very preliminary - please do not cite

A Laboratory Study of Advertising and Price Competition*

by

John Morgan[1], Henrik Orzen[2] and Martin Sefton[3]

October 22, 2002

* We are grateful to the ESRC for funding the experiments under grant R000 22 3344. Orzen's work was supported by the Gottlieb Daimler- und Karl Benz- Foundation.

1. Haas School of Business, University of California, Berkeley. e-mail: morgan@haas.berkeley.edu

2. School of Economics, University of Nottingham, Nottingham, NG7 2RD, United Kingdom. e-mail: lexho@nottingham.ac.uk

3. School of Economics, University of Nottingham, Nottingham, NG7 2RD, United Kingdom. e-mail: martin.sefton@nottingham.ac.uk

Abstract

We consider a model of costly advertising and price competition in which sellers face two types of buyer. "Brand-loyal buyers" are attached to a particular seller, while "Bargain-hunters" buy from the seller who advertises the lowest price. Equilibrium requires two levels of mixed strategy: sellers randomize between advertising and not advertising, and advertisers randomize over prices. In this model, advertising fees play a facilitating role, reducing the amount of competition and allowing sellers to make higher profits.

We use a laboratory experiment with automated buyers to investigate the effects of advertising fees and the number of competitors on advertising and pricing decisions. As in previous experiments involving mixed strategy predictions, there are some discrepancies between subjects' observed and predicted decisions, but comparative static predictions are borne out by the data. In particular, in most treatments subjects advertise too much and set advertised prices too low, but, as predicted, advertising propensities and advertised prices move in the predicted direction as underlying market parameters change. Predictions about how advertising fees affect profits, however, require not just that advertising propensities and advertised prices move in the right direction as fees change, but also that they move by the right magnitudes. Our experimental data offer, at best, limited support for the prediction that advertising costs play a facilitating role. In some of our sessions sellers do succeed in earning higher profits when advertising fees are higher, but in many of our sessions the opposite effect is observed.

Keywords: Experiments, Price Dispersion

JEL Classification Numbers: C72, C92

1. Introduction

In contrast to the classical theory of perfectly competitive markets, many markets are characterized by imperfect price information. Consumers often do not know the prices charged by all sellers in a market, and buyers and sellers must often incur substantial costs to discover or transmit this information. A now well-established literature analyzes how various search frictions can generate imperfect information, and hence affect market performance, and this literature shows that many of the properties of perfectly competitive markets do not carry over to markets with imperfect price information. For example, in a market where some consumers are better informed than others about what prices are available, the Law of One Price may not hold. That is, in equilibrium, different sellers may charge different prices for a homogeneous product. Moreover, changes in the underlying structure of such markets often have implications that differ quite strikingly from the perfectly competitive case.

For example, Varian's 'Model of Sales' (Varian, 1980) analyzes price competition among identical sellers supplying a homogeneous product. Demand for the product comes from two types of consumer. Informed consumers know the prices charged by different sellers and buy at the lowest price (as long as this does not exceed their reservation price). Uninformed consumers do not know what prices are available, and simply choose a seller at random and buy from this seller (again, supposing that seller's price does not exceed their reservation price). In this model the unique symmetric equilibrium involves price dispersion as sellers use mixed strategies to generate prices.

In a previous paper (Morgan, Orzen and Sefton, 2001) we derived some comparative static implications of this model, and reported an experiment designed to test them. The data showed substantial and persistent price dispersion, and although empirical price distributions deviated from the theoretical distributions, the model was successful in predicting how average prices varied with changes in the underlying market structure. In particular, we noted that, rather intuitively, prices are predicted to decrease as the proportion of informed consumers increases. This prediction was strongly supported by our experimental data. We also noted that, less intuitively, prices are predicted to increase with the number of competing sellers. Again, this prediction found strong support in our experimental data.

One feature of Varian's model is that the composition of the market faced by a seller - the numbers of informed and uninformed consumers - is exogenous. In a more recent paper Baye and Morgan (2001) study a market where informed consumers receive their information through an "information gatekeeper" - this might be a newspaper or internet site that publicizes prices. In order to advertise their price, a seller must pay an advertising fee set by the gatekeeper. The model captures the idea that even when some consumers are looking for bargains, sellers must incur costs in order to attract their attention. When the advertising fee is zero, the model is identical to Varian's: all sellers advertise and face a mix of informed consumers and uninformed consumers. How does introducing a cost of advertising affect the market?

In theory the answer is quite clear-cut. Higher advertising fees reduce sellers' propensity to advertise and result in less competitive pricing, so that both informed and uninformed consumers can expect to pay higher prices when the advertising fee is higher. Thus, higher advertising fees hurt consumers. On the other hand, sellers benefit from higher advertising fees - the reason being that the higher prices charged in equilibrium outweigh the direct, negative, impact on profits of higher advertising costs. Thus, advertising fees can be viewed as a facilitating device, reducing the incentive of sellers to undercut one another, and enabling sellers to attain higher profits.

This theoretical prediction relies on a quite complicated equilibrium in which two levels of mixed strategy are employed. First, sellers randomize between advertising and not advertising. Second, sellers that advertise used mixed strategies to generate prices. Thus, while there is a unique symmetric equilibrium, the complexity of computing equilibrium strategies make it far from obvious that subjects will use equilibrium strategies, or even be able to approximate them.

In previous experiments studying the predictive power of mixed strategy equilibria in normal form games systematic departures from equilibrium predictions are commonly observed (e.g. see the review in Walker and Wooders, 2001). Nevertheless, comparative static predictions are often borne out - behavior changes with variations in game payoffs and equilibrium does a good job of predicting these changes (for example see Binmore et al., 2001, particularly Figure 5). A similar pattern is observed in applied settings, such as posted price markets, where equilibria only exist in mixed strategies - behavior deviates from theoretical predictions in terms of levels, but conforms with comparative static predictions (Brown-Kruse *et al*., 1994, Morgan *et al*. 2001).

This paper reports an experimental test of whether increases in advertising fees allow sellers to post higher prices and achieve higher profits in a simple version of the Baye and Morgan model. The results from our study are consistent with the pattern from previous experiments on mixed strategy equilibria in that (i) we observe systematic differences between observed and predicted behavior for any given treatment, (ii) the theory correctly predicts the direction of changes in behavior across treatments - as advertising costs are increased, subjects advertise less frequently and post higher prices.

However, support for comparative static predictions at the level of subject decisions does not imply support for comparative static predictions at the level of market outcomes -- i.e. expected profits and expected prices paid by different types of consumer. In fact we find that while the prices paid by consumers increase with advertising costs, sellers' profits do not always increase with advertising costs.

The remainder of the paper is organized as follows. In Section 2 we briefly describe the theoretical model, in Section 3 the experimental design and procedures, and in Section 4 the results. Section 5 concludes.

2. Theory

2.1 A Model of Advertising and Price Competition

In this section, we study a variant of Baye and Morgan's (2001) model of information gatekeepers. Consider a homogeneous product market in which there are n competing sellers. Each seller is identical and has zero cost of production and no capacity constraints. Each seller simultaneously chooses a price, p, and decides whether or not to advertise. A seller that decides to advertise pays an advertising fee, f, to do so.

Demand in this market comes from N consumers, each of whom has an identical reservation price, r, for one unit of the good. We suppose that m consumers are "bargain-hunters" and make purchasing decisions by comparing advertised prices. These consumers each purchase one unit from the seller advertising the lowest price, provided that this price does not exceed r. In the event that no sellers advertise, or in the event that all advertised prices exceed r, these consumers then choose a seller at random and buy from this seller as long as its price does not exceed r. Otherwise, these consumers will not buy the product. The remaining $N - m$ consumers are "brand-loyal" consumers who are evenly divided among the n sellers and simply purchase from the seller to whom they are loyal, as long as that seller's price does not exceed r. Otherwise, they do not purchase the item.

2.2 Equilibrium

For sellers to advertise we require that the advertising fee not be set too high. Following Baye and Morgan, we assume that $f < \frac{rm(n-1)}{n}$. Using techniques parallel to Baye and Morgan, one can show that the model has a unique symmetric equilibrium. In equilibrium sellers advertise with probability

(1) $$a = 1 - \left(\frac{nf}{(n-1)mr} \right)^{\frac{1}{n-1}}.$$

When a seller does not advertise, it chooses a price equal to r. When it does advertise, a seller prices according to the cumulative distribution function

(2) $$F(p) = \frac{1}{a}\left(1 - \left(G(p)\right)^{\frac{1}{n-1}}\right)$$

with support $[p_0, r]$ where

(3) $$G(p) = \frac{(r-p)(N-m)/n + (1-a)^{n-1}mr}{mp}$$

and

(4) $$p_0 = \frac{mr(1-a)^{n-1} + r(N-m)/n}{m + (N-m)/n}.$$

Equations 1-4 define the equilibrium strategies in terms of the exogenous parameters N, m, r, n and f. Sellers' expected profits and the expected price paid by each type of consumer can be calculated from these strategies. In equilibrium, each seller earns an expected profit of

(5) $$Ep = r\frac{N-m}{n} + \frac{f}{n-1}.$$

Loyal consumers expect to pay a price of

(6) $$\overline{p} = (1-a)r + a\int_0^r (1 - F(p))dp = \int_0^r (1 - aF(p))dp.$$

Bargain-hunters expect to pay a price of

(7) $$p_{\min} = (1-a)^n r + \sum_{k=1}^{n} \frac{n!}{k!(n-k)!}(1-a)^{n-k} a^k \int_0^r (1-F(p))^k dp = \int_0^r (1-aF(p))^n dp .$$

2.3 Comparative Statics

The effect of a change in the advertising fee is clear from equation (1):

$$\frac{da}{df} = \frac{-1}{f(n-1)}\left(\frac{nf}{(n-1)mr}\right)^{\frac{1}{n-1}} < 0 .$$

Thus, an increase in the advertising fee reduces the probability with which a seller advertises. Sellers that do not advertise should charge a price of r -- this is unaffected by any change in advertising fee. As the advertising fee changes the effect on advertised prices depends in a complicated way on the specific values of the parameters. For example, with two competitors the distribution of advertised prices shifts to the right when f increases, implying a higher expected advertised price:

$$\frac{dF(p)}{df} = \frac{n(1-aF(p))^{2-n}}{apm(n-1)}\left(pF(p)\left(\frac{1-aF(p)}{1-a}\right)^{n-2} - 1\right)$$

which is negative for n = 2. However, with a sufficiently large number of sellers the expected advertised price can decrease with f.

Changes in advertising fees have unambiguous implications for the expected payoffs of buyers and sellers. In identifying these, it is of course important to recognize that the final effect of a change in market conditions will depend on how sellers alter their behavior in response to such a change. For example, the direct effect of an increase in advertising costs is to lower profits: if sellers did not alter their advertising propensities or prices they would simply incur higher costs. But there is also an indirect effect to account for since in equilibrium sellers respond to an increase in advertising costs by advertising less frequently and posting higher prices. Indeed, in this model the indirect effect always outweighs the direct effect, since equilibrium profits are pinned down by the profits that can be obtained from charging a price of one and not advertising. From equation (5) it is clear that:

$$\frac{dEp}{df} = \frac{1}{n-1} > 0 .$$

Thus, costly advertising plays a facilitating role - a higher advertising fee deters sellers from competing with one another, and results in higher expected profits. In equilibrium, it is also straightforward to show that the expected price paid by loyal consumers increases with f

$$\frac{d\overline{p}}{df} = \int_{p_0}^{r} \frac{n}{(n-1)^2 mp} G(p)^{\frac{1}{n-1}-1} dp > 0 ,$$

as does the expected price paid by bargain-hunters:

$$\frac{dp_{\min}}{df} = \int_{p_0}^{1} \frac{n^2}{(n-1)^2 mp} G(p)^{\frac{n}{n-1}-1} dp > 0 .$$

In our experiment we manipulate the advertising fee across treatments in order to test these comparative static implications. With the higher fee advertising propensities are predicted to be lower, the expected prices paid by both types of consumer are predicted to be higher, and expected profits are predicted to be higher. Our experiment tests these predictions separately for two different values of *n*: *n* = 2 and *n* = 4.

This design also allows us to compare outcome variables across different numbers of competitors. However, in most cases comparative static predictions with respect to *n* are parameter specific. Advertising propensities are predicted to change with *n* according to

$$\frac{da}{dn} = \left(\frac{nf}{(n-1)mr}\right)^{\frac{1}{n-1}} \frac{1}{(n-1)^2} \left(\frac{1}{n} + \ln\left(\frac{nf}{(n-1)mr}\right)\right).$$

The last term in parentheses is decreasing in *n*, so if this term is negative when *n* = 2 it is negative for all *n*. Therefore, if $f < \frac{mr}{2} e^{-\frac{1}{2}}$, then advertising propensities decrease with *n*.

However, if $f > \frac{mr}{2} e^{-\frac{1}{2}}$ advertising propensities at first increase, and then decrease, with *n*.

Similarly, it can be shown that, depending on the specific parameter values, expected advertised prices can either increase or decrease with *n*. For example, when $f = 0$ the model reduces to the simpler model examined in Morgan et al., and so expected advertised prices (and therefore the expected price paid by uninformed consumers) increase with *n*. On the other hand, it can be shown that for other values of f an increase in the number of competitors decreases both expected advertised prices and the expected price paid by uninformed consumers (we give an example later). The effect on sellers' profits with respect to changes in the number of competing firms, is, however, unambiguous:

$$\frac{dEp}{dn} = -\frac{r(N-m)}{n^2} - \frac{f}{(n-1)^2} < 0 .$$

Thus the model allows us to predict the effect of a change in market parameters on equilibrium outcomes. In the case of changes in advertising fees the qualitative predictions are usually robust in the sense that they hold across the entire range of parameter values. In the case of changes in the number of competitors, the qualitative predictions are more delicate, and usually rely upon the precise values of market parameters.

3. Experimental Design and Procedures

The experiment consisted of 12 sessions conducted at the University of Nottingham during Fall 2001 and Spring 2002. Subjects were recruited from a distribution list comprised of undergraduate students from across the entire university who had indicated a willingness to be paid volunteers in decision-making experiments, where participants earn on average between £6 and £12 per hour, depending on the experiment. For this experiment subjects were sent an e-mail invitation promising to participate in a session lasting approximately 90 minutes, for which they would receive a £3 show-up fee plus “an additional amount that would depend on decisions made during a session.”[2]

Twelve subjects participated in each session, and no subject appeared in more than one session. Throughout the session, no communication between subjects was permitted, and all choices and information were transmitted via computer terminals. At the beginning of a session, the subjects were seated at computer terminals and given a set of instructions, which were then read aloud by the experimenter.[3]

The session then consisted of three phases of twenty periods each. At the beginning of each period, subjects were randomly assigned to groups of either two (Two-seller sessions) or four (Four-seller sessions) sellers, and then simultaneously chose prices from the set {0, 1, 2, …, 100} and whether or not to advertise their price.

Each group faced twenty-four computerized buyers who bought one unit each. Twelve of these buyers corresponded to bargain-hunters. These buyers were programmed to buy a unit from whichever seller advertised the lowest price. (In the case of ties for the lowest price, purchases were divided equally among the tied sellers; in the event that no seller advertised its price, purchases were evenly divided among all sellers in a group.) The remaining twelve buyers corresponded to loyal consumers. These buyers were evenly divided among all sellers in a group.

After all subjects had submitted their prices, profits for each seller, denominated in ‘points’, were calculated as (price × quantity) less any advertising fees incurred. At the end of each period a ‘Results Screen’ was displayed on each terminal. This screen listed all twelve prices and advertising decisions submitted in the period, together with the associated quantities, highlighting the prices and quantities for that subject’s competitor(s). The screen

---

[2] At the time of the experiment the exchange rate was approximately £1 = $1.50.
[3] Appendix A contains copies of the instructions.

also informed subjects of their own point earnings for that period, the previous five periods, as well as accumulated point earnings.

At the end of the session, subjects were paid the show-up fee plus 1p per 30 points accumulated over all sixty periods. Earnings averaged £18.97 (Two-seller sessions) and £8.93 (Four-seller sessions) for sessions lasting between 50 and 80 minutes.

In six of the sessions, which we refer to as our 'Hi-Lo-Hi' sessions, sellers faced a high cost of advertising (ϕ = 400) in the first phase of twenty periods, a low cost of advertising (ϕ = 200) in the second phase (periods 21-40), and a high cost (ϕ = 400) in the third phase (periods 41-60). This ordering was reversed in the other six sessions, where sellers faced low costs of advertising in Phases I and III and a high cost in Phase II. We refer to these sessions as 'Lo-Hi-Lo' sessions. Six sessions employed a Two-seller treatment, while the other six employed a Four-seller treatment. The design is summarized in Table 1 below.

Table 1. Summary of Experimental Design

	Two sellers	Four sellers
Hi-Lo-Hi	3 sessions	3 sessions
Lo-Hi-Lo	3 sessions	3 sessions

This design enables the effect of changes in advertising fees to be assessed by within-session comparisons. These comparisons can be made both for the case of n = 2, and for the case of n = 4. The effect of changes in the number of competitors can be also be assessed by making the relevant across-session comparisons. For the experimental parameters, the predictions for our various treatments are summarized in Table 2.

Table 2. Theoretical predictions				
	n = 2		n = 4	
	ϕ = 200	ϕ = 400	ϕ = 200	ϕ = 400
advertising propensity	0.67	0.33	0.39	0.24
expected advertised price	73	88	64	75
expected unadvertised price	100	100	100	100
expected profits	800	1000	367	433
expected price paid by loyals	82	96	86	94
expected price paid by bargain-hunters	73	93	63	82

4. Results

4.1 Overview

Table 3 presents summary statistics for the various treatment combinations, averaging over all relevant phases and sessions. A comparison of the entries in this table with those in Table 2 reveal that while the exception of the two-seller low advertising fee treatment there are some quite substantial discrepancies. Most noticeably, subjects advertise too frequently relative to the predictions.

In terms of the comparative static predictions, the decisions made by subjects vary across treatments in the predicted direction. Advertising propensities appear to decrease with both f and n, as predicted, while average advertised prices appear to increase with f and decrease with n, as predicted for our experimental parameters. The implications of these decisions for expected profits and the expected prices paid by each type of consumer, however, get rather mixed support.

Table 3. Results: Summary Statistics

	n=2		n=4	
	ϕ = 200	ϕ = 400	ϕ = 200	ϕ = 400
advertising propensity (%)	67	59	53	39
average advertised price	75	82	63	69
average unadvertised price	96	97	98	97
average profits	803	794	292	301
average price paid by loyals	82	89	80	86
average price paid by bargain-hunters	74	83	53	66

4.2 Testing Comparative Static Predictions

To assess the significance of these results we use non-parametric tests applied to *session-level* data. The advantage of this approach to analyzing the data is that it does not rely on any assumptions about the underlying data generation process within a session. We expect subjects to learn as they make decisions repeatedly, and we expect subjects to react to other subjects' decisions. Thus, while we have many observations per session, these observations should not be viewed as independent. On the other hand, we regard any summary statistic constructed from a single session to be independent from those constructed from other sessions. Thus, our approach yields exact tests without imposing strong assumptions about how subject choices are related to one another. Our null hypotheses state that changes in the market structure (i.e. in n or f) have no impact on decisions or market outcomes, and we test these against one-sided alternative hypotheses suggested by the model.

Our first result concerns subjects' decisions about whether to advertise or not:

Result 1. Sellers advertise less frequently when (i) the advertising cost is higher, and (ii) the number of sellers is higher.

Support for Result 1(i): Table 4 presents the proportion of decisions to advertise by phase for each session. In all twelve sessions, advertising propensities change from phase to phase in the direction predicted by the model - sellers advertise more frequently when the advertising fee is low. Under the null hypothesis the probability of observing advertising propensities move in the predicted direction between phases one and two twelve times out of twelve is 0.0002. Thus the null hypothesis that advertising cost does not affect advertising propensities can be rejected in favor of the directional alternative implied by the model, at any conventional significance level.

Table 4. Advertising Propensities

	Lo-Hi-Lo Sequence					Hi-Lo-Hi Sequence		
	Phase 1	Phase 2	Phase 3			Phase 1	Phase 2	Phase 3
Two	0.663	0.542	0.679		Two	0.671	0.792	0.752
Seller	0.646	0.513	0.596		Seller	0.579	0.675	0.638
Sessions	0.596	0.467	0.638		Sessions	0.533	0.754	0.617
Four	0.542	0.438	0.583		Four	0.396	0.558	0.388
Seller	0.538	0.433	0.550		Seller	0.392	0.554	0.354
Sessions	0.504	0.375	0.513		Sessions	0.367	0.433	0.325

Support for Result 1(ii): Table 4 also reveals that advertising propensities vary with the number of competing sellers in a systematic way. Consider phase one of the Lo-Hi-Lo sessions. The lowest advertising propensity of the three two-seller sessions is higher than the highest advertising propensity of the four-seller sessions. Under the null hypothesis such an extreme ordering would occur with probability 3!3!/6! = 0.05. Thus we can reject the hypothesis that advertising propensities do not vary with the number of competitors in favor of the alternative that advertising propensities fall with *n* for phase one of these sessions at a 5% significance level. In fact, this is also true for all the other phases, and for all the phases of the Hi-Lo-Hi sequence, as well.

Next we consider advertised prices:

Result 2. Sellers advertise higher prices when (i) the advertising cost is higher, and (ii) the number of sellers is lower.

Support for Result 2. As advertising cost is manipulated between phase 1 and 2, the average advertised price moves in the predicted direction in all but two cases (indicated with asterisks in Table 5), so that we can reject the hypothesis that behavior is invariant to f ($p = 0.019$). Between phases 2 and 3 the average advertised price moves in the predicted direction in every single case. Also, note that in all phases of both advertising sequences advertised prices fall as the number of competitors is increased, in line with the theoretical prediction.

Table 5. Average Advertised Prices

	Lo-Hi-Lo Sequence					Hi-Lo-Hi Sequence		
	Phase 1	Phase 2	Phase 3			Phase 1	Phase 2	Phase 3
Two Seller Sessions	68.84	83.11	76.60		Two Seller Sessions	81.82	76.65	82.71
	76.98	87.04	74.66			79.69	76.51	82.94
	68.15	86.69	81.07			78.16	76.75	80.19
Four Seller Sessions	56.78	72.30	65.44		Four Seller Sessions	65.27	63.49	70.32
	61.56	69.84	64.82			66.39*	69.15*	70.94
	58.18	71.14	65.89			59.81*	61.28*	70.92

Any seller that does not advertise should set a price of 100 in any treatment of our experiment. In fact a similar pattern is observed in all treatments - in the first few periods there are a few unadvertised prices below 100, but average unadvertised prices converge toward 100 as the session progresses. This is the only significant factor affecting average unadvertised prices:

Result 3. Average unadvertised prices do not vary systematically across treatments. In all treatments average unadvertised prices are somewhat lower than predicted (100) in phase one, but converge to within 1% of the predicted level by phase three.

Support for Result 3: We report average unadvertised prices in Table 6. There is a significant *n* effect in phases one and three of the Lo-Hi-Lo sessions (p = 0.10 for each case on the basis of separate two-sided tests). Against this however, one should note that there is no evidence of a significant *n* effect for phase two, or for any phase of the Hi-Lo-Hi sessions. Within session changes in unadvertised prices between phases one and two clearly indicate a learning effect - in 11 of 12 sessions unadvertised prices increase. However, this effect covers both transitions from low to high and from high to low advertising fees. A somewhat smaller effect is observed in the changes form phase two to three: in 9 of 12 sessions average unadvertised prices increase (and in one unadvertised prices remain at 100).

Table 6. Average Unadvertised Prices

	Lo-Hi-Lo Sequence					Hi-Lo-Hi Sequence		
	Phase 1	Phase 2	Phase 3			Phase 1	Phase 2	Phase 3
Two Seller Sessions	92.23	98.33	99.23		Two Seller Sessions	91.91	99.98	99.74
	87.68	98.68	99.54			92.13	99.04	99.44
	92.79	99.09	99.54			96.56	96.07	99.61
Four Seller Sessions	95.21	99.65	100.00		Four Seller Sessions	91.67	99.91	100.00
	95.14	98.64	100.00			92.25	100.00	100.00
	95.18	98.68	99.97			91.77	99.73	99.44

Taken together, these results indicate that while subjects decisions cannot be explained completely by the model, when the model predicts a change in subjects' decisions in response to changes in the market environment, the data exhibit a corresponding change.

Loyal consumers can expect to pay a price of 100 if their firm does not advertise, and a somewhat lower price if their firm attempts to attract bargain-hunters. Since firms are less likely to advertise when the advertising fee is high, and since advertised prices tend to be higher when the fee is high, we should expect from our earlier results that loyal consumers pay higher prices, on average in the f = 400 treatment. Indeed, this is the case:

Result 4. Average prices paid by loyal consumers increase with advertising fees.

Support for Result 4: As shown in Table 7, between phases one and two the average price paid by loyals moves in the predicted direction with respect to f in all but the two cases (indicated by asterisks) where advertised prices move in the wrong direction. Thus we can reject the hypothesis that f does not affect the expected price paid by loyals in favor of the directional alternative predicted by the model (p = 0.019). Between phases two and three the average price paid by loyals moves in the predicted direction in all twelve sessions.

Table 7. Average Prices paid by Loyals

	Lo-Hi-Lo Sequence					Hi-Lo-Hi Sequence		
	Phase 1	Phase 2	Phase 3			Phase 1	Phase 2	Phase 3
Two Seller Sessions	76.73	90.08	83.86		Two Seller Sessions	85.14	81.51	87.39
	80.77	92.71	84.71			84.93	83.83	88.92
	78.11	93.30	87.72			86.75	81.50	87.63
Four Seller Sessions	74.40	87.69	79.84		Four Seller Sessions	81.22	79.58	88.50
	77.09	86.16	80.65			82.12*	82.90*	89.71
	76.53	88.35	82.50			80.05*	83.07*	90.18

What about the effect of increasing *n*? As shown in Table 2, the model predicts a small increase in the expected price paid by loyals when the advertising fee is low, and a small decrease when the advertising fee is high. Here our results are quite mixed. In phase 2 of Lo-Hi-Lo and phase 1 of Hi-Lo-Hi an increase in the number of competitors leads to a significant decrease in the average price paid by loyals ($p = 0.05$), but in the remaining four cases we fail to reject the null against the one-sided alternative suggested by the model.

Referring again to Table 2, the expected price paid by bargain-hunters is predicted to be higher when (i) the advertising fee is higher, and (ii) the number of sellers is lower. These predictions are confirmed by our data:

Result 5. Average prices paid by bargain-hunters increase when (i) the advertising fee is higher, and (ii) when the number of competitors is lower.

Support for Result 5: As shown in Table 8, in every session the average price paid by bargain-hunters moves in the predicted direction as the advertising fee changes from phase to phase. Likewise, in every case, average prices paid by bargain-hunters are lower for $n = 4$ than for $n = 2$.

Table 8. Average Price paid by bargain hunters

	Lo-Hi-Lo Sequence					Hi-Lo-Hi Sequence		
	Phase 1	Phase 2	Phase 3			Phase 1	Phase 2	Phase 3
Two Seller Sessions	68.41	84.21	74.75		Two Seller Sessions	80.69	73.37	82.99
	73.66	88.64	76.27			78.60	77.82	83.49
	69.57	89.56	80.20			79.20	73.34	80.58
Four Seller Sessions	47.80	70.36	52.86		Four Seller Sessions	63.76	52.24	65.43
	50.94	58.75	50.93			61.85	59.61	69.49
	47.31	69.60	58.51			58.47	56.19	73.90

In order for sellers to benefit from higher advertising fees, it is not sufficient that they reduce their propensity to advertise and charge higher advertised prices. The magnitude of the shifts must also be sufficiently close to those prescribed by equilibrium analysis. Our final result summarizes the extent to which sellers are able to exploit the facilitating role of advertising costs.

Result 6. Average profits increase when there are fewer sellers, but are not systematically effected by advertising costs.

Support for Result 6. As seen in Table 9, for any phase and any advertising fee sequence, profits are significantly lower when four sellers compete. However, there is no clear pattern to how advertising fees affect profits. Comparing phases one and two, profits move in the predicted direction in six of twelve sessions (four of six two-seller sessions and two of six four-seller sessions). Comparing phases two and three, profits move in the predicted direction in eight of twelve sessions (three of six two-seller sessions and five of six four-seller sessions).

Table 9. Average Profit per Seller

	Lo-Hi-Lo Sequence					Hi-Lo-Hi Sequence		
	Phase 1	Phase 2	Phase 3			Phase 1	Phase 2	Phase 3
Two Seller Sessions	738.23	828.96	815.87		Two Seller Sessions	726.57*	770.87*	721.50*
	797.41	882.93	846.66			749.61*	834.90*	779.25*
	766.86	910.33	879.93			782.49	778.22*	762.48*
Four Seller Sessions	258.19	298.96	281.50		Four Seller Sessions	276.55*	283.86*	306.60
	276.48*	261.52*	284.75*			275.11*	316.73*	336.01
	270.73	323.85	320.44			268.77*	331.18*	362.23

5. Conclusion

As in previous experiments with (simpler) games involving mixed strategy equilibria, equilibrium does not describe the data particularly well. Nevertheless, at the level of subjects' decisions, the comparative static predictions based on equilibrium analysis are matched by the data. As advertising fees and numbers of competitors are manipulated, subjects' decisions change systematically in the direction predicted.

Note however that for some comparative static predictions of the model, concerning how a change in market environment affects sellers' profits, it is not sufficient that subjects' change their behavior in the correct direction - they must also change their behavior by the right magnitude. These comparative static predictions get at best mixed support from our experimental data.

In all our sessions sellers experienced two changes in advertising fees. In six sessions the first change induced a change in sellers' profits in the predicted direction, while in the other six sessions seller's profits moved in the opposite direction. In two-thirds of our sessions the second change in advertising fee induced a change in sellers' profits in the predicted direction.

One obvious question is whether there are alternative experimental environments that might deliver data that are closer to equilibrium play. For example, while our subjects advertised too frequently, there are clearly learning effects (for example in the development of unadvertised prices) and it may be the case that twenty periods in each phase is not sufficient for subjects to get a clear picture of what constitutes a good or bad decision. A design that enhances learning opportunities might well paint a different picture. However, while naturally occurring environments often involve ample opportunities to learn, it is rare to find cases where decision-makers face such a cleanly defined task and get to repeat it in a stationary environment. Finding the design parameters that are most conducive to equilibrium play might tell us little about whether equilibrium analysis is applicable to a naturally occurring market. A more promising avenue for further research may be to ask whether there are alternatives to equilibrium models, perhaps based on adaptive or boundedly rational behavior, that do a better job of organizing our data.

References

Baye, M. and Morgan, J. (2001). Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets. *American Economic Review*, **91**, 454-74.

Binmore, K., J. Swierzbinski and C. Proulx (2001). Does Minimax Work? An Experimental Study, *Economic Journal*, **111**, 445-464.

Brown-Kruse, J., S. Reynolds and V. Smith (1994). Bertrand-Edgeworth Competition in Experimental Markets, *Econometrica*, **62**, 343-371.

Cason, T. and Friedman, D. (2000). Buyer Search and Price Dispersion: A Laboratory Study. mimeo.

Morgan, J., H. Orzen and M. Sefton (2001). An Experimental Study of Price Dispersion. *mimeo*, http: //www.nottingham.ac.uk/economics/cedex/papers/mos2.pdf.

Varian, H. (1980). A Model of Sales. *American Economic Review*, **70**, 651-59.

Walker, M. and J. Wooders (2001). "Minimax Play at Wimbledon," *American Economic Review*, **91**, 1521-1538.

Appendix A - Instructions

General rules

This session is part of an experiment in the economics of decision making. If you follow the instructions carefully and make good decisions, you can earn a considerable amount of money. At the end of the session you will be paid, in private and in cash, an amount that will depend on your decisions.

There are twelve people in this room who are participating in this session. It is important that you do not talk to any of the other people in the room until the session is over.

The session will consist of 60 periods. You begin with an initial balance of 9,000 points, and in each period you can earn additional points (in some periods these may be negative). At the end of the experiment you will be paid based on your total point earnings, i.e. your initial balance plus points earned from all 60 periods. Points will be converted to cash using an exchange rate of 30 points = 1p. Notice that the more points you earn, the more cash you will receive at the end of the session.

Description of a period

Each person in the room has been designated as a seller. In each period {six} [three] markets will operate, and you will be randomly allocated to one of these. Similarly, the other sellers will be randomly allocated to markets. You will be competing with the [three] other {seller} [sellers] who {is} [are] randomly allocated to your market. Your point earnings will depend on the decisions in your market. Because sellers are randomly allocated to markets at the beginning of each period, the identity of your [three] competitor[s] will change from period to period.

At the beginning of each period you must decide what price to charge. You make your decision by entering a price (any whole number between 0 and 100) on your terminal. Then, after entering this price, you decide whether or not to pay a fee to advertise it.

After all sellers have made their decisions the computer will calculate your point earnings for the period. Your point earnings will be equal to the price you charge times the number of units you sell, minus the advertising cost:

$$\text{point earnings} = (\text{price} \times \text{number of units sold}) - \text{advertising cost}$$

If you decided not to advertise your price, you will incur zero advertising cost. If you decided to advertise your price you will incur an advertising cost, as explained later.

The number of units sold are calculated as follows:

(A) You automatically sell [3] {6} units whether you advertise or not.

(B) If no one in your market advertises, you sell an additional [3] {6} units.

(C) If you advertise and charge the lowest *advertised* price in your market, you sell an additional 12 units. (If you tie for the lowest advertised price in your market, the 12 extra units will be evenly divided between you and [the competitor(s) you are tied with] {your competitor}.)

At the end of each period the decisions of each seller will be displayed on your terminal. (You will be informed of the decisions of all sellers in all markets; the rows corresponding to your market will be highlighted.) The terminal will also display the number of units sold, and the point earnings, of each seller. Your terminal will also display your point earnings for that period, your point earnings from the previous five periods, and your accumulated point earnings.

Differences between periods

All periods are identical except that your competitors will be changing from period to period, and the advertising cost will be changing from phase to phase.

In Phase One (periods 1 to 20), the advertising cost is 400.

In Phase Two (periods 21 to 40), the advertising cost is 200.

In Phase Three, (periods 41 to 60), the advertising cost is 400.