CARESS Working Paper 97-12
Experimentation and Competition

Arthur Fishman
University of Haifa
Mt. Carmel
31905 Haifa,
ISRAEL
e-mail: afishman@post.tau.ac.il

Rafael Rob
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA. 19104-6297
USA
e-mail: rrob@econ.sas.upnn.edu

September 24, 1997

Abstract

This paper examines the effect of competition on firms efforts to experiment and learn about market demand. Consumers are assumed to know prices only at sellers they have actually visited, but must bear search costs to find the prices of other sellers. Under these conditions we show that firms incentives to experiment are diluted by comparison with the monopoly case and that this effect is stronger the smaller the search cost. The learning environment we portray gives rise to several time paths which have been empirically documented, including penetration pricing, cream-skimming and cyclical pricing.

1. Introduction

This paper studies dynamic equilibrium price formation when rms are imperfectly informed about demand but can increase their information by experimenting with prices. The optimal behavior of a single optimizing agent (a monopoly seller) in this setting is already well understood (e.g., Prescott [22], Grossman et. al. [16], McLennan [19], Easley and Kiefer [12], Aghion et. al. [1]): Experimentation can increase future expected pro ts by generating valuable information about market demand. On the other hand, by distorting the price away from the myopically optimal level, experimentation sacri ces current pro ts. The monopoly s optimal strategy must balance this current cost against the expected future pro ts from additional information.

The objective of this paper is to determine the effect of *competition* on the equilibrium level of experimentation and learning. We explore this issue in the context of search markets where rms uncertainty about demand is paralleled on the consumer side by uncertainty about prices. More speci cally, we consider markets in which consumers are only informed about the prices of stores they have actually visited but must bear search costs to verify prices at other stores. The cost of search is a natural measure of competition in this setting, the market being considered as more competitive the lower the cost of consumer search.

The main conclusion we reach under these circumstances is that competition leads to less experimentation and learning and to lower prices. This occurs because equilibrium prices diverge, which induces consumers search. This forces rms to lower their prices (so as not to lose sales), and this dilutes rms incentives to experiment in the rst place. We identify two channels through which this force operates.

The rst channel, termed the public good aspect of experimentation, is operative to the extent that information learned from experimentation is publicly observable. That promotes a tendency for only one rm at a time to experiment with relatively high prices while its competitors contemporaneously adopt a more conservative stance and stick with safer , lower prices (which consumers are more likely to accept), pending the revelation of new information. But this generates price dispersion between rms, encouraging consumer search. That, in turn, lowers the value of learning more about demand and leads to less experimentation, an effect which tends to be more pronounced the lower the cost of consumer search.

The second channel through which learning is reduced stems from the *opposite* consideration, namely, the extent to which rms are not informed about each others market experience - the private good aspect of experimentation. The more private information gleaned from experimentation is, the more likely it is that different rms will reach different conclusions about the underlying state of demand, even if their initial assessments were identical. Price dispersion occurs then as a result of different sellers reaching different conclusions about demand. As above, this increases the consumers incentive to search for a lower price in the (realistic) expectation that competing sellers may have formed more pessimistic beliefs about market demand and a consequent preference for lower prices. And, as above, increased search reduces experimentation and learning.

Thus, both the public good and the private good aspect of experimentation act to dampen learning and dilute market power. Rothschild [27] has referred to the possibility that a single optimizing agent may, as a result of a chain of unlucky experiences, incorrectly arrive at an unduly pessimistic conclusion about its environment, as the persistence of error. In our setting, competition compounds this phenomenon; the lower the cost of search, the less rms learn, and the greater the probability of settling into an erroneous belief. When the cost of search is small enough to eliminate experimentation altogether, the probability of error becomes a certainty.

Our work thus provides an important link between the experimentation and learning literature, on the one hand, and the equilibrium search literature on the other (e.g., Butters [8], Reinganum [23], Carlson and McAfee [10], Burdett and Judd [7], Rob [25], Fishman and Fershtman [13]). Of particular relevance is the comparison to Diamond s [11] result, that a small market friction on the demand side (a small search cost) leads to monopoly prices even if there are many sellers offering the same product. Our work shows that the addition of uncertainty on the supply side (rms uncertain about demand) serves to erode this monopoly power. The net result is that prices are more competitive once uncertainty is admitted on both sides of the market.

There are several other literatures to which this paper relates. One literature is the signal-jamming literature, which includes Riordan [24], Fudenberg and Tirole [15], Aghion et. al. [2], Mirman et. al. [20], [21]. In that literature rms learn about market demand (as they do in our framework) when each rm can affect through the choice of prices the extent of learning by the other. However, the signal-jamming literature takes the demand side as exogenous and given, whereas here consumers decide strategically where to buy, and rms pricing policies can affect consumers decisions.

Also related to this paper is the work on informational externalities. See, for example, Rob [26], Caplin and Leahy [9], Zeira [30], Harrington [17], Bolton and Harris [6] and Banerjee and Fudenberg [4]; and the related work on herd behavior and information cascades, for example Bikhchandani et. al. [3]. That literature emphasizes the private and social usefulness of observing others' actions (or market information, in general) and the resulting possibility of market failures (due to the fact that firms don't "pay" for the observations). While some of these ideas come into play here, the main difference is the endogeneity of consumer behavior, and the related effect of price dispersion.

Another literature to which this paper relates is the marketing literature on price dynamics of new products (see Mahajan et al [18] or Shapiro [28]). That literature classifies price paths into penetration pricing, cream-skimming and cyclical pricing. All these paths are feasible in our framework, although they arise here from the dynamics of learning and experimentation and not from price discrimination or introductory offers (as in the marketing literature).

The paper proceeds as follows. The following section analyzes a model of noisy demand. To unravel the public and private good effects of search on experimentation, we separately analyze equilibrium experimentation and learning under two opposing informational regimes. Under the first, firms are perfectly informed about each other's past prices and sales so that whatever is learned by one firm is publicly available to its competitors.. Under the second informational assumption, experimentation is a *fully* private good; each firm quotes prices secretly (e.g, law or consultancy firms) so that an individual firm knows neither its competitors' past sales nor past prices (though it perfectly knows its own price and sales history). We show that under both regimes, lower search costs lead to less learning by firms. In section 4 we analyze a special example of fully deterministic demand. In this example, firms reach a definite conclusion about demand in finite time, which subsequent market experience never reverses. The analysis of this example yields further results about the qualitative properties of the time-path of prices (e.g., whether they increase, decrease or cycle over time). Section 5 concludes.

2. The Model

The market opens at an infinite number of discrete time periods and includes two infinitely-lived, risk-neutral firms which produce a homogenous product at zero marginal cost, and compete in prices. One consumer arrives at each firm in each period. Consumers have either a low valuation or a high valuation for the product. A low-valuation consumer values a unit at u_L and a high-valuation consumer at

u_H, $u_H > u_L$. The probability that a randomly selected consumer has a high valuation is $\tilde{a}$. The firms know u_L and u_H but not the true value of $\tilde{a}$. Specifically, from the firms' perspective, $\tilde{a}$ is a random variable which can assume one of two values: a_L or a_H, $1 > a_H > a_L \geq 0$. We refer to a_H as the high-demand state and to a_L as the low-demand state. The realization a_i $(i = L, H)$ of $\tilde{a}$ remains constant through time. We assume the non-trivial case, $a_H u_H > u_L > a_L u_H$. This implies that u_H is the monopoly price under a_H while u_L is the monopoly price under a_L, so that learning about the true value of $\tilde{a}$ is potentially valuable.

Denote the firms' prior probability (in advance of any market participation) that $\tilde{a} = a_H$ by ρ_0, i.e., $\rho_0 \equiv \Pr(\tilde{a} = a_H)$, and let ρ be the subsequent posterior (which depends, of course, on subsequent information). Consumers are assumed to know both the state of demand (a_L or a_H) as well as their own realized willingness to pay (u_L or u_H).[1]

Upon entering the market, a consumer is randomly matched with one of the two firms, whose price she costlessly observes. She must, however, pay a search cost of $s > 0$ to learn the price of the other firm and buy from it. If she learns the price of both firms, she can buy from either one (i.e., search is with recall). If she observes the same price at each firm, one firm is chosen randomly (and with equal probabilities). A consumer knows the calendar date – the number of periods which have elapsed since the market opened – and the firms' prior probability, ρ_0, but not the current posterior probability of any firm (nor its price or sales history).

It is well known that if firms are perfectly informed about demand, the fact that all consumers have positive search costs, however small, implies that the unique equilibrium price for the duopoly is the monopoly price (Diamond, 1971). This will not be the case for our imperfectly informed firms, as is shown below.

3. Analysis

3.1. Monopoly Experimentation

The analysis of the monopoly problem provides a convenient reference point. If the monopolist sets his price equal to u_L no new information is generated, since that price is accepted by both low and high-valuation consumers. By contrast, if the price is above u_L, a sale increases the likelihood of a_H while failure to sell increases

[1] This suppresses the (seemingly intractable) case where consumers indirectly learn from sellers' prices about other consumers' demands and, therefore, about what other sellers may have learned.

the likelihood of a_L. Therefore charging more than u_L is called experimentation . An experiment is called a success if it ends in a sale and a failure otherwise. Given the prior, ρ, a sale at a price greater than u_L updates the belief in a_H to:

$$R_Y(\rho) = \frac{\rho a_H}{\rho a_H + (1-\rho)a_L}, \tag{3.1}$$

while failure to sell at that price updates the belief in a_H to:

$$R_N(\rho) = \frac{\rho(1-a_H)}{\rho(1-a_H) + (1-\rho)(1-a_L)}, \tag{3.2}$$

where Y and N correspond to the consumer buying and not buying, respectively. Clearly, $R_Y(\rho) > \rho > R_N(\rho)$.

Let $V(\rho)$ be the value function the monopoly s expected discounted pro t from following the optimal pricing rule when its current belief in a_H is ρ. As noted above, u_L cannot provide information. Therefore, if it is ever optimal to stop experimenting i.e., to charge u_L then u_L remains optimal at every future date. Thus the value of charging u_L is $\frac{u_L}{1-\delta}$. On the other hand, if the monopoly experiments, the continuation value depends on whether the experiment is a success or a failure. Thus

$$V(\rho) = Max\{\frac{u_L}{1-\delta}, \widetilde{\rho}[u_H + \delta V(R_Y(\rho))] + (1-\widetilde{\rho})\delta V(R_N(\rho))\}, \tag{3.3}$$

where $\delta \in (0,1)$ is a discount factor and $\widetilde{\rho} \equiv \rho a_H + (1-\rho)a_L$.
The following is a well known result. (See Berry and Friedstedt 1985, p. 38)

Proposition 0: The optimal pricing rule for the monopoly is characterized by a cutoff value -call it ρ^*- so that u_H is optimal $V(\rho) = \widetilde{\rho}[u_H + \delta V(R_Y(\rho))] + (1 - \widetilde{\rho})\delta V(R_N(\rho))-$ if $\rho > \rho^*$, and u_L is optimal $V(\rho) = \frac{u_L}{1-\delta}$ – if $\rho < \rho^*$.

To avoid the trivial cases we assume that $1 > \rho_0 > \rho^*$. Since $\rho_0 > \rho^*$, it is optimal for the monopolist to experiment at least once and, since $1 > \rho_0$, it is optimal to stop experimenting if the evidence that comes in is sufficiently unfavorable.

The determination of ρ^* re ects not only the impact of pricing on the current payoff but also its potential effect on future pro ts. More speci cally, let

$$\rho^m \equiv \frac{(u_L/u_H) - a_L}{a_H - a_L}. \tag{3.4}$$

ρ^m is the *myopically* optimal cutoff point: u_L (u_H) maximizes the one-period expected pro t if $\rho < \rho^m$ ($\rho > \rho^m$). ρ^m ignores, however, the impact of experimentation on *future* pro t. As a consequence, $\rho^* < \rho^m$, implying that when $\rho \in (\rho^*, \rho^m)$ the monopolist charges u_H although u_L is myopically optimal.

3.2. Duopoly Experimentation when Information is a Private Good

We now turn to the analysis of duopoly pricing and experimentation. As promised in the introduction, we shall analyze the duopoly scenario under two different information regimes. In this subsection it is assumed that whatever a rm learns about demand is its private information. More speci cally, it is assumed that prices are quoted secretly to consumers. Thus each rm is fully informed about its own price and sales history but is ignorant of the price and sales history of its competitor.

This assumption may seem extreme and in some contexts it is indeed not the most realistic. There are, however, many instances in which this assumption *is* realistic. One instance is where the product is customized to consumers needs. A classical example is where the product is a turbo-generator and the customer is an electric utility. In this instance, the product and its price are adjusted to the buyer s speci c needs, which makes it difficult to ascertain (or interpret) prices at rival rms.[2] Another instance is when the product is a service, e.g., legal or consultancy services. Here, again, the service and the price vary from customer to customer (either because of the nature of the service or because of the customer s willingness to pay), and are not public information. A third instance is where the rm produces a whole variety of products, which makes it difficult to assemble detailed information about individual prices. For example, AT&T recently went to Federal court to demand that its chief rivals in the long distance business disclose their prices in public lings with the government. According to the New York Times (Feb 11, 1993, page D5), Information about prices charged by competitors can be invaluable in any eld. But it is particularly so in telecommunications where the array of services purchased by big corporations come in so many different combinations that it is often difficult for a company to discern the pricing strategy of its rivals without hard information. The same article quotes the president of AT&T s Business Communications services division as saying I

[2] This case has been extensively discussed in the oligopoly literature (see Stigler (1964)), asserting that the difficulty of obtaining reliable price information would lead to the instability of collusive arrangements.

have to spend an enormous amount of money to nd out what s happening in the market place . In accordance with these examples, we now assume that prices and sales histories are private information.

We con ne attention to pure strategy Markov equilibria (**PSME**) - equilibria in which (i) Prices are chosen deterministically and (ii) Prices depend only on variables which affect pro ts directly, and not on irrelevant past history[3]. In our setting, these variables are the rm s current belief about demand, as summarized by its posterior probability, and the consumers search behavior.

To appreciate the effect of demand uncertainty on competition, consider the benchmark case in which rms are perfectly informed (i.e., they know $\widetilde{a}$ with certainty). As remarked above, for any s, the unique equilibrium price for a duopoly rm is then identical with that of the monopoly: u_H if $\rho > \rho^*$ and u_L if $\rho < \rho^*$. Since the rms uncertainty about demand applies only to prices above u_L, that logic directly implies:

Lemma 1: The duopoly price is never less than u_L.

An immediate consequence of lemma 1 is that low-valuation consumers who are charged more than u_L leave the market without search: Since no rm charges less than u_L, the expected return from costly search can only be negative.

As it turns out, however, demand uncertainty *does* affect the prices with which the duopoly *experiments*. In fact, it will turn out that, depending on the cost of search, a duopolist may be unable to charge u_H no matter how high a probability it assigns to a_H. The reason is that when a rm as a result of a favorable market history charges more than u_L, there is every possibility that its competitor, having experienced a less favorable history, has formed a more pessimistic belief and is simultaneously charging u_L. This possibility creates price dispersion in the market and provides the high-valuation consumers an incentive to search, in proportion to the search cost, s. We refer to this incentive as the search effect. If the search effect is sufficiently strong (i.e., if s is sufficiently small), even high-valuation customers will not accept u_H without search. Accordingly, we denote by $\overline{p_t} \leq u_H$ the highest acceptable price to a high-valuation consumer at period t. Our main result below is that for sufficiently low s and for a sufficiently large t, $\overline{p_t} < u_H$.

In analyzing this case note that high-valuation consumers incentives to search

[3]The history *is* relevant if rms are trying collude by means of punishment strategies which depend on past behavior. However, our focus here is on experimentation and not on collusion.

may depend on what they know about the true state of demand: If consumers know that $\widetilde{a} = a_L$, their inducement to search is stronger than under $\widetilde{a} = a_H$ since they consider it more likely that an unsampled store has already reached a pessimistic conclusion and is consequently charging a low price. Therefore, the pricing policies of rms will have to account for potentially divergent search behavior by different high-valuation consumer types (here types means high-valuation consumer who know a_L or a_H). The next lemma deals with this issue.

Lemma 2: In any PSME for the duopoly, at any period, the price charged by rms is such that high-valuation consumers (u_H) of either type (a_L or a_H) do not search.

Proof: We argue by contradiction. Assume that some equilibrium price, p_t, triggers search. Let π_t be the probability that a rm s price at period t is u_L and let 1-π_t be the probability that its price is p_t. (This is the belief held by one rm about the other, given its own sales history). Since a rm can be visited by an uninterrupted sequence of t u_H-customers with positive probability (at least ρa_H^t), we must have $\pi_t < 1$. Then a rm which charges p_t makes the sale to a high valuation consumer with probability 0.5(1-π_t) and loses the sale with probability 1-0.5(1-π_t) (if buyers observe the same price at both rms they choose one at random). By reducing its price to p_t - ε, it makes a sale with probability 1 when its rival charges p_t. Thus the deviation changes its expected pro t from 0.5(1-π_t)p_t to (1-π_t)(p_t -ε), which represents an increase if ε is sufficiently small (here we use the fact that $\pi_t < 1$). When the consumer does not show up or his valuation is u_L, the two prices yield equal revenue (zero). Hence, $p_t - \varepsilon$ gives an expected pro t higher than p_t.

Furthermore, ε can be chosen sufficiently small that any consumer who was searching under p_t continues to search under p_t-ε. The reason is that the reservation price of such consumer must have been strictly below p_t which means that an intermediate price between p_t and the reservation price can be found. Since consumers continue to search under $p_t - \varepsilon$, the deviating duopolist gains revenue and does not lose information. So he is better off with $p_t - \varepsilon$, which means that p_t was not an equilibrium strategy in the rst place. ■

This pooling result establishes that any price which is charged in equilibrium is accepted by high-valuation consumers (whether $\widetilde{a} = a_L$ or $\widetilde{a} = a_H$). A fortunate consequence of this fact is that the updating of information along the duopoly path is the same as along the monopoly path for prices above u_L. The duopoly and monopoly paths may differ only with respect to the prices above u_L that may be charged, and in the consequent incentive to experiment. Therefore we con ne attention to equilibria with the following properties.

A PSME for the duopoly consists of three sequences, $\{\overline{p}_t, q_t, \rho_t^*\}_{t=1}^{\infty}$, where:

(i) $\overline{p}_t$ is the highest price which high-valuation consumers accept without search.

(ii) ρ_t^* is the cutoff posterior probability so that rms price policies are:

$$p_t = \begin{cases} \overline{p_t} & if\ \rho_{t-1} > \rho_t^* \\ u_L & \rho_{t-1} < \rho_t^*, \end{cases}$$

where ρ_{t-1} is the posterior after a $t-1$ period long sales history (ρ_t is, of course, a random variable at any $\tau < t$),

(iii) $q_t =$ Prob $(\widetilde{\rho}_{\tau-1} < \rho_\tau^*$ for some $\tau \leq t \mid \widetilde{a} = a_L$), i.e., the probability that a rm has reached a pessimistic conclusion at any time before t conditional on $\widetilde{a} = a_L$. Since q_t covers *all* periods up to t, the sequence $(q_t)_{t=1}^{\infty}$ is increasing.

Such that:

(E.1) Consumers search optimally[4]:

$$\overline{p}_t = \min(u_H, u_L + \frac{s}{q_t}). \tag{3.5}$$

(E.2) Firms maximize their expected pro ts:

$$V(\rho, t; \overline{p}) = Max\{\frac{u_L}{1-\delta}, \widetilde{\rho}[\overline{p}_t + \delta V(R_Y(\rho), t+1; \overline{p})] + (1-\widetilde{\rho})\delta V(R_N(\rho), t+1; \overline{p})\}, \tag{3.6}$$

where $V(\rho, t; \overline{p})$ is the duopolist s value function[5], $\overline{p} = (\overline{p}_1, \overline{p}_2, ...)$ is the sequence of maximum prices and t is calendar time. The maximum above is attained by the rst (second) term when $\rho < \rho_t^* (\rho > \rho_t^*)$.

We can now state and prove:

[4] Given that each rm s price is u_L with probability q_t, the price which makes the consumer indifferent between accepting p_t and searching is $u_L + \frac{s}{q_t}$.

[5] We abuse notation here since V was previously used to denote the *monopoly* value function

Proposition 1: There exists an $s^* > 0$, so that for $s < s^*$: (i) The maximum price along *any* PSME duopoly path is lower than u_H from some date T and onwards; $\overline{p}_t < u_H$ for every $t \geq$ some T. (ii) The cutoff posterior probability, above which the duopolist experiments and below which it charges u_L, is higher than that of the monopoly, at each period: $\rho_t^* > \rho^*,\ t \geq 1$; thus, a dupolist experiments less than a monopolist.

Proof: Let $R(\rho, r, t)$, $\rho \geq \rho^*$, be the posterior on a_H conditional on r successes among t experiments ($0 \leq r \leq t$). Then:

$$R(\rho, r, t) = \frac{\rho a_H^r (1-a_H)^{t-r}}{\rho a_H^r (1-a_H)^{t-r} + (1-\rho) a_L^r (1-a_L)^{t-r}}.$$

Now for any t there exists an integer $r(t)$, $0 \leq r(t) \leq t$, so that

$R(\rho, r, t) \geq \rho^*$ for $r \geq r(t)$, and $R(\rho, r, t) < \rho^*$ for $r < r(t)$.

(When $r(t) = 0$, the posterior is above ρ^* even under the worst possible evidence). Note that $R(\rho, r, t)$ is diminishing in t and, thus, that $r(t)$ is increasing in t. Let q_t^m be the probability that the *monopolistic* posterior has eroded below ρ^* at some time prior to t. Then

$$q_1^m = 0, \text{ since we start with } \rho_0 > \rho^*, \text{ and}$$

$$q_{t+1}^m = \begin{cases} q_t^m & \text{if } r(t+1) = r(t) \\ q_t^m + \sum_{k=r(t)}^{r(t+1)} \begin{pmatrix} t \\ k \end{pmatrix} a_L^k (1-a_L)^{t+1-k} & \text{if } r(t+1) > r(t). \end{cases}$$

Since $(q_t^m)_{t=1}^\infty$ is an increasing sequence of numbers between zero and one, it converges to a limit. Denote this limit by q^* ($0 < q^* \leq 1$), and de ne

$$s^* = q^*(u_H - u_L).$$

Then for any $s < s^*$ there exists a T so that:

$$s < q_t^m (u_H - u_L),\ t \geq T. \tag{3.7}$$

Consider now the duopoly equilibrium under $s < s^*$. Since the high price under duopoly is never higher than u_H, the incentive to experiment is certainly not stronger, i.e., $\rho_t^* \geq \rho^*$. Thus the probability of eroding below the cutoff point

ρ_t^* is no smaller than q_t^m. Then by (3.5), $\overline{p}_t \leq u_L + \frac{s}{q_t^m}$, and by (3.7) the RHS is $<u_H$. Hence $\overline{p}_t < u_H$ for $t \geq T$. So according to Lemma 2, the duopoly price in periods $t \geq T$ must be no higher than $\overline{p}_t$ (regardless of what evidence comes in). Furthermore, since the duopolistic updating is identical to the monopolistic updating, the only difference between the monopoly value function, (3.3), and the duopoly value function, (3.6), is in the prices charged. This implies that the duopoly value function is smaller, the difference re ected in the experimentation term. Therefore, $\rho_t^* > \rho^*$, for $t \geq T$.

But the fact that $V(\rho, t; \overline{p})$ is smaller than $V(\rho)$ for $t \geq T$ implies (through its de nition) that it is also smaller for $t < T$, which implies that $\rho_t^* > \rho^*$ for $t < T$. ■

The idea of the preceding proof is that once enough time has elapsed since the market s inception, the probability that an unsampled rm s belief in high demand has eroded below the cutoff level is sufficiently high to make search attractive if search costs are low enough. Then the credible threat of search restricts experimentation to prices below u_H. This in turn reduces the pro tability of experimentation, which is why $\rho_t^* < \rho^*$ from the start, leading to less learning and to a greater probability of error.

Rothschild [27] has argued that a single optimizing agent may permanently settle into an erroneously pessimistic belief, as a consequence of a chain of bad luck, even though it could get closer to the true state by experimenting a sufficiently large number of times, and has termed this phenomenon the persistence of error . It occurs in our model under a monopoly which reaches the erroneous conclusion that demand is low when it is actually high. Our analysis shows that competition compounds this phenomenon; the probability that the duopoly erroneously comes to believe that demand is low when it is actually high is greater the lower the cost of consumer search.

3.3. Duopoly Experimentation when Information is a Public Good

This subsection analyzes the duopoly equilibrium under the opposite information regime, when all information learned from experimentation is public: Each rm is perfectly and costlessly informed about its competitor s price and sales volume at the preceding period. Thus, the rms have the same belief, ρ, at every period[6].

[6] However, we shall continue to assume that consumers know only calendar time and the prior belief, ρ_0, but not the rms market history (nor their current belief at any period save

Despite the radically different environment (information being public rather than private), proposition 1 holds here as well.

Proposition 2: Suppose experimentation is a public good. Then, for any $s < \overline{s} \equiv u_H - u_L$, there exists a range (ρ^*, ρ^{**}) so that the monopolist experiments for any ρ_0 in this range whereas both duopolists charge u_L.
Proof: We rst show that there exists a $\rho^{***} > \rho^*$, such that if $\rho < \rho^{***}$, at least one rm charges u_L with probability 1, for *any* s. Suppose rm 1 s strategy is to experiment with u_H. Let $V_L(\rho)$ $(V_H(\rho))$ be the value to rm 2 from currently charging u_L (u_H) and then behaving optimally, and let $V = max\{V_L, V_H\}$[7]. Then:

$$V(\rho) = Max\{V_L(\rho), V_H(\rho)\}, \tag{3.8}$$

$$V_L(\rho) = u_L + \delta\{[\widetilde{\rho}V(R_Y(\rho)) + (1-\widetilde{\rho})V(R_N(\rho))]\}, \tag{3.9}$$

$$\begin{aligned} V_H(\rho) = {} & \widetilde{\rho}[u_H + \delta(\widetilde{\rho}\, V(R_{YY}(\rho)) + (1-\widetilde{\rho})V(R_{YN}(\rho)))] \\ & +(1-\widetilde{\rho})\delta[\widetilde{\rho}\, V(R_{NY}((\rho)) + (1-\widetilde{\rho})V(R_{NN}(\rho))], \end{aligned} \tag{3.10}$$

where $R_{ij}(\rho)$ is the posterior on ρ following 2 experiments resulting in i, j, where $i, j = N, Y$.

Consider $\rho = \rho^*$. Then, given that rm 1 experiments, non experimentation is *strictly* more pro table for rm 2 than experimentation; $V_L(\rho^*) > V_H(\rho^*)$. This is because u_L is myopically optimal, and from the perspective of learning, rm 2 is better off to proceed sequentially and use the information it gets for free from rm 1: In the event that rm 1 succeeds, rm 2 can always experiment in the next period; in the event that rm 1 fails, neither will want to experiment any more. Thus, if rm 1 experiments and $\rho = \rho^*$, rm 2 s best response to rm 1 s strategy is to charge u_L, even if $s > s^*$.[8]

Since the solution to the system (3.8) (3.10) is continuous in ρ, there also exists a $\rho^{***} > \rho^*$ such that $V_L(\rho) > V_H(\rho)$ for $\rho < \rho^{***}$. Thus for $\rho < \rho^{***}$, one rm at most experiments with a price greater than u_L, no matter how large s is.

Suppose $\rho_0 \in (\rho^*, \rho^{***})$. Then at least one rm, call it rm 2, charges u_L. It is easy to verify that Lemma 2 applies to the public-information regime as well.

the rst).

[7] We abuse notation again since V was previously used to denote the value function of a monopoly.

[8] This reasoning is predicated on the experimental price being u_H. However, even if the experimental price is lower, this can only *reduce* the incentive to experiment. Therefore, the result that only one rm experiments holds for *any* experimental price.

Therefore, if $s < \overline{s}$, the experimental price can be no higher than $u_L + s$, because a u_H-consumer is certain to find the price u_L at the store he has not sampled as yet and, according to Lemma 2, the experimental price must be acceptable to all u_H-consumers. Consequently, the incentive of firm 1 to experiment is weaker than the monopolistic incentive (a monopolist can charge u_H without having to worry about search). This implies the existence of a ρ^{**}, $\rho^{***} \geq \rho^{**} > \rho^*$ so that firm 1 does not experiment in the range (ρ^*, ρ^{**}). Therefore, neither duopolist experiments in this range while a monopolist does. ■

Thus, despite the radically different environment, the search effect, and its attendant reduction of market power and learning, identified in the private information regime, applies under the public information regime as well. The logic is quite different, however. In the private information regime, the search effect results from the tendency of firms to form different beliefs about demand, as a reflection of their diverse experience in the market place. In the public information regime, on the contrary, firms always hold identical beliefs. As the proof of Proposition 2 shows, it is precisely this feature which, for a range of beliefs greater than ρ^*, allows one firm to free ride on the experimentation of its competitor, promoting a tendency away from *simultaneous* towards unilateral experimentation even though all information is shared by firms. Again, this has the effect of increasing price dispersion, which limits the extent to which the duopoly experiments, relative to a monopoly, *even when unilateral experimentation is individually optimal for each firm*. But this limits the market power of firms and reduces their motivation to learn about demand.

4. An Example: Deterministic Demand

This section analyzes a special example which illustrates the principles developed in the preceding sections and is simple enough to yield a more precise characterization of the solution. The solution adds further insights about the qualitative nature of the time-path of prices, and how it relates to the information structure; namely, it shows whether prices increase, decrease or fluctuate over time. In the example, the uncertainty is about demand at high prices, so that search reduces market power and experimentation (same as in section 3).

In this example, u is identical for all consumers and constant over time. Consumers know the value of u but firms consider u to be a random variable whose value is either $u = u_L > 0$ or $u = u_H > u_L$. We say that demand is high (low) if

$u = u_H$ ($u = u_L$). The firms initially assign the prior probability of ρ_0 to high demand.

If this were all, a single trial with a price greater than u_L would resolve all uncertainty: A sale would prove that $u = u_H$, and failure to sell that $u = u_L$. To avoid such a trivial solution, it is assumed, in addition, that demand fluctuates randomly between "boom" and "slump" periods. Consumers arrive in the market only in boom periods; in a slump period, no consumers enter the market. A boom occurs with probability $\beta, \beta \in (0, 1)$, and the occurrence of booms is independent over time. At any period, a firm can only observe the number of sales it made, but not how many consumers entered the market at that period. For example, the firm may sell its merchandise via a mail catalogue so it cannot observe the "physical" number of shoppers. Alternatively, even if firms can observe whether consumers physically entered their stores, they may be unable to distinguish between "serious" customers and window shoppers. Thus failure to sell at a price greater than u_L, while increasing the likelihood that demand is low, is inconclusive because it can be attributed either to *permanently* low demand ($u = u_L$) or to a temporary slump (in which case it is not possible to sell at *any* price). On the other hand, a sale at a price greater than u_L proves that $u = u_H$. As before, u_L is uninformative and experimentation consists of charging more than u_L.

4.1. Monopoly Regime

Again, we first consider the problem of a monopoly which is visited by a single customer in each period with probability β.

Following each failure to sell at a price greater than u_L, the belief in high demand is revised downwards to $R(\rho)$, where:

$$R(\rho) \equiv \rho(1-\beta)/[\rho(1-\beta) + (1-\rho)].$$

Therefore, the monopoly's discounted profit when it charges u_H and the prior is ρ is

$$\rho\beta[u_H + \delta\beta u_H/(1-\delta)] + (1-\rho\beta)\delta V(R(\rho)). \tag{4.1}$$

The discounted profit from charging u_L (forever) is

$$\frac{\beta u_L}{1-\delta}. \tag{4.2}$$

Again the monopoly's optimal pricing policy is characterized by a cutoff belief ρ^* such that its price is u_H whenever $\rho > \rho^*$ and u_L if $\rho < \rho^*$. For the deterministic demand example, ρ^* can be explicitly derived.

Proposition 3: The cutoff belief for the monopolist in the deterministic demand case is:

$$\rho^* = \frac{(1-\delta)u_L}{[(1-\delta)u_H + \delta\beta(u_H - u_L)]}. \tag{4.3}$$

Proof: The high price is charged until the value from charging a low price exceeds the value from charging the high price, i.e., until (4.1) is less than (4.2). Let ρ^* be the belief at which the monopolist is indifferent between the two. Then, $V(R(\rho^*)) = \beta u_L/(1-\delta)$ (since $R(\rho^*)<\rho^*$, the monopolist strictly prefers to charge u_L at $R(\rho^*)$). Substituting this into (4.1), equating with (4.2) and solving for ρ^* gives the expression in (4.3). ■

A useful, alternative characterization of the monopoly optimum is as follows. There exists an integer $M, M \geq 1$, such that u_H is charged for the rst M periods with probability 1. This is the experimentation period . If at least one success occurs during the rst M periods, the high price continues to be charged forever. Otherwise, by the end of the M-th period, the posterior erodes below ρ^*, experimentation stops, and the price settles down to u_L from period $M+1$ and onwards. The length of the experimentation period can be determined as follows. Let $R^t(\rho)$ be the t-th iterate of $R(\rho)$:

$$R^t(\rho) \equiv \rho(1-\beta)^t/[\rho(1-\beta)^t + (1-\rho)].$$

$R^t(\rho_0)$ is the posterior at the beginning of period $t+1$ following t successive failures. Since $R^t(\cdot)$ decreases monotonically to zero as $t \to \infty$, there must be an M with the property that $R^t(\rho_0) \geq \rho^*$ for $t \leq M$ and $R^{M+1}(\rho_0) < \rho^*$. This M de nes the experimentation period.

4.2. Duopoly Regime, Private Good Case

We now turn to the analysis of the duopoly under the private information regime: Each rm is fully informed about its own price and sales history but does not know the price and sales history of its competitor; consumers do not know the price histories of either rm. Here we assume that the consumer arrival probabilities at the two rms are statistically independent. That is, a consumer arrives at rm i with probability β which is independent of whether one arrives at rm j. This assumption is different from the assumption in section 3 that one consumer visits each rm in each period. However, under the previous assumption (and with deterministic demand), the rms would have the same information at all dates (since all consumers are identical) and, consequently, there would never be private information about demand. So the assumption of independent arrivals is needed to retain the idea of divergent beliefs and divergent pricing.

Obviously Lemma 1 applies to this setting. It is veri able that Lemma 2 applies as well. Thus at any period there are two possible prices, u_L and the experimental price p_t, where $u_H \geq p_t \geq u_L$.

Clearly, like a monopoly, a duopolist which has experienced M successive failures charges u_L from then on (where M is derived in section 4.1). Thus if demand is actually high, the probability that rm i s price is u_L at period $M+1$ is at least (1-β)^M, the probability of M successive failures (If the duopolist stops experimenting before experiencing M failures, the probability is higher yet). Thus the expected return from search to a high valuation consumer who is charged $p > u_L$ at period $M+1$ is at least (1-β)$^M[(u_H - u_L) - (u_H - p)] =$ (1-β)$^M(p - u_L)$. The consumer, therefore, optimally rejects p to search if: (1-β)$^M(p - u_L) \geq s$. Let $\overline{p}_{M+1}$ satisfy the above weak inequality with equality: (1-β)$^M(\overline{p}_{M+1} - u_L) = s$. This gives:

$$\overline{p}_{M+1} = u_L + s/(1-\beta)^M.$$

Then any price above $\overline{p}_{M+1}$ triggers search. In other words, the *effective* reservation price of high-valuation buyers is at most $\min(\overline{p}_{M+1}, u_H)$. Thus, by Lemma 2, the high price at period $M+1$ cannot be above $\overline{p}_{M+1}$. Let $\widetilde{s}$ be the search cost for which $\overline{p}_{M+1} = u_H$, i.e.,

$$\widetilde{s} = (u_H - u_L)(1-\beta)^M.$$

If $s \geq \widetilde{s}$, $\overline{p}_{M+1} \geq u_H$. In that case the duopolist is effectively a monopolist and prices accordingly. If, however, $s < \widetilde{s}$, then $\overline{p}_{M+1} < u_H$. Similarly, $\overline{p}_t < u_H$ for all $t > M+1$. Thus, if $s < \widetilde{s}$, the highest price which can be charged from period $M+1$ and onwards is $\overline{p}_{M+1}$ *even if the duopolist is certain that demand is high.*

But that is not all. The fact that market power is diminished in future periods diminishes the incentive to experiment in earlier periods, which further erodes market power. The net result is stated in the following proposition.

Proposition 4: Suppose experimentation is a private good. Let $\widetilde{s} = (u_H -$

$u_L)(1-\beta)^M$. If $s \geq \widetilde{s}$, the equilibrium price path of the duopoly is identical to that of the monopoly. If $s < \widetilde{s}$, the duopoly price path is given as follows. Define:

$$\overline{p}_t \equiv u_L + \frac{s}{(1-\beta)^{t-1}} \quad \text{for} \quad t \leq M+1, \tag{4.4}$$

and

$$W_t = \rho_{t-1}\beta[u_H + \delta\frac{\beta\overline{p}_{t+1}}{1-\delta}] + (1-\rho_{t-1}\beta)\delta\frac{\beta u_L}{1-\delta}. \tag{4.5}$$

Let $0 \leq N \leq M$ be the largest integer for which $W_N > \dfrac{\beta u_L}{1-\delta}$. Then the equilibrium duopoly price path is to charge u_H for the first N periods. If none of the first N experiments is a success, the firm charges u_L forever after. If one of the first N experiments is a success, the firm charges $\overline{p}_{N+1}$ in period $N+1$ and onwards. If for every $t \geq 1, W_t < \dfrac{\beta u_L}{1-\delta}$, the duopoly charges u_L from the very first period (i.e., never experiments).

Proof: As established above, the highest possible price at $t \geq M+1$ is $\overline{p}_{M+1}$. Therefore the duopoly value function at the beginning of period M is

$$W_M = Max\{\frac{\beta u_L}{1-\delta}, \rho_{M-1}\beta[u_H + \delta\frac{\beta\overline{p}_{M+1}}{1-\delta}] + (1-\rho_{M-1}\beta)\delta\frac{\beta u_L}{1-\delta}\}.$$

$W_M < V(\rho_{M-1})$ because the termination value upon success is only $\dfrac{\beta\overline{p}_{M+1}}{1-\delta}$, and not $\dfrac{\beta u_H}{1-\delta}$ as is the case under monopoly. The fact that the continuation value is diminished, relative to that of a monopoly, raises the attractiveness of switching to a low price after only $M-1$ failures, i.e., already at period M. Whether a duopolist will indeed switch over at the beginning of period M depends on the values of $\rho_{M-1}$and s; the smaller is ρ_{M-1}(which occurs when ρ_0 is small) and the smaller is s, the stronger the incentive to switch over earlier. The intuition is that a smaller ρ_{M-1} corresponds to a more pessimistic outlook on consumers' valuation, whereas a smaller s raises the degree of competition in the market by increasing consumers' incentive to search for lower prices. Either way, firms' market power is weakened, leading to lower prices and, hence, to a lower payoff from experimentation.

Assume now that the configuration (ρ_0, s) induces the duopolist to stop experimenting by period M. Once high-valuation consumers realize this, their willingness to pay in period M (and certainly beyond period M) is reduced not only below u_H but even below $\overline{p}_{M+1}$. This follows from the fact that an erroneously pessimistic conclusion about the state of demand is more likely with a shorter experimentation period (the probability of an error is $(1-\beta)^{M-1}$, instead of $(1-\beta)^M$). This increases the probability of nding a lower price at the competing rm, and increases consumers incentive to search. To counter this incentive, sellers must lower prices so as to make them acceptable to the high valuation consumers. But this feeds back into sellers incentive to experiment, leading them to experiment over a shorter period yet and reducing the high price below $\overline{p}_{M+1}$.

To analyze the end result of this de ne $\overline{p}_t$ by:

$$\overline{p}_t \equiv u_L + \frac{s}{(1-\beta)^{t-1}} \quad \text{for} \quad t \leq M+1,$$

and let

$$W_t = \rho_{t-1}\beta[u_H + \delta\frac{\beta\overline{p}_{t+1}}{1-\delta}] + (1-\rho_{t-1}\beta)\delta\frac{\beta u_L}{1-\delta}.$$

W_t represents the value of charging u_H in the t-th period (after $t-1$ failures), charging $\overline{p}_{t+1}$ thereafter in case of a success and charging u_L thereafter in case of a failure. This expression is predicated on the opponent switching over to u_L in period $t+1$ (upon failure), which makes the maximum price the rm can charge at that time (even if it succeeds) $\overline{p}_{t+1}$. The variable t here represents the length of the experimentation period, and in order for it to be consistent with rms incentives we must have $W_t > V(0)$ and $W_{t+1} \leq V(0)$.[9] Since $W_{M+1} \leq V(0)$ and since $\overline{p}_t < u_H$ (when $s < \widetilde{s}$) for all $t < M+1$, there must be a t, in the range $0 \leq t \leq M$, so that $W_t > V(0)$ and $W_{t+1} \leq V(0)$ (this integer need not be unique since W_t is not necessarily monotonic). This integer de nes the experimentation period under duopoly with private information. ■

An implication of the preceding proposition is:

Proposition 5: There exists an $\widehat{s} > 0$ such that for $s < \widehat{s}$: (i) If $\rho^* < \rho_0 < \rho^m$, the duopoly charges u_L from the rst period onward (i.e., never experiments) with probability one, while the monopoly charges u_H at the rst period and continues to charge u_H forever with positive probability. (ii) If $\rho^m < \rho_0$, the duopolist charges u_H only until ρ_t drops below ρ^m, and then switches to u_L. In this instance the duopolist charges u_H only because of his short-term pro t motive; once that

[9] Corresponding to this t there is a value function, analogous to (3.6), de ned in every period $1 \leq \tau \leq t$.

motive is removed he does not experiment.

Proof: Consider expressions (4.4) and (4.5) for $t = 1$ and $s = 0$. These expressions show that $\rho_0 u_H < u_L$ implies $W_1 < \frac{\beta u_L}{1-\delta}$. Therefore if $\rho_0 < \rho^m$ the duopolist nds it optimal to charge u_L from the very start. Since $\overline{p}_1$ is linear in s, there must be an $\widehat{s}$ so that for $s < \widehat{s}$, $W_1 < \frac{\beta u_L}{1-\delta}$ continues to hold.

On the other hand, if $\rho_0 u_H > u_L$, we have $W_1 > \frac{\beta u_L}{1-\delta}$ so that it is optimal for a duopolist to charge u_H. An uninterrupted sequence of failures which reduces ρ below ρ^m, places it in the set of circumstances described by (i), and the same conclusion applies.

Therefore in the set of circumstances described by (ii), information is generated as a by-product, but the duopolist does not actively invest in information generation beyond the period at which this by-product is freely available. ■

Thus, the duopoly stops experimenting after only N periods (with $N \leq M$). It is also clear from the above that generally N is smaller, the lower the search cost. The probability that a duopolistic rm reaches the wrong conclusion about demand is $(1-\beta)^N$. Moreover, by proposition 5, if the search cost is sufficiently small, experimentation ceases altogether. Then the probability that the rms erroneously conclude that demand is low becomes certain.

The pricing path that results under these circumstances resembles what has been termed cream skimming , i.e., during the early history of the industry the price is high and then it decreases. Here, if the initial ρ_0 is in the experimentation range, *both* rms experiment, which results in an average price of u_H. Then, even if a rm succeeds to sell during the experimentation period it *lowers* its price to $\overline{p}_t$; and if a rm fails to sell during the experimentation period it will, a fortiori, lower its price to u_L. So either way, by the end of the experimentation period, the average price goes down and remains low thereafter. Hence, cream skimming in our context results from the fact that rms compete and are privy to their own sales information, forcing them to offer low prices late in the history of the industry.

4.3. Duopoly Regime, Public Good Case

Now consider the duopoly equilibrium when information is a public good: Each rm is perfectly and costlessly informed about its competitor s price and sales volume at each preceding period; therefore rms have the same posterior belief, ρ, at each point in time.

Since in the public good regime both rms always have identical beliefs, it is no longer necessary to assume independent arrivals, as in the preceding subsection. Accordingly, we resume the assumption of sections 2 and 3, that the arrival of consumers across rms is *perfectly correlated* (though the qualitative features derived under this assumption also hold if arrivals are independent). That is, with probability β one consumer arrives at each rm, while with probability 1-β (a slump period), no consumer arrives at either rm. We continue to assume that consumers only know the calender time.

Proposition 6: Let $\overline{s} \equiv u_H - u_L$ and for any $s < \overline{s}$, let $T^*(s)$ be the rst t for which $s > (1-\beta)^t(u_H - u_L)$. Then, there exists a $\rho(s) \geq \rho^*$ with the following property. If $\rho(s) \leq \rho < \rho^m$, one rm charges u_L while the other rm experiments until either it makes a sale, which implies that demand is high and a high price is charged from that point onwards by both stores; or u_L sells at the same time that u_H doesn t, implying that demand is low and u_L is charged from that point onwards in both stores . If $s > \overline{s}$, the experimental price is u_H and $\rho(s) = \rho^*$. If $s < \overline{s}$, $\rho(s) > \rho^*$,the experimental price is less than u_H during periods 1,..., $T^*(s) - 1$ and is u_H at $T^*(s)$ and thereafter. If $\rho < \rho(s)$, both stores charge u_L.

Proof: In analogy with the proof of proposition 2, in the general noisy demand model, the rst step of the proof is to observe that, if $\rho < \rho^m$, one rm at the most experiments (with a price greater than u_L), while the other (contemporaneously) charges u_L.

Suppose one rm experiments and consider the other rm s best response. Simultaneous experimentation by both rms cannot provide more information than unilateral experimentation: When only one rm experiments, both rms become perfectly informed if either the high-priced rm sells or if the high-priced rm fails to sell at the same time that the low-priced one does. In the former case, the rms learn that demand is high. In the latter case, the fact that the low price sells rules out the possibility of a slump, forcing the conclusion that the high priced rm s failure is the result of low demand. Finally, failure to sell by both rms implies a slump period, is therefore perfectly uninformative, and leaves the rms initial beliefs about demand intact.

On the other hand, when both rms experiment simultaneously, only a success (at both rms[10]) gives precise information. Failure (by both) is ambiguous since it can be attributed to either a demand slump or to low demand. So given that one rm already experiments, the other can only *lose* information by experimenting at the same time. On the other hand, myopically, the optimal price is u_L, since $\rho < \rho^m$. Thus experimenting at the same time is not only less informative than not experimenting, but is also more costly (in terms of foregone expected pro ts at the present period). Thus if $\rho < \rho^m$, one rm at the most experiments.

If $s > \overline{s} = u_H - u_L$, it is too costly for consumers to search (even if one rm charges u_H while the other charges u_L, which is the best scenario for search). Then an experimenting duopolist can charge u_H, so he experiments whenever the monopoly does. Thus, if $\rho^* < \rho < \rho^m$, one rm charges u_L while the other experiments with u_H until either there is a sale at u_H, (which implies that demand is high) and both rms charge u_H from that point onwards; or u_L sells at the same time that u_H doesn t (implying that demand is low), and u_L is charged by both rms from that point onwards. (of course, if $\rho < \rho^*$, both rms stop experimenting).

Consider now $s < \overline{s}$ and $\rho < \rho^m$ and assume one rm is experimenting with u_H in period t. Then, the consumer who visits it has a probability of $(1-\beta)^t$ of nding the price u_L at the other rm ($(1-\beta)^t$ is the probability that no consumer has shown up, in which case the other rm has not yet concluded that $\widetilde{u} = u_H$.) Therefore the payoff to search is $\sigma = (1-\beta)^t(u_H - u_L)$. If $t < T^*(s)$, then $\sigma > s$, so the consumer will search, which contradicts Lemma 2. Consequently, during periods $1, ...T^*(s) - 1$, the experimental price must be $u_L + s/(1-\beta)^t < u_H$. On the other hand, if $t \geq T^*(s)$, $\sigma < s$, so the experimental price is u_H. Therefore the price dynamics are such that one rm experiments with either $u_L + s/(1-\beta)^t$ or with u_H (depending on the calendar date) until uncertainty is resolved, or ρ is sufficiently low that experimentation is abandoned. Given that experimentation is with a price $< u_H$, the return from experimentation is smaller than it is under the monopoly scenario. Consequently, there exists a $\rho(s) > \rho^*$ so that experimentation does not take place in $(\rho^*, \rho(s))$. This completes the proof. ■

As in the noisy demand model, the ability to free ride on one s competitor s experimentation efforts promotes a tendency towards unilateral experimentation, which increases price dispersion and the consumers propensity to search, diluting the incentive to experiment even when *individual* experimentation (with u_H) is optimal.

[10] By assumption, rms that charge identical prices make the same number of sales.

The resulting price path under these circumstances resembles what has been termed penetration pricing, i.e., prices during the early history of the industry are low and then they rise. Here, when the initial ρ_0 is in the experimentation range $\rho_0 \in (\rho(s), \rho^m)$, one rm experiments while the other rm does not. So during the experimentation phase the average price in the industry is either $(1/2)u_L + (1/2)(u_L + s/(1-\beta)^t)$ or $(1/2)u_L + (1/2)u_H$. Either way, it is less than u_H. Then, if demand is high and there is a boom period, both rms learn the state of demand and raise their prices to u_H which they sustain thereafter. Hence penetration in this context means rms are receiving favorable news about the state of demand which enables them to raise their prices.[11]

Another possibility is cyclical pricing. Suppose that initially $\rho_0 > \rho^m$ so that both rms charge u_H. If there are no successes, the belief in high demand erodes to ρ, $\rho(s) < \rho < \rho^m$, whence only one rm experiments and, moreover, the experimenting rm s price is no higher than u_H, so that the average market price drops below u_H. A success at this phase leads each rm to increase its price back up to u_H. Thus prices are observed to decrease and subsequently increase.[12]

5. Concluding Remarks

In the model considered above, the nature of the uncertainty applies to demand at the high price. But it is equally plausible that rms are uncertain about demand at *low* prices. For example, suppose type 1 consumers demand only one unit if the price is less than or equal to u_H, while type 2 consumers demand two units if the price is less than or equal to u_L, and one unit if the price is greater than u_L (but not greater than u_H). The proportion of type 2 consumers is $\widetilde{a}$ and, again, the rms are uncertain about $\widetilde{a}$: With probability γ, the proportion $\widetilde{a}$ is a_L, with probability 1-γ, $\widetilde{a} = a_H$. In this case, it is the *low* price which is informative; experimentation involves charging u_L (and observing the number of sales). Assume the non-trivial case in which the monopoly price is u_L if $\widetilde{a} = a_L$ and u_H if $\widetilde{a} = a_H$.

The analysis of this case is completely analogous to the previous case; there is a cutoff belief at which the monopoly stops experimenting, and subsequently charges u_H. And again the search effect limits market power for the duopoly, if search costs are sufficiently low. Consider the private information regime. Here,

[11] Of-course, if rms do not receive favorable news they will *decrease* their prices; however, the point is that there *exists* a sample path under which prices go up over time.

[12] Pricing at the experimental stage thus resembles a sale.

the observation of a price above u_L means the firm has grown *pessimistic* and *stopped* experimenting. The incentive to search results from the expectation that its competitor is still optimistic enough to experiment - i.e. is still charging u_L. This restricts the highest price which can be charged below u_H, if s is sufficiently small, no matter what the posterior belief is - the analogue of part (i) of proposition 1. But, in contrast to the previous model, here, the effect is to make experimentation (with u_L) *more*, rather than less profitable. This is because narrowing the gap between the "sure", high price, and the experimental low price reduces the cost of experimentation, and hence there is more experimentation; i.e. the duopoly cutoff belief, say ρ', satisfies $\rho' < \rho^*$. Thus, in both cases uncertainty increases competition in proportion to the cost of search. But in one case - when uncertainty applies to demand at high prices - this has the effect of reducing experimentation, while if uncertainty is about demand at low prices, incentives to experiment increase.

To summarize, then, this paper examines the issue of learning and experimentation when firms *and* consumers are imperfectly informed. The principal conclusion under these conditions is that the degree of experimentation is different from the monopoly case. This results, although for different reasons, both when firms are privy to their own sales experience and when the sales experience is publicly known. It is also true both when uncertainty is about demand at high and low prices, although the effect on experimentation is different. We have also shown that a whole variety of price paths is possible: When information is a private good the price path is necessarily downward sloping over time, which corresponds to cream skimming. On the other hand, when information is a public good the price path might be upwards sloping which corresponds to penetration pricing, or it might be cyclical. Therefore, learning in this model and the resulting price path are shown to depend on the market structure, the nature of uncertainty, and on the degree to which firms can keep their information private.

References

[1] P. Aghion, P. Bolton, C. Harris and B. Jullien, Optimal learning by experimentation, *Rev. Econ. Stud.* 58 (1991), 621-654.

[2] P. Aghion, R. Espinoza and B. Jullien, Dynamic duopoly with learning through market xperimentation, *Econ. Theory* 3 (1993), 517-539.

[3] S. Bikhchandani, D. Hirschleifer and I. Welch, A theory of fads, fashion, custom and cultural change as informational cascades, *J. Polit. Econ.* 100 (1994), 992-1026.

[4] A. Banerjee and D. Fudenberg, A simple model of bayesian word of mouth, mimeo, Harvard University, 1994.

[5] D. Berry and B. Fristedt, Bandit Problems: Sequential Allocation of Experiments, London: Chapman and Hall (1985).

[6] P. Bolton and C. Harris, Strategic Experimentation, Discussion Paper TE/93/261, LSE, London (1993).

[7] K. Burdett and K. Judd, Equilibrium price dispersion, *Econometrica* 51 (1983), 955-970.

[8] G. Butters, Equilibrium price dispersions, *Rev. Econ. Stud.* 44 (1977), 465-492.

[9] A. Caplin and J. Leahy, Sectoral shocks, learning and aggregate uctuations, *Rev. Econ. Stud.* 60 (1993), 777-794.

[10] A. Carlson and P. McAfee, Discrete equilibrium price dispersion, *J. Polit. Econ.* 91 (1983), 480-493.

[11] P. A. Diamond, Model of Price Adjustment, *J. Econ. Theory* 3 (1971), 156-168.

[12] D. Easley and N. Kiefer, Controlling a stochastic process with unknown parameters, *Econometrica* 56 (1988), 1045-1064.

[13] A. Fishman and C. Fershtman, Price cycles and booms: dynamic search equilibrium, *Amer. Econ. Rev.* 82 (1992), 1221-1233.

[14] A. Fishman and Rob, R., Experimentation and Competition, CARESS working paper , University of Pennsylvania, (1996).

[15] D. Fudenberg and J. Tirole, A duopoly theory of signal jamming, *Rand J. Econ.* 17 (1986), 366-376.

[16] S. Grossman, R. Kihlstrom and L. Mirman, A Bayesian approach to the production of information and learning by doing, *Rev. Econ. Stud.* 44 (1977), 533-548.

[17] J. Harrington, Experimentation and learning in a differentiated product duopoly, *J. Econ. Theory* 66 (1995), 275-279.

[18] V. J. Mahajan, E. Muller and F. M. Bass, New product diffusion models in marketing: a review and directions for research, *J. Marketing* 54 (1990), 1-26.

[19] A. McLennan, Price dispersion and incomplete learning in the long run, *J. Econ. Dynam. Cont.* 7 (1984), 331-347.

[20] L. Mirman, L. Samuelson and A. Urbano, Monopoly experimentation, *Int. Econ. Rev.* 34 (1993), 549-564.

[21] L. Mirman, L. Samuelson and E. Schlee, Strategic information manipulation in duopolies, *J. Econ. Theory* 62 (1994), 363-384.

[22] E. Prescott, The multiperiod control problem under uncertainty, *Econometrica* 40 (1972), 1043-1058.

[23] J. Reinganum, A simple model of equilibrium price dispersion, *J. Polit. Econ.* 87 (1979), 851-858.

[24] M. Riordan, Imperfect information and dynamic conjectural variations, *Rand J. Econ.* 16 (1985), 41-50.

[25] R. Rob, Equilibrium price distributions *Rev. Econ. Stud.* 44 (1985), 493-510.

[26] R. Rob, Learning and capacity expansion under demand uncertainty, *Rev. Econ. Stud.* 58 (1991), 655-675.

[27] M. Rothschild, A two armed bandit theory of market pricing, *J. Econ. Theory* 9 (1974), 185- 202.

[28] C. Shapiro, Optimal pricing of experience goods, *Bell J. Econ.* 14 (1983), 497-507.

[29] G. Stigler, A theory of oligopoly, *J. Polit. Econ.* 72 (1964), 44-61.

[30] J. Zeira, Informational cycles, *Rev. Econ. Stud.* 61 (1994), 31-44.