

Endogenous Differentiation of Information Goods under Uncertainty

Robert S. Gazzale
rgazzale@umich.edu
Department of Economics

Jeffrey MacKie-Mason*†
jmm@umich.edu
School of Information and Department of
Economics

University of Michigan
Ann Arbor, MI 48109

ABSTRACT

Information goods can be reconfigured at low cost. Therefore, firms can choose how to differentiate their products at a frequency comparable to price changes. However, doing so effectively is complicated by uncertainty about customer preferences, compounded by the fact that the search for a good product niche is carried out in competition with other searching firms.

We study two firms that differentiate their information goods. The firms simultaneously compete in product configuration and price. We assume a non-uniform distribution of consumers: the largest number prefer a product located at a “sweet spot,” but the rate at which the customer density falls off away from this product configuration is unknown. Our characterization reflects the standard tradeoff between exploitation (current profit) and exploration (learning to enhance future profit). In our model firms balance current profits from competing for a mass and a niche market, while learning about the profitability of these alternative strategies.

We show that the amount of learning that firms will undertake depends on the convexity or concavity of the profit function in the rate of demand fall-off. In our model firms have an incentive to learn, and can use both price and product configuration in order to explore. We show that the ability to explore in product characteristic space leads to a previously unidentified consequence of learning: attenuation of competition. The incentive to learn induces firms to differentiate their products more than they would if the value of learning were ignored. This leads to decreased direct com-

petition with rivals, and thus higher prices and profits than if the firms were acting myopically. Thus, we might expect that when firms are not well informed about consumer preferences for information goods — as might be especially true in new markets for innovative products — product diversity will be higher and direct competition will be smaller than might otherwise be expected.

1. INTRODUCTION

Information goods can be reconfigured at low cost. For example, information aggregators (newspapers, databases) can unbundle and re-bundle information objects in a variety of ways. In the print-on-paper world, low-price bundles (like daily newspapers) generally are offered in one standard edition (with perhaps a small number of minor variants). Extensive customization is provided by information services at a high cost. With electronic publication, the cost of customizing a standard edition can approach zero.

There has been little research on how firms choose to differentiate their information goods. This problem is especially challenging because firms rarely have complete information about the preferences of potential customers over product characteristics. Thus, over time they make their price and product configuration decisions based not only on expected current profits but also based on the value of the learning they expect from each period’s offering. To further complicate things, this search for a good product niche is carried out in competition with other searching firms.

We consider two firms competing in two dimensions: product configuration and price. We model product configuration as a one-dimensional space: a line on which firms choose a location. In certain markets it is clearly technologically feasible, and perhaps optimal, for a provider of information goods to customize its offerings so that it in effect occupies multiple locations in product space.¹ We limit, however, the firms in our model to choosing one location in any period for a few reasons. First, we do so in order to focus on the ability of firms to control the degree of product differentiation in an environment where firms need to learn about the attractiveness of differentiation. Second, even if firms could completely customize their offerings based on certain customer characteristics, it remains quite likely that firms will attempt to differentiate their offerings from those of its competitors in other ways. Thus, our model might be

*MacKie-Mason gratefully acknowledges support from an IBM Partner Fellowship, and from NSF grant CSS-9988715.

†The authors appreciate helpful comments from audiences at Johns Hopkins University, the University of Maryland Robert H. Smith School of Business, and Ford Motor Research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC’01 October 14-17, 2001, Tampa, Florida, USA
Copyright 2001 ACM 1-58113-387-1/01/0010 ...\$5.00.

¹See, for example, Farag and Van Alstyne [7].

interpreted as one in which firms choose a *brand identity*. Generalizing the model to firms that offer multiple product configurations is a worthwhile task for future research.²

The largest number of customers most prefer a product located at a “sweet spot”, with the density of customers preferring other products falling off with distance from the sweet spot. The firm’s optimal product configuration needs to balance the rewards from selling to the many customers near the sweet spot against the dual costs of losing customers in the less densely populated tails and of lower prices due to fiercer competition near the sweet spot. This is intended to suggest the choice between competing for a mass market and a niche market.³

To introduce uncertainty about consumer preferences we assume that the firms know the location of the sweet spot, but not the rate at which demand falls off with distance from the sweet spot. We use a two-period model to allow the firms an opportunity to learn about preferences from their experience. Now we have a problem of *exploitation* versus *exploration*: The locations and prices firms choose each period will determine current profits, but (in the first period) will also reveal information that might increase their ability to extract profits in future periods. The most informative location/price combination will not generally yield the highest expected current profits. Therefore, the optimal product configuration and pricing decision generally balances the value of learning against the cost of foregone current profits.

Grossman et al. [9] are among the first to study have identified the *exploration* versus *exploitation* tradeoff in an economic problem.⁴ As an example, they consider an individual’s consumption of an item whose value is unknown. Each time the consumer tries the item, the value she receives is equal to the underlying value plus a stochastic shock. Thus the more she experiments with an item, the better she knows its true value. Under the conditions outlined, the non-myopic consumer makes larger purchases of this item in order to learn its value and make better decisions in future periods. Subsequent authors, such as McLennan [17] and Aghion et al. [1] study experimentation by a monopolist uncertain about the demand for its product, and derive conditions under which there will be adequate learning.

In a related paper, Harrington [10] considers duopolists competing in price in a differentiated products market with firms uncertain about the degree of substitutability among products. However, in contrast to our model with endoge-

nous product differentiation, Harrington’s firm locations are fixed. He shows that under certain demand conditions firms wish to learn in the first period, while under other conditions they do not wish to learn. With price the only strategic variable in his model, greater learning follows from a greater price difference between the two firms. In our model, with firms choosing both price and product configuration, learning can be increased by lowering price (thereby attracting more niche customers far from the sweet spot) or by differentiating products. Our model also differs because Harrington’s firms are uncertain about the degree of differentiation between their products, whereas ours are uncertain about the distribution of consumer preferences. An implication of this difference is that, for a given price decrease (holding everything else constant) a firm in Harrington’s model knows the number of new customers who enter the market, but not how many customers the firm takes from its rival. In our model, neither the number of new customers in the market nor the number of customers taken from its rival is known with certainty.

Our model is also related to the Hotelling literature on endogenous product differentiation.⁵ The standard model in that large literature has two firms locating on a line, but consumer preferences are distributed uniformly on a segment rather than more densely around a sweet spot. We work with a richer model of consumer preferences because the uniform distribution has only one parameter (the width of the line segment), and to model uncertainty would have required that we suppose firms did not know how far uniform consumer preferences over product configurations extended, which does not readily map to familiar information goods markets. Our approach allows for uncertainty in a natural way: firms are not sure how rapidly consumer demand falls off away from the sweet spot.

The standard Hotelling model also fixes price and lets firms compete only in location (in contrast to Harrington, who reverses the state and control variables). We endogenize both price and product configuration. Finally, although some work on the Hotelling problem incorporates firm uncertainty, to our knowledge we are the first (other than Harrington) to study learning in a location model of endogenous product differentiation.

Our work is also related to the growing literature, using both empirical methods and simulations, that studies the product positioning of information goods. Clay et al. [6] find that as new firms entered online book selling, prices remained flat or rose. They document a wide degree of heterogeneity product and pricing strategies. They conclude that “the real puzzle is the stores with wide selection and average prices”, but in a new market with substantial learning, our model suggests that experimenting with this and various other configurations may not be so puzzling after all.

Segev and Beam [18] report on some of the practices of electronic brokerages, who provide prices for other goods or services, and potential matches to trading partners. They find tremendous uncertainty about profit maximizing strategies, and that in response experimentation with prices and product configurations is greater than might be expected. Through a simulation they find that in this environment brokers will do best to differentiate widely, for example by

²Some authors studied firms in Hotelling models that can sell more than one product, each with a different “location” or configuration. These authors make the extremely limiting assumption that price is fixed exogenously, so that competition is only in location, as well as the other restrictive assumptions of the Hotelling models identified above. Even in these highly stylized models results are hard to obtain and are inconsistent. For example, Gabszewicz and Thisse[8] find that two firms spread their products across the space but locate each of their varieties right next to the competing firm’s most similar variety. But Martinez-Giralt and Neven [16], with only one minor change in assumptions, finds that firms locate all of their products in a cluster, yet locate those clusters as far from the competitor’s cluster as possible.

³MacKie-Mason et al. [15] analyze the effect that Internet service architecture can have on the choice between mass market and niche product configuration.

⁴[11] presents an early discussion of exploration versus exploitation in his formalization of the adaptive learning problem.

⁵See Anderson et al. [2] for a thorough survey.

either focusing on serving buyers (charging high fees to sellers and low fees to buyers), or focusing on serving sellers.

In section 2 we present our model, with details on the information goods market, firm behavior and consumer behavior. We then solve for the subgame perfect equilibrium of the two-stage game in section 3. We discuss the results and possible generalizations in section 4. Our primary result is that firms will use first-period price and product configuration in order to increase learning. However, in contrast to standard models of firm learning, this is not at the expense of first-period profits. Firms are able to increase learning by increasing the level of differentiation between their products. This reduction of competition enables firms to increase prices and thus increase short-term as well as long-term profits.

2. THE MODEL

2.1 The Market

We consider a market for an information good that can be differentiated in one dimension. An example would be Web sites that provide news content, differentiated by the ratio of national to international news. A more general model would permit differentiation in multiple dimensions. We represent this dimension as a line on the real numbers. The product offered by each firm is characterized as a location on this line.




Figure 1: Distribution of consumers' most preferred product configurations over range of product possibilities

We characterize a consumer by the product configuration (location) she most prefers. We then map the distribution of consumers over the product space line, with the vertical height above the line representing the number of consumers for whom that location represents their most preferred product configuration (see figure 1). We assume that there is a single product configuration that is most preferred by the largest number of consumers. We call the location of these α consumers the “sweet spot”, normalized to be the zero on the horizontal axis. Along the product space axis, the distance from the most popular location is represented by l , which can be either positive or negative. The number of consumers decreases as one moves away from the sweet spot

at a rate of γ . Thus the number of consumers located at l is $\alpha - \gamma|l|$.

2.2 Consumer Behavior

We assume consumers purchase at most one of the two competing goods in each period. There is no cost to evaluating the options and choosing a provider. A consumer receives a utility of r if she consumes her most preferred good, and an amount that decreases at rate c the further the consumed good is from her most preferred configuration. Letting (l, p) represent a product's configuration and price, a consumer of type t receives $r - p - c * |t - l|$. Consumers select the good that provides the greater utility, or neither if utility would be negative. That some consumers may choose to purchase nothing is another factor that distinguishes our model from the standard Hotelling approach, which assumes that all consumers purchase. In our model there is both an intensive and extensive margin: A firm can lose (niche) customers to the “outside option” or (mass market) customers to “head-to-head” competition with the other firm.

We assume that the distance cost is linear for analytic convenience. The constant cost c could be interpreted as the loss in utility per article as a bundled information good offers fewer articles of the type the consumer wishes to read (e.g., less national news). In a more general representation of preferences the distance cost might be nonlinear.

The density of consumers who purchase a given firm's good according to the behavioral rule above constitute that firm's demand. For two reasons we add a stochastic component to each firm's demand. First, it is unreasonable to model a world with firm uncertainty but to then assume that every consumer makes exactly the right decision every period. The second reason we provide below, after we explain the information available to the firm and its behavior. To implement stochastic demand we assume that each firm i 's demand is subject to an additive random variable, ϵ_i , whose cdf $G_i()$ has a mean of zero and variance of σ_ϵ .

2.3 Firm Behavior

Two firms compete in this market for two periods. The firms are *ex ante* identical. In each period, at zero cost, each firm can differentiate its product by choosing a location on the line, at the same time announcing a price. The firm's objective is to maximize the sum of discounted profits, which are equal to revenues because we assume that location and production costs are zero to capture the easy reconfigurability and reproduction of information goods.

We assume the values α , c , r , and the distributions of the ϵ_i and of γ are known to both firms. The need for learning arises because they do not know the value of γ . However, the firms have the same distribution of prior beliefs over γ , denoted by the CDF $F(\gamma)$, and thus the same expected valuation ($\hat{\mu}_0$).

After the first period of trade, the prices, locations and number of consumers served by each firm is common knowledge. Conditional on this knowledge and the prior belief $\hat{\mu}_0$, firms update their beliefs about the value of γ . Our primary goal is to investigate how the opportunity to learn about the value of γ affects the conduct of the firm in the first period.

Above we gave one reason why we assume each firm's demand has an additive stochastic component, ϵ_i . The second reason is that given the common knowledge assumed above, almost any location and price in the first period would re-

veal the true value of γ to each firm. The intuition is that the density function is piecewise linear with each slope having the same absolute magnitude, and firms already know one point on the function (α). If they could observe demand from this density perfectly then they could solve for a second point on the density function and could perfectly calculate the slopes. In no realistic problem can firms perfectly infer all relevant consumer preference information from a single experiment, so we add a noise term to ensure incomplete inference.

3. SUBGAME PERFECT EQUILIBRIUM

In this section we solve the model for a subgame perfect equilibrium. Our two-period subgame perfect framework does enable us to draw valuable inferences about the more realistic case where the number of periods is larger. First, we use second period behavior as a “no learning” or myopic benchmark against which to compare the actions of firms who take into account the consequences of current period actions on subsequent period profits. Second, adding additional periods does little alter the incentives of the game. We can thus view the first period as representing periods under which the firms act under uncertainty and the final period as the limiting case as the value to learning goes to zero.

We believe there is substantial value to the study of the equilibria of a tractable but reasonably realistic model of the dynamics of learning and product configuration. First, by knowing the equilibria of the game, we will know something about the behavioral incentives facing firms that find themselves out of equilibrium in a more realistic setting. Second, we are able to obtain explicit analytic results, which enables us to establish general predictions about the comparison firms that strategically learn and those that do not. In future research these predictions can then be tested against empirical data. Further, the predictions of the game-theoretic equilibria can be used as a guide to the design of intelligent heuristics; we discuss this possibility in relation to our own research on software agent heuristics in section 4.

Since the game is finite, we use backward recursion: We first solve for optimal play by the two firms in the second period, conditional on their updated expectation, $\hat{\mu}_1$, from the first period. In the subgame we look for Nash equilibria, in which if each firm makes the best play conditional on the choices of the other firm, the choices will be mutually consistent. Then, given the solutions for prices and locations in the second period as a function of $\hat{\mu}_1$, we solve for the optimal price and location choices by the firms in the first period. Since their objective is to maximize the sum of discounted profits over the two periods, their first period choices will take into account not only profits in the first period, but also the effect of these first period actions on expected second period profits due to their learning about the slope of the customer preference density.

We denote the leftmost firm as firm 1, and the rightmost as firm 2, and their locations as l_1 and l_2 respectively. Given the consumer choice rule, for any l_1 and l_2 we can identify the leftmost and rightmost consumer types who purchase one of the goods as follows:

$$\begin{aligned} t_l &= l_1 + \frac{p_1 - r}{c}, \\ t_r &= l_2 + \frac{r - p_2}{c}. \end{aligned}$$

The consumer type which is indifferent between the offerings of the two firms, t_m , will be

$$t_m = \frac{l_1 + l_2}{2} + \frac{p_2 - p_1}{2c}.$$

See figure 2.




Figure 2: Illustration of market division between firms

PROPOSITION 1. *In a pure strategy equilibrium, all customers located between the two firms are served. Thus there exists a unique t_m .*

(Proofs for the results are given in an appendix.)

Without loss of generality, assume that $t_m \geq 0$. Then demand for each firm is

$$D_1(p_1, p_2, l_1, l_2) = \int_{t_l}^0 (\alpha + \gamma l) dl + \int_0^{t_m} (\alpha - \gamma l) dl + \varepsilon_1, \quad (1)$$

$$D_2(p_1, p_2, l_1, l_2) = \int_{t_m}^{t_r} (\alpha - \gamma l) dl + \varepsilon_2. \quad (2)$$

Profits for each firm are:

$$\begin{aligned} \pi_1(p_1, p_2, l_1, l_2) &= p_1 D(p_1, p_2, l_1, l_2) \\ \pi_2(p_1, p_2, l_1, l_2) &= p_2 D(p_1, p_2, l_1, l_2). \end{aligned}$$

3.1 Second Period Equilibrium

Given $\hat{\mu}_1$, their expectation of γ after period 1, firms maximize total expected profit.⁶ Taking the other firm’s price and location as given, each firm calculates the first order conditions for its profit function subject to two constraints.

⁶The profit functions are linear in γ , so we can replace γ by its expected value $\hat{\mu}$ when calculating expected profits.

The first is that all consumers who purchase receive non-negative utility. The second is that $l_2 \geq l_1$. We assume for the moment that neither constraint binds. This yields four best response functions in four unknowns:

$$\begin{aligned} l_1(p_1; p_2, l_2) &= 0 \\ p_1(l_1; p_2, l_2) &= 0 \\ l_2(p_2; p_1, l_1) &= 0 \\ p_2(l_2; p_1, l_1) &= 0. \end{aligned}$$

which are then solved to find the Nash equilibrium.⁷ Only one of the sixteen solutions to this system satisfies the second-order conditions, so in the unique equilibrium firms set price and location as follows:

$$p_1^* = \frac{3c\alpha}{8\hat{\mu}} \quad (3)$$

$$p_2^* = \frac{3c\alpha}{8\hat{\mu}} \quad (4)$$

$$l_1^* = \frac{r}{c} - \frac{7\alpha}{8\hat{\mu}} \quad (5)$$

$$l_2^* = -\frac{r}{c} + \frac{7\alpha}{8\hat{\mu}}, \quad (6)$$

which will yield the following expected profit

$$\begin{aligned} E[\pi_i|\hat{\mu}] &= p_i^* \int_0^{l_2^* + \frac{r}{c} - \frac{p_2^*}{c}} (\alpha - \hat{\mu}l) dl \\ &= \frac{9c\alpha^3}{64\hat{\mu}^2}. \end{aligned}$$

We can provide some economic interpretation to the best response functions and the resulting equilibrium. If we look solely at the location decision of firm 1, setting marginal benefit equal to marginal cost implies that for an incremental move closer to its opponent, the number of customers that firm 1 gains from its rival equals the number lost on the outside margin. By differentiating the bounds of integration of the demand function with respect to location, we see that the former is equal to $\frac{1}{2}$ the height at t_m and the latter is equal to the height at t_r or t_l . The best response in terms of price alone is more complicated, but still involves balancing the internal and external margins. For our symmetric equilibrium where the middle indifferent consumer is at the sweet spot, there are a continuum of price and location pairs that satisfy the condition that the external and internal margins must be equal. Whether the firms locate near the middle at a relatively low price or closer to the midpoint of 0 and t_r (or t_l) at a relatively high prices depends on how much firms desire to fight for the mass market. The desirability of the mass market in turn depends on the expected slope of consumer density, $\hat{\mu}$: the lower is $\hat{\mu}$, the more valuable are the niche markets relative to the (more competitive) mass market. This can be seen from the effect of $\hat{\mu}$ on equilibrium prices and locations in equations (3)-(6).

We next examine at how expected profit depends on $\hat{\mu}$:

$$\partial E[\pi_i|\hat{\mu}] / \partial \hat{\mu} = -\frac{9c\alpha^3}{32\hat{\mu}^3} < 0 \quad \forall \hat{\mu} > 0 \quad (7)$$

$$\partial^2 E[\pi_i|\hat{\mu}] / \partial \hat{\mu}^2 = \frac{27c\alpha^3}{32\hat{\mu}^4} > 0 \quad \forall \hat{\mu} > 0. \quad (8)$$

Expected profits are decreasing in $\hat{\mu}$. This is due to the fact that, when we increase γ , demand falls off more sharply as we move away from the sweet spot. Equation (8) implies that expected profits are convex in $\hat{\mu}$. That expected profits are convex in $\hat{\mu}$ implies the firms behave as if risk-loving, i.e., they prefer more variability in their posterior mean on γ . That is, they prefer to learn more information about the actual value of γ , as this permits them to do a better job optimizing in period 2. We discuss this point in further detail in Section 4.

At this point, we must return to the neglected constraints, which force us to put bounds on the range of beliefs for which the above equilibrium holds. As shown in Proposition 1, in any pure-strategy equilibrium there are no unserved consumers between the two firms. This will only be true in a symmetric equilibrium if a consumer located at the sweet spot has non-negative utility. From the consumer utility function and the equilibrium prices in (3)-(4) this will be true in our equilibrium only if $\hat{\mu} \geq \frac{5\alpha c}{8r}$. Likewise, to ensure that $l_2 \geq l_1$, it is necessary that $\hat{\mu} \leq \frac{7\alpha c}{8r}$.⁸ The above symmetric equilibrium thus holds for $\hat{\mu} \in [\frac{5\alpha c}{8r}, \frac{7\alpha c}{8r}]$.⁹ We can show that for values of $\hat{\mu}$ outside of this range, the unique equilibrium will be asymmetric.

The “reasonable” width of this range for $\hat{\mu}$ is an empirical question, since there are no constraints other than non-negativity on the free parameters. The important question is whether this region captures an economically interesting set of problems. We believe that it does. This region is important because it is precisely the region in which firms desire to compete *both* for the mass market and remain attractive to many in the niche market. This is an accurate description for many markets of interest.

Outside of this range for $\hat{\mu}$ the model corresponds to different types of markets. For example, consider the effect of decreasing $\hat{\mu}$ inside the symmetric range: the niche markets become more attractive. As one would expect, our equilibrium conditions show that the firms are less inclined to compete for the mass market: the firms increase product differentiation and prices increase as firms act more like local monopolists in the niches. As we continue to decrease $\hat{\mu}$ past the symmetric range, an interesting thing happens. We can show that if one firm is serving the sweet spot, the best response of the other firm is to not compete at all for consumers served by the former. In short, the niche becomes so attractive that it is preferable to be a local monopolist for this niche. In this situation there is no head-to-head competition for the mass market, and thus the firms do not balance their appeal to mass market and niche customers.

⁸That $l_2 \geq l_1$ is arbitrary, but we used it to specify the demand functions facing each firm in our solution for the equilibrium. The same parameter restriction would hold if we reversed the firms and imposed $l_1 \geq l_2$. More generally, what is required is that firms compute their expected demand consistently with their equilibrium location.

⁹In order to ensure that $\hat{\mu}$ is in this range in the second period, it suffices that the firms place zero probability on any γ outside of this range in the first period.

⁷The best response functions are extremely long so we do not reproduce them here. They are available from the authors upon request.

Whether we are in this case, the case where $\hat{\mu}$ is so steep as to make the niches unattractive, or the intermediate case in which firms balance between selling to both mass and niche markets of course depends on the market in question. In this paper, we limit our attention to the case where $\hat{\mu}$ is inside the specified range, and leave an analysis of the dynamics of the other two cases to future work.

3.2 First Period Equilibrium

Having solved for the equilibrium in the last period, we now characterize the Nash solution for the first period, taking into account the effect of first period choices on expected second period prices, locations and profits. The link between periods is through learning: realized first period demand provides information about the value of γ , so that generically $\hat{\mu}_1 \neq \hat{\mu}_0$, and second period prices and locations are functions of $\hat{\mu}_1$ (see equations (3)-(6)).

Firms update their beliefs on γ based on the implication of their locations and the total realized demand for the preference density slope, taking into account that demand has a mean zero stochastic component. Making use of the symmetry of the demand density and that firms locate equidistant from the sweet spot, we derive the following total demand equation for $D(p_1, p_2, l_1, l_2) \equiv D_1(p_1, p_2, l_1, l_2) + D_2(p_1, p_2, l_1, l_2)$:

$$\begin{aligned} D(p_1, p_2, l_1, l_2) &= 2 \int_0^{t_r} (\alpha - \gamma l) dl + \varepsilon \\ &= 2\alpha t_r - \gamma t_r^2 + \varepsilon \\ &= 2\alpha \left(l_2^* + \frac{r-p_2}{c} \right) - \gamma \left(l_2^* + \frac{r-p_2}{c} \right)^2 + \varepsilon. \end{aligned} \quad (9)$$

Rearranging equation (9) gives:

$$\frac{2\alpha}{l_2 + \frac{r-p_2}{c}} - \frac{D(p_1, p_2, l_1, l_2)}{\left(l_2^* + \frac{r-p_2}{c} \right)^2} = \gamma - \frac{\varepsilon}{\left(l_2 + \frac{r-p_2}{c} \right)^2}. \quad (10)$$

PROPOSITION 2. *The left-hand side of (10) is an unbiased estimator for γ with variance equal to $\frac{\sigma_\varepsilon^2}{\left(l_2 + \frac{r-p_2}{c} \right)^4}$.*

After observing first period total demand, firms apply Bayes' rule to combine their prior beliefs with this new unbiased estimate to obtain updated beliefs on γ .

We can view the choice of prices and locations followed by a demand observation as an experiment. An experiment is more informative if it reduces the variance of the estimator. As the denominator of the variance is equal to $t_r^4 = t_l^4$, reducing the variance is accomplished by increasing the number of niche consumers served. The intuition is that as the firms increase their reach (i.e., by moving t_l and t_r further away from the sweet spot), demand is more affected by γ , and the relative effect of ε diminishes. Firms increase their reach by either lowering their prices or moving away from the sweet spot.¹⁰

¹⁰That the distribution of ε is independent of total demand is an analytic convenience. Increasing reach will assuredly provide a more informative experiment if the ratio of variance of ε to total demand is non-increasing as $t_r = -t_l$ increases.

We define the posterior cdf of γ as $F(\cdot | D(p_1, p_2, l_1, l_2))$.¹¹ We apply Bayes' Rule to get:

$$F'(\gamma | D(p_1, p_2, l_1, l_2)) = \frac{G'(D - 2\alpha t_r + \gamma t_r^2) F'(\gamma)}{H'(D)},$$

where

$$H'(D) = \int G'(D - 2\alpha t_r + \gamma t_r^2) F'(\gamma) d\gamma.$$

Because both firms have access to the same information about the outcome of the first-period experiment, they arrive at the same updated distribution of beliefs about γ and thus the same and expected value, $\hat{\mu}_1$. We have shown previously that with the same bounded beliefs about γ , the unique second-period equilibrium is symmetric. Therefore, for any first-period equilibrium the two firms have the same expected second-period profit equal to

$$W(p_1, p_2, l_1, l_2) = \int E \left[\pi_i \int F'(\gamma | D(p_1, p_2, l_1, l_2)) d\gamma \right] H'(D) dD.$$

Both firms maximize cumulative profits discounted at rate δ , so the first-period value function for firm 2 is

$$V_2(p_1, p_2, l_1, l_2) = \begin{cases} \int_{t_m}^{t_r} (\alpha - \gamma l) dl + \delta W(p_1, p_2, l_1, l_2) & \text{if } t_m \geq 0 \\ \int_{t_m}^0 (\alpha + \gamma l) dl + \int_0^{t_r} (\alpha - \gamma l) dl + \delta W(p_1, p_2, l_1, l_2) & \text{if } t_m \leq 0. \end{cases}$$

A firm's best reply function, ϕ_i , is a pair of price-location values defined by

$$\phi_i(p_j, l_j) \in \arg \max_{p_i, l_i} V_i(p_i, l_i; p_j, l_j).$$

PROPOSITION 3. *ϕ_i exists.*

A symmetric subgame perfect equilibrium exists iff there exists $\hat{l}_1 = -\hat{l}_2$ and $\hat{p}_1 = \hat{p}_2$ such that

$$\begin{aligned} \{\hat{p}_1, \hat{l}_1\} &= \phi(\hat{p}_2, \hat{l}_2) \\ \{\hat{p}_2, \hat{l}_2\} &= \phi(\hat{p}_1, \hat{l}_1) \end{aligned}$$

PROPOSITION 4. *A symmetric subgame perfect equilibrium exists.*

We now establish our main result. We wish to establish the effect that the opportunity to learn has on first-period price and product differentiation decisions. To do this, we compare the equilibrium first-period prices and locations to those that would be an equilibrium if both firms ignored the value of learning.

Since both prices and locations can be changed costlessly between periods, without learning there is no link between the periods, and optimal first-period behavior is purely exploitative: that is, the maximizing the sum of discounted two-period profits degenerates into separately maximizing

¹¹To reduce clutter we do not label variables with a period index. In this section prices and locations refer to first period activity.

profits in each period based on the prior expectation $\hat{\mu}_0$ on the unknown slope γ . Consequently, the best response functions in the first period are the same as in the second period. Denoting first-period equilibrium price and location values for a firm that ignores learning by $\{\check{p}_i, \check{l}_i\}$, these values are the same as the second-period values: $\{\check{p}_1, \check{l}_1, \check{p}_2, \check{l}_2\} = \{p_1^*(\hat{\mu}_0), l_1^*(\hat{\mu}_0), p_2^*(\hat{\mu}_0), l_2^*(\hat{\mu}_0)\}$. Therefore, from the results of section 3.1 we know that a unique and symmetric first-period equilibrium exists for these non-learning firms.

PROPOSITION 5. *A firm that takes the value of learning into account in the first period will choose a location further from the sweet spot, and a price higher than would a firm that ignores the value of learning. That is, for $i = 1, 2$,*

$$\begin{aligned}\hat{p}_i &> \check{p}_i \\ |\hat{l}_i| &> |\check{l}_i|.\end{aligned}$$

The consequence is that consumers will face more product diversity, but higher prices, in an information goods market described by our assumptions.

3.3 Consumer Welfare

We analyze the effect of the learning process on consumers for two reasons. First, while the process results in a short run increase in market power for the firm, it also results in an increase in product diversity and in the number of consumers served, so that aggregate consumer welfare may actually increase. Second, understanding the effect on consumer welfare sheds light on the manner in which firms conduct their learning. We find that even within the range of beliefs about γ where the symmetric equilibrium holds, their beliefs about the attractiveness of the niche will determine how much competition is relaxed for a given experiment.

The effect of the learning process on consumer welfare is a complicated affair. Looking solely at the first period effect, how an individual consumer fares will depend on her type. Those located near the sweet spot will face higher prices and products less tailored to their tastes when firms locate further apart but raise their prices. Consumers located to the outside of the firm locations will receive a more desirable product, albeit at an increased price. Finally, as the number of consumers served increases in a learning environment, these new consumers clearly benefit. In this section, we look to resolve some of this ambiguity.

Using notation developed in Section 3.2, we define the no-learning equilibrium prices and locations as \check{p}_i and \check{l}_i . Making use of the symmetry of equilibrium demands, we look solely at the expected surplus of consumers to the right of the sweet spot. Their expected surplus in the no-learning case, $\check{C}S$ is:

$$\check{C}S = \int_0^{\check{t}_r} (\alpha - \hat{\mu}l)(r - c|\check{l}_2 - l| - p_2)dl. \quad (11)$$

We now look at how consumer surplus changes as firms increase their reach. As we show in section 5 at equations (17)-(19), for any given reach, $\tilde{t} = t_r = -t_l$, we can write the profit maximizing prices and locations as follows:

$$\tilde{p}_2 = c\tilde{t} - \frac{\tilde{t}^2 c \hat{\mu}}{2\alpha} \quad (12)$$

$$\tilde{l}_2 = -\frac{r}{c} + 2\tilde{t} - \frac{\tilde{t}^2 \hat{\mu}}{2\alpha}. \quad (13)$$

Armed with a characterization of the equilibrium prices and locations as firms increase their reach, we can now gauge their effects on consumers. In equation (11), we change the outer bound of integration to \tilde{t} , and substitute for prices and locations as detailed in equations (12) and (13). Differentiating with respect to \tilde{t} gives us:

$$\partial \check{C}S / \partial \tilde{t} = \frac{3}{128} \left(16r\alpha + \frac{64r^2 \hat{\mu}}{c} - \frac{47c\alpha^2}{\hat{\mu}} \right). \quad (14)$$

There is thus a region where consumer surplus is increasing in expectation, and one in which it is decreasing. We can solve equation (14) to find the threshold, which we shall call $\ddot{\mu}$,

$$\ddot{\mu} = \frac{c\alpha(4\sqrt{3}-1)}{8r}, \quad (15)$$

where consumer surplus is increasing in expectation for $\hat{\mu} > \ddot{\mu}$. Our threshold is approximately at the midpoint of $\left[\frac{5\alpha c}{8r}, \frac{7\alpha c}{8r}\right]$.

We can gain some insight towards interpreting this condition by looking at how a firm changes price as it changes its reach. In Section 5, we show that $\partial \tilde{p}_2 / \partial \tilde{t} = c(1 - \tilde{t}\hat{\mu}/\alpha)$. Thus the greater $\hat{\mu}$, the smaller any price increase for a given “unit” of learning (i.e. change in reach). Likewise, the greater $\hat{\mu}$, the less firms move their locations towards the tails. We can thus see how the manner in which firms experiment is affected by their beliefs. Even though firms desire to explore the niches, if their beliefs about the attractiveness of the tail are “pessimistic” enough, the mass market is more worth fighting over, and this moderates their move towards the niche for the sake of learning.

Thus, for low $\hat{\mu}$ (more valuable niches) firm learning reduces consumer welfare, while for higher $\hat{\mu}$ (more competition for the mass market) learning increases consumer welfare.

4. DISCUSSION

Rather than charge prices or differentiate goods to maximize current expected profits, firms may choose different prices or product configurations in order to create better experiments to improve their estimates of consumer preferences. Experimentation is usually thought to be undertaken at the expense of short-run profits. We have shown that this need not be the case. In a model of competition under uncertainty, in which firms have the ability to decrease direct competition, firms’ desire to resolve uncertainty can lead to short-run profits higher than would be the case if firms did not care about subsequent periods.

In our model of endogenous product differentiation with uncertainty about consumer preferences, firms are trying to learn the rate at which consumer preferences fall off away from the “sweet spot”, in order to choose the right balance between competing with low prices for the mass market of consumer and competing with higher prices for a niche market. What firms learn in the first period about the distribution of consumer preferences changes their second-period price and product configuration choices. Thus, first-period price and configuration choices affect expected second-period profits.

The amount of learning that a firm desires to undertake (at the cost of foregone current profits) depends on the convexity or concavity of the profit function in the firm’s belief about the unknown parameter, $\hat{\mu}$. Due to Jensen’s In-

equality, a concave utility function induces risk aversion: the decision-maker prefers a given value with certainty to a gamble with the same expected payoff.

In our model of endogenously-differentiated information goods future profits are convex in beliefs about γ . Consequently, firms prefer a gamble to that gamble's expected payoff, and they are willing to alter first period actions to gamble on what they will learn about γ . In order to learn more about γ , firms set first-period prices and product configurations to better explore the tails of consumer space than they would if they ignored the opportunity to learn.

This "risk loving" behavior arises even though the firms are by assumption risk neutral. For given prices and locations, profits are linear in γ , and firms thus maximize profits based on the expected value of γ . Only when the true state of the world is $\hat{\mu}$, however, are the firms actually choosing the optimal actions for that state. The effect of γ on profits is not linear because if the firms "knew" that the state of the world was not $\hat{\mu}$, they would choose better actions. Firms thus desire better information about what γ actually is, i.e. have a desire to experiment, so that they can better tailor their actions to the actual state of the world.

The manner in which a firm's desire to experiment affects prices and product configurations is relatively straightforward. As a firm's "reach" increases its sales are more affected by the value of γ , and the stochastic component of demand becomes relatively less important. Thus, locating further away provides a more informative experiment.

While it is true that firms could decrease price to increase demand and thus increase the informativeness of the experiment, we actually see the opposite effect on prices in this model. To understand why, consider a given experiment, which is to say an expansion of the outer bounds of consumers who buy (t_l and t_r). While a firm could serve this customer base by lowering its price, it could also move further away from its direct competition, which allows it to raise price. Clearly the latter strategy is superior, as it allows the firm to serve the same number of customers at a higher price.

Our main result suggests that when there is uncertainty about consumer preferences for information goods, there will be substantial experimentation in the form of product diversity. This seems consistent with casual observation of the past several years of commerce in information and other electronically transacted goods. With many new goods and services uncertainty about preferences has been high. Correspondingly the rate of introduction of new products and differentiation amongst them has been quite high.

Whether prices have been high or low for new products is not as obvious. In some markets the desirability of charging higher prices has perhaps been mitigated by other factors, such as the desire to build a brand reputation or to lock in customers. However, evidence from Bailey [3] suggests that as new firms entered in various electronic commerce markets, prices increased.

Whether our results are robust requires further investigation. For example, we are currently analyzing the dynamics that arise in the asymmetric equilibria that arise when γ is lower than the range studied in this paper (when niches are expected to be either more or less desirable). There are other directions in which one could generalize our model. For example, firms might be heterogeneous in one of several ways: they might start with different beliefs, or they might start at

different locations in product space and have nonzero costs of re-location. We also might learn more from a model in which there are multiple dimensions along which products can be differentiated, or in which there are more than two firms that sell imperfectly substitutable information goods (or in which each firm can sell multiple different goods). We also wonder whether the effect of valuable learning opportunities on pricing and product differentiation would be the same if there were more than one unknown parameter of the consumer preferences distribution. For example, a firm might not know the slope γ and also might not know the disutility cost c consumers incur as offered product configurations get further away from their most preferred product.

In a series of papers ([12, 4, 14, 13, 5]) we and our co-authors have studied the out-of-equilibrium behavior of software agents that search price and product configuration spaces under uncertainty about consumer preferences. In those papers the agents representing firms selling information goods face environments too complex to explicitly solve for optimal strategies even in a single firm environment. Instead, they pursue various search heuristics. We adopted relatively generic (uninformed) search heuristics due to the relative paucity of prior literature on the theory of optimal product and price configuration in an information goods environment. The results in the present paper, by characterizing some of the properties of optimal learning strategies in a particular setting, provide guidance for the design of informed search strategies in more complex (and thus realistic) settings. In separate research, we are pursuing the implications of the present paper for computational analyses of behavior off the equilibrium path.

5. APPENDIX: PROOFS OF PROPOSITIONS

PROOF. 1 Assume that this is not the case in the second period. Then there are unserved consumers located between the two firms. Without loss of generality assume that some of these consumers are of type $t > 0$. Due to the fact that consumer density decreases as we increase t , the number of consumers at the right boundary of the leftmost firm ($\alpha - \hat{\mu}(l_2 - \frac{r-p_2}{c})$) is greater than the number of consumers at the leftmost boundary of the rightmost firm ($\alpha - \hat{\mu}(l_2 + \frac{r-p_2}{c})$) for all γ . Therefore the rightmost firm can profitably deviate by moving its location to the left, so this cannot be a pure strategy equilibrium. The same logic holds if some of these consumers are of type $t < 0$.

The proof is similar for period 1. In addition to the increase in profits in period 1, the deviation also increases expected profits in period 2. To see this, note that from Proposition 2 the deviation decreases that variance of the estimator. As expected second period profits are convex in the expectation of γ , a more informative experiment increases expected profits.

PROOF. 2 The result is transparent: the variance of the estimator is the variance of ε divided by a constant, which is σ_ε^2 divided by the squared constant.

PROOF. 3 From equations (1) and (2) we have that aggregate demand is continuous in $\{p_1, p_2, l_1, l_2\}$, and thus that the value function is continuous in the same arguments. Since $p_i \in [0, r]$ and $|l_i| \in [0, \frac{\alpha}{\gamma}]$, ϕ_i exists by the Weierstrass Theorem.

PROOF. 4

We first prove that if an equilibrium exists it is symmetric. Note that the expected second-period profit function $\delta W(p_1, p_2, l_1, l_2)$ is the same in the value function for both firms, and symmetric because second-period profits are symmetric. Other than this additive expression, the first-period value functions are identical to the second period expected profit functions. Thus, a firm's first-order conditions are identical in the two periods except for the addition of a partial derivative of δW with respect to the choice variable of interest; that partial derivative will be symmetric for the two firms because the function W is symmetric. Therefore, if a solution to the system of first-order conditions exists, a symmetric solution must exist.

We now show the existence of the equilibrium. We first note that firms affect expected second period profits, $W(\cdot)$, solely through their choice of t_r and t_l . As any prices and locations yielding the same t_l and t_r are equally informative, prices and locations will be such that they maximize current profits for a selected t_r and t_l . For any $\tilde{t} = t_r = -t_l$, there exists a unique price and location pair $(\tilde{p}_i, \tilde{l}_1 = -\tilde{l}_2)$ that maximizes current period profits. These values are given by¹²:

$$\tilde{p}_1 = c\tilde{t} - \frac{\tilde{t}^2 c \hat{\mu}}{2\alpha} \quad (16)$$

$$\tilde{p}_2 = c\tilde{t} - \frac{\tilde{t}^2 c \hat{\mu}}{2\alpha} \quad (17)$$

$$\tilde{l}_1 = \frac{r}{c} - 2\tilde{t} + \frac{\tilde{t}^2 \hat{\mu}}{2\alpha} \quad (18)$$

$$\tilde{l}_2 = -\frac{r}{c} - 2\tilde{t} - \frac{\tilde{t}^2 \hat{\mu}}{2\alpha}. \quad (19)$$

We can thus characterize maximal one-period expected profits for any \tilde{t} as follows:

$$E[\pi_i | \tilde{t}, \hat{\mu}] = \frac{c\tilde{t}^2(\hat{\mu}\tilde{t} - 2\alpha)^2}{4\alpha}$$

The firms' maximization problem thus reduces to finding \hat{t} to maximize total discounted expected profits. A firm's first period value function is thus:

$$V(\tilde{t}) = \frac{c\tilde{t}^2(\hat{\mu}\tilde{t} - 2\alpha)^2}{4\alpha} + \delta W(\tilde{t})$$

and symmetric equilibrium is $\hat{t} \in \arg \max_{\tilde{t}} V(\tilde{t})$. Existence of \hat{t} follows the same reasoning as presented in preceding proof.

PROOF. 5

Define \tilde{t} as the "reach" that maximizes expected first period profits, i.e. $\tilde{t} = \tilde{l}_2 + \frac{r-\tilde{p}_2}{c}$, and $\hat{t} \in \arg \max_{\tilde{t}} V(\tilde{t})$. The informativeness of first period prices and locations are increasing in \tilde{t} . Thus, due to the convexity of $E[\pi | \hat{\mu}]$ in $\hat{\mu}$, $W(\cdot)$ is increasing in \tilde{t} . This combined with the fact that expected first-period profits are less than $\frac{9c\alpha^3}{64\hat{\mu}_0^2}$ (i.e. the best myopic profits) for all $t < \tilde{t}$ implies that $\hat{t} \geq \tilde{t}$.

To see the direction in which prices and locations move as we increase \tilde{t} , we differentiate equations (17) and (19) with respect to \tilde{t} and get

¹²Simple algebraic substitution reveals that the solution to the one-period maximization problem given by equations (3)-(6) is the solution to equations (16)-(19) for $\tilde{t} = \tilde{l}_2^* + \frac{r-\tilde{p}_2^*}{c}$.

$$\partial \tilde{p}_2 / \partial \tilde{t} = c \left(1 - \frac{\tilde{t} \hat{\mu}}{\alpha} \right) > 0 \quad \forall \tilde{t} < \frac{\alpha}{\hat{\mu}} \quad (20)$$

$$\partial \tilde{l}_2 / \partial \tilde{t} = 2 - \frac{\tilde{t} \hat{\mu}}{\alpha} > 0 \quad \forall \tilde{t} < \frac{\alpha}{\hat{\mu}}. \quad (21)$$

Thus, $\hat{p}_i \geq \tilde{p}_i$ and $|\hat{l}_i| \geq |\tilde{l}_i|$.

6. REFERENCES

- [1] Philippe Aghion, Patrick Bolton, Christopher Harris, and Bruno Jullien. Optimal learning by experimentation. *Review of Economic Studies*, 58(4):621–54, July 1991.
- [2] Simon P. Anderson, Andre de Palma, and Jacques-Francois Thisse. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, Massachusetts, 1992.
- [3] Joseph P. Bailey. Electronic commerce: Prices and consumer issues for three products: Books, compact discs, and software. Technical Report OCDE/GD(98)4, Organisation for Economic Co-Operation and Development, 1998. Available from <http://www.oecd.org/dsti/sti/it/ec/prod/ie98-4.pdf>.
- [4] Christopher H. Brooks, Edmund Durfee, and Rajarshi Das. Price wars and niche discovery in an information economy. In *EC'00: Proceedings of the Second ACM Conference on Electronic Commerce*. ACM Press, October 2000.
- [5] Christopher H. Brooks, Scott Fay, Rajarshi Das, Jeffrey K. MacKie-Mason, Jeffrey O. Kephart, and Edmund Durfee. Automated strategy searches in an electronic goods market: Learning and complex price schedules. In *EC'99: Proceedings of the ACM Conference on Electronic Commerce*. ACM Press, November 1999.
- [6] Karen Clay, Ramayya Krishnan, and Eric Wolff. Pricing strategies on the web: Evidence from the online book industry. In *EC'00: Proceedings of the 2nd ACM Conference on Electronic Commerce*. ACM Press, October 2000.
- [7] Neveen Farag and Marshall Van Alstyne. Information technology—a source of friction?—an analytical model of how firms combat price competition online. In *EC'00: Proceedings of the 2nd ACM Conference on Electronic Commerce*. ACM Press, October 2000.
- [8] J. J. Gabszewicz and J.-F. Thisse. Spatial competition and the location of firms. In J. J. Gabszewicz, J.-F. Thisse, M. Fujita, and U. Schweizer, editors, *Fundamentals of Pure and Applied Economics. Volume 5: Location Theory*. Harwood Academic Publishers, Chur, Switzerland, 1986.
- [9] Sanford J. Grossman, Richard E. Kihlstrom, and Leonard J. Mirman. A bayesian approach to the production of information and learning by doing. *Review of Economic Studies*, 44(3):533–47, Oct. 1977.
- [10] Joseph Harrington. Experimentation and learning in a differentiated-products duopoly. *Journal of Economic Theory*, 66:275–288, 1995.
- [11] John H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI, first edition, 1975. Second edition (1992), published by MIT Press, Cambridge, MA.

- [12] Jeffrey O. Kephart, Christopher H. Brooks, Rajarshi Das, Jeffrey K. MacKie-Mason, Robert Gazzale, and Edmund Durfee. Pricing information bundles in a dynamic environment. In *EC'01: Proceedings of the Third ACM Conference on Electronic Commerce*. ACM Press, October 2001.
- [13] Jeffrey O. Kephart, Rajarshi Das, and Jeffrey K. MacKie-Mason. Two-sided learning in an agent economy for information bundles. In Frank Dignum and Carlos Sierra, editors, *Agent-Mediated Electronic Commerce*, Lecture Notes in Computer Science. Springer, Heidelberg, Germany, 2001.
- [14] Jeffrey O. Kephart and Scott A. Fay. Competitive bundling of categorized information goods. In *EC'00: Proceedings of the Second ACM Conference on Electronic Commerce*. ACM Press, October 2000.
- [15] Jeffrey K. MacKie-Mason, Scott Shenker, and Hal Varian. Network architecture and content provision: An economic analysis. In Gerald Brock and Greg Rosston, editors, *The Internet and Telecommunications Policy*. Lawrence Erlbaum Associates, Mahway, NJ, 1996. Available from <http://www-personal.umich.edu/~jmm/papers/tprc.pdf>.
- [16] X. Martinez-Giralt and D. J. Neven. Can price competition dominate market segmentation? *Journal of Industrial Economics*, 36:431–442, 1988.
- [17] Andrew McLennan. Price dispersion and incomplete learning in the long run. *Journal of Economic Dynamics and Control*, 7(3):331–47, Sept. 1984.
- [18] Arie Segev and Carrie Beam. Broker strategies in electronic commerce markets. In *EC'99: Proceedings of the ACM Conference on Electronic Commerce*. ACM Press, November 1999.