, & (12)\\ p_{LH} &= \frac{p_{HL} + cq_L - \underline{\theta}(q_H - q_L)}{2}, & (13)\\ (p_{LH} - cq_L)\frac{n}{\Delta\theta}\left[\frac{p_{HL} - p_{LH}}{q_H - q_L} - \underline{\theta}\right] &\geq (p_{HH} - cq_L)n, & (14)\\ (p_{HL} - cq_H)\frac{n}{\Delta\theta}\left[\overline{\theta} - \frac{p_{HL} - p_{LH}}{q_H - q_L}\right] &\geq (p_{LH} - cq_H)n, & (15)\\ (p_{LH} - cq_L)\frac{n}{\Delta\theta}\left[\frac{p_{HL} - p_{LH}}{q_H - q_L} - \underline{\theta}\right] &\geq (p_{HL} - cq_L)\frac{n}{2\Delta\theta}\left[\overline{\theta} - \frac{p_{HL} - v}{q_L}\right]. & (16) \end{aligned}$$

Expressions (11) and (12) say that when both firms have the same quality, we get the Bertrand outcome of marginal-cost pricing. Given consumers' beliefs, no firm has an incentive to unilaterally deviate: in both cases, choosing a higher price will lead to zero demand because whatever the deviating price, consumers' beliefs will be that both firms offer low quality, while choosing a lower price is unprofitable. Expression (13) says that when one firm has low and the other high quality, p_{LH} must be a best response to p_{HL} (and thus given by (5)). This is because no price chosen by the low-quality firm will change consumers' beliefs that the firm offering price p_{HL} has high quality and the other firm low quality, except for p_{LL} (which is not an attractive deviation) and p_{HH} (which is the subject of the following condition). Thus, the low-quality firm is free to choose any price without affecting beliefs, which means that the equilibrium price must be a best response. Expression (14) says that the low-quality firm must not want to deviate to p_{HH}. In that case, consumers would believe that both firms offer low quality, and since $p_{HH} < p_{HL}$, the low-quality firm will get the entire demand. Expression (15) says that the high-quality firm must not want to deviate to p_{LH} $(-\varepsilon)$, in which case it would grab the entire demand but consumers would believe it offers low quality. Finally, expression (16) says that the low-quality firm must not want to deviate to p_{HL}, in which case it would it would grab half the demand that arises at this price given that consumers believe that both firms offer low quality. (Note that (16) is based on p_{HL} being such that some low-valuation consumers buy and others don't. In general, it could be the case that none or all consumers are buying.)

Without deriving general conditions under which (11) through (16) can be simultaneously satisfied, we now give an example based on our particular parameter values to show that it is possible that they are. Suppose $q_L = 1$ and $q_H = 6$. Then, the following prices constitute an equilibrium: $p_{LL} = 1$, $p_{HH} = 6$, $p_{LH} = 10$, $p_{HL} = 24$. As previously discussed, firms have no incentive to deviate unilaterally when they have the same qualities. When they have different qualities, the low-quality firm obtains an equilibrium profit of $(10 - 1)(4/3)[(24 - 10)/(6 -$

$1) - 1] = 21.6$. By deviating to $p = p_{HH} = 6$, it would obtain $(6-1) \cdot 4 = 20$. By deviating to $p = p_{HL} = 24$, it would obtain zero. The high-quality firm obtains an equilibrium profit of $(24-6)(4/3)[4-(24-10)/(6-1)] = 28.8$. By deviating to $p = p_{LH} = 10$, it would obtain $(10-6) \cdot 4 = 16$. No other deviation is potentially profitable.

This example does not generalize to cases where the difference in quality is smaller, however. The intuition why signaling is only possible when the quality differential is sufficiently large is the following. Consumers have out-of-equilibrium beliefs that make it unattractive for firms to deviate to prices such that the resulting *pair* of prices is not part of an equilibrium. To make sure that the low-quality firm does not want to deviate to p_{HL}, this price must be sufficiently high for demand to be low when consumers believe quality is low. But at the same time, a high p_{HL} means that the equilibrium demand (and payoff) for the high-quality firm will not be large, in which case it becomes attractive to deviate and steal the demand of the low-quality firm. If the quality differential is small, these two requirements cannot be simultaneously met. Notice that in the above example, the price charged by the high-quality firm (24) is higher than the full-information equilibrium price (which, from (6), is 16). This illustrates that separation occurs through an increase in price.

A.2 The *subset* treatment

Let n^I denote the number of informed and n^U the number of uninformed consumers (with $n^I + n^U = n$). Informed consumers behave exactly as the consumers under full information, analyzed in Section 2. That is, expected demand from the informed consumers given prices and qualities is given by (2) and (3), replacing n by n^I. Suppose moreover that uninformed consumers' beliefs are exactly the same as under no information (see Appendix A.1). For a symmetric fully revealing equilibrium, we once again need prices $(p_{LL}, p_{HH}, p_{HL}, p_{LH})$ defined as above. Suppose prices satisfy the following conditions:

$$p_{LL} = cq_L, \tag{17}$$

$$p_{HH} = cq_H, \tag{18}$$

$$p_{LH} = \frac{p_{HL} + cq_L - \underline{\theta}(q_H - q_L)}{2}, \tag{19}$$

$$(p_{LH} - cq_L)\frac{n}{\Delta\theta}\left[\frac{p_{HL} - p_{LH}}{q_H - q_L} - \underline{\theta}\right] \geq (p_{HH} - cq_L)\left[n^U + \frac{n^I}{\Delta\theta}\left(\frac{p_{HL} - p_{HH}}{q_H - q_L}\right)\right], \tag{20}$$

$$(p_{HL} - cq_H)\frac{n}{\Delta\theta}\left[\overline{\theta} - \frac{p_{HL} - p_{LH}}{q_H - q_L}\right] \geq (p_{LH} - cq_H)(n^U + n^I), \tag{21}$$

$$(p_{LH} - cq_L)\frac{n}{\Delta\theta}\left[\frac{p_{HL} - p_{LH}}{q_H - q_L} - \underline{\theta}\right] \geq (p_{HL} - cq_L)\frac{n^U}{2\Delta\theta}\left[\overline{\theta} - \frac{p_{HL} - v}{q_L}\right]. \tag{22}$$

This profile of strategies and beliefs is an equilibrium. To interpret the conditions on prices, notice first that expressions (17) through (19) are exactly the same as (11) through (13).

The crucial conditions for existence of a separating equilibrium, however, are expressions (20) through (22). Expression (20) again formalizes the idea that the low-quality firm must not want to deviate to p_{HH}. Now though, such a deviation is less profitable because only the uninformed consumers will buy in their entirety from the low-quality firm. The informed consumers only buy from it if their taste for quality is sufficiently low, and the deviation is sure to reduce the low-quality firm's profits from informed consumers since p_{HH} is not a best response to p_{HL}. Thus, the right-hand side of (20) is lower than that of (14). Expression (21) says that the high-quality firm must not want to deviate to $p_{LH} - \varepsilon$. With such a move, it could capture both the uninformed and the informed consumers: the uninformed, because they would believe that both firms offer q_L and thus buy from the one charging the lower price, the informed, because they observe who has the higher quality. The result is the same as before: because $n^I + n^U = n$, the right-hand side of (21) is the same as that of (15). Expression (22) says that the low-quality firm must not want to deviate to p_{HL}. Again, though, such a deviation is less profitable than under no information because none of the informed consumers would buy from the low-quality firm in that case. The only consumers buying from it are half of those whose willingness to pay, given they believe both firms offer q_L, is sufficiently high. Thus, the right-hand side of (22) is lower than that of (16).

In sum, both (20) and (22) are easier to satisfy than the corresponding conditions under no information; moreover, the right-hand sides of both expressions increase with n^U and decrease with n^I. At the same time, (19) remains unchanged. We conclude that, the greater the number of informed consumers and the smaller the number of uninformed ones, the easier it is to sustain a separating equilibrium, for a given quality differential.

With the parameter values used in our experiment, where there are two informed and two uninformed consumers, and again fixing $q_L = 1$, it is sufficient that $q_H = 2$ to support separation. The following prices then constitute an equilibrium: $p_{LL} = 1$, $p_{HH} = 2$, $p_{LH} = 3$, $p_{HL} = 6$. Firms have no incentive to deviate unilaterally from these prices when they have the same qualities. When they have different qualities, the low-quality firm obtains an equilibrium profit of $(4/3)(3-1)[(6-3)/(2-1)-1] = 5.33$. By deviating to $p = 2 = p_{HH}$, it would obtain $(2-1)(2+(2/3)(6-2)/(2-1)) = 4.67$. By deviating to $p = p_{HL} = 6$, it would obtain $(6-1)(2/2) = 5$. The high-quality firm obtains an equilibrium profit of $(4/3)(6-2)[4-(6-3)/(2-1)] = 5.33$. By deviating to $p = 3 = p_{LH}$, it would obtain $(3-2)\cdot 4 = 4$. By deviating to $p = 4.5$, which is the best response to p_{LH}, it would obtain $(4.5-2)(2/3)(4-(4.5-3)/(2-1)) = 4.17$. No other deviation is potentially profitable. Finally, notice once again that in the above example, the price charged by the high-quality firm (6) is higher than the full-information equilibrium price (which, from (6), is 4).

References

Akerlof, George (1970): The Market for Lemons: Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics*, Vol. 89, pp. 488–500.

Fischbacher, Urs (2007): z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, Vol. 10(2), pp. 171–178.

Fluet, Claude and Garella, Paolo G. (2002): Advertising and Prices as Signals of Quality in a Regime of Price Rivalry. *International Journal of Industrial Organization*, Vol. 20, pp. 907–930.

Fréchette, Guillaume (2007): Session Effects in the Laboratory. Working Paper, New York University.

Holt, Charles and Sherman, Roger (1990): Advertising and Product Quality in Posted-Offer Experiments. *Economic Inquiry*, Vol. 28, pp. 39–56.

Schuett, Florian (2010): Network Neutrality: A Survey of the Economic Literature. *Review of Network Economics*, Vol. 9(2), pp. 1–13.

Shaked, Avner and Sutton, John (1982): Relaxing Price Competition Through Product Differentiation. *Review of Economic Studies*, Vol. 49(1), pp. 3–13.

Smith, Vernon L. (1994): Economics in the Laboratory. *Journal of Economic Perspectives*, Vol. 8(1), pp. 113–131.

Yehezkel, Yaron (2008): Signaling Quality in an Oligopoly When Some Consumers Are Informed. *Journal of Economics and Management Strategy*, Vol. 17(4), pp. 937–972.