Managing Change and the Success of Niche Products

Kieron J. Meagher

SANTA FE INSTITUTE

Managing Change and the Success of Niche Products*

Abstract

This paper presents a model of a firm preparing to launch a product into a market where consumer preferences change over time and cannot be directly observed. Market research is used in order to determine what type of product to produce. The employees making decisions based on the market research data are of bounded rationality, hence organizational structure plays a key role in the analysis of decision making. The implications for profits of the interactions of organizational structure and the stochastic environment are analyzed. Particular attention is paid to the choice of organizational size and profitability for niche appeal products and mass appeal products.

Kieron J. Meagher
Economics Program, RSSS,
Australian National University,
Canberra, ACT 0200, AUSTRALIA.
Fax: + 61-6-249-0182.
Email: kmeagher@coombs.anu.edu.au

1   Introduction

This paper develops a model of management structure and function for a firm which is developing a product for a predetermined launch date. The example which typifies the situation being considered is the development of a feature film.

There are standard distribution deals for many films and in the short term the price that consumers pay for different films remains fixed. Hence film producers compete not on price but the characteristics of their films. The preferences of the film going public are not directly observable and change over time. However market research can and does influence the content of films during the production process.

This paper attempts to address two main questions motivated by this example. What kind of management structure should a firm in this kind of market adopt in order to try and maximize profits? What effects do differing management structures have on performance?

The product is one that can vary in its characteristics, hence the firm has to pick a combination of characteristics for their product which they believe will maximize their returns. The preferences of consumers change over time and are not directly observable. To describe this situation a standard model of one dimensional preferences is expanded to a dynamic stochastic setting where a firm consists of boundedly rational agents. In particular, agents are limited in the amount of computation they can perform in each time period.

As a result of the stochastic nature of the problem, a firm is forced to make a decision on what product to produce, based on a forecast of the unknown market conditions at the launch date. Since the firm has no competitors it simply decides to locate at which ever point its forecast tells it will be the centre of the market when the product is launched.

It is shown in [5] and [7], that the time to calculate a forecast is an increasing function of the amount of data considered and decreasing, bounded function of the number of people employed in making the decision. By applying earlier work on hierarchies to the stated economic problem, this paper has succeeded in endogenizing the cost of delay (the time taken to calculate the forecast).

Each decision making procedure (method for producing a forecast) corresponds to a different management structure and associated management function, which are modeled here as information processing hierarchies. This methodology allows analysis of the minimum delay a certain organizational structure takes to make a decision based on a given quantity of market research.

The principal intuition of the paper is, increasing sample size gives a more accurate estimate of a historical situation, however it also increases delay making this estimate a more out of date basis for predicting the future. In formalizing this intuition it also shown that the optimal size of a firms management structure depends on the speed of change that occurs in the product market and also, surprisingly, on the size of the firms market niche.

2   The Formal Model

2.1   The Market.

The traditional location model of competition in differentiated products is modified to become a dynamic stochastic model.We begin by describing the situation during period t, and

then present the process by which market conditions change over time.

Consumers derive satisfaction from the consumption of the intrinsic characteristics of the goods they purchase. For convenience we assume that the goods under consideration has 1 relevant characteristic, which can be represented by a real number. A good produced by firm k can be represented by the unique point x_k on the real line (the real line is referrd to as the charcteristic space). For ease of exposition, a good is generally identified with the point representing its characteristics, so that the good itself is said to be "located" at the point on the real line.

Individuals differ in their preference for characteristics. Each individual has a unique ideal type of good, which is also represented by a point, c_i, also on the real line. Consumers prefer products that are closer to their ideal points. In particular an individuals utility from a product is strictly decreasing in the distance between the product and the individual in the characteristic space (similar to goods we refer to consumers as being located in the characteristic space at their ideal point).

Individuals differ in their location and are represented as a continuum of consumers, whose locations are uniformly distributed over the interval $[M(t) - \frac{1}{2}, M(t) + \frac{1}{2}]$. Consumers have unit demand and buy which ever product gives them the highest utility, which will be the closest product to their ideal point. They only buy if the utility from a product satisfies some reservation utility constraint (which would be determined exogenously by outside options).

We assume that the reservation utility level u^* is the same for each consumer and that the utility functions are identical functions of the distance of a product from an individuals ideal point. The utility for consumer i with ideal point c_i, from good x_k is given by

$$U_i(x_k, c_i) = U(x_k, c_i) = V(|x_k - c_i|) \tag{1}$$

where $V(|x_k - c_i|)$ is strictly decreasing in $|x_k - c_i|$.

Notice that the utility function of each individual is symmetric about that individual's ideal point. Much of the following could be generalized to non symmetric utility functions. The functional form in 1 is used for its expositional ease, and because the main focus of this paper is to elucidate the inter-relationship between internal information processing and market behavior in as general and concise a way as possible. The complexities of more specific functional forms may well be needed when this model is fitted to data.

In the rest of paper we will consider a firm in isolation. The firm will have buyers whose ideal points satisfy $V(d_{k,i}) \geq u^*$. Since the utility function is strictly decreasing and symmetric it follows that there exists some $r \geq 0$, such that consumer c_i will only purchase good x_k if $|x_k - c_i| \leq r$. Hence the potential customers for product x_k are those that lie in the closed interval $[x_k - r, x_k + r]$, referred to as the product interval of appeal, or the firm's interval (as opposed to the market interval).

Consumers are represented by a probability distribution (with density $f(z)$) over the characteristic space. Integrating the density function over the interval will give the proportion of the total population of consumers who would buy product x_k if it were the only product on the market. Since consumers are uniformly distributed on $[M(t) - \frac{1}{2}, M(t) + \frac{1}{2}]$, the proportion of consumers purchasing x_k (referred to as the market share $s(x_k, t)$) is given by

$$s(x_k, t) = \int_{x_k - r}^{x_k + r} f(z) dz = \mu([x_k - \frac{w}{2}, x_k + \frac{w}{2}] \cap [M(t) - \frac{1}{2}, M(t) + \frac{1}{2}]). \tag{2}$$

2

Where $\mu(.)$ is the measure of a set (which is simply the length on the overlap in this case) and $w = 2r$.

It should be noted that this is a share of the potential market, not a firm's share of actual sales in the market. Hence this definition of market share differs from the one in common use.

The midpoint $M(t)$ of the consumer distribution is assumed to be a function of time. For expositional ease we assume that the data generating process for $M(t)$ is given by the random walk

$$M_t = M_{t-1} + \varepsilon_t. \tag{3}$$

with $\varepsilon_t \sim N(0, \sigma^2)$.

2.2 The Firm

A firm has to decide where to locate its product in order to maximize expected profits. The launch date is assumed to be fixed in advance. The management problem is to attempt to maximize profits by choosing a sample size, a time frame for processing the sample and an organizational structure with which to process the sample. These choices are not independent and are determined by the production and management technology available to the firm.

The technological situation a firm faces is considered next.

2.2.1 Decision making by the Firm

This is a model of product development for a fixed launch date in the face of evolving consumer preferences. Firms only gain information about what is happening in the market by surveying consumers. The product location decision (i.e. which x_k to produce) must be made prior to launch so there is no sales data from the product on which to base inference about the market. Instead a firm must go out and conduct some form of market research if it is to locate and track the population of consumers in the characteristic space. This market research might be in the form of questionnaire surveys or product pre-release tests. For example, feature films often have test screenings and their endings or offensive scenes are changed depending on viewer reactions. Producers of computer software generally distribute beta versions of programs to test out new features before release.

In order to capture as large a market share as possible a firm really needs to know about the shape and location of the market when it releases its product. Conceivably, firms might be using sample information to infer the distribution of consumer preferences, the form of the process by which this distribution changes over time and the history of shocks that have a persistent influence on the location of the market. Similar to most of the literature on monopolists learning about their demand conditions, we simplify the problem a firm faces by assuming that it knows all the relevant functional forms and the values of all coefficients and variables, except the history of actual outcomes of $M(t)$'s and the ε_t's, which are the shocks which occur to the market.

Firms will be able to guess the correct location of $M(t+1)$ in period t with probability 0 because there is a white noise shock term. However, a firm will get a positive payoff as long as its product interval overlaps with the market. The market transition process, Equation

3, is autoregressive, hence knowledge about past locations of the market will be useful in producing an estimate of where the market will be located at launch time.

A firm is interested in acting to maximize expected profit. The profit from launching product x_k in period t is denoted by $\pi(x_k, t)$. The profit will be the revenue, $R(x_k, t)$, generated from the market share the product captures, less the costs. Costs are of two types: a constant marginal cost of production c and a fixed cost, F.

Assume that price p and the size of the market Q are constants. Then the quantity sold is $Qs(x_k, t)$, and $R(x_k, t) = pQs(x_k, t)$. Hence expected profit is given by

$$
\begin{aligned}
E[\pi(x_k, t)] &= E[R(x_k, t) - cQs(x_k, t) - F] \qquad (4) \\
&= (p - c)QE[s(x_k, t)] - F.
\end{aligned}
$$

Potentially there are two sources of fixed cost, fixed costs in the production of the good and management overheads. In this model, management overheads could be of two types: the cost of collecting information and the labor cost of processing that information in order to make a decision. We are concerned here with how internal structure and the behavior of firms are related, hence the labor cost of processing the information is of more interest. However no method for endogenizing wages (and hence labor cost) has yet been devised for this type of model. An arbitrary choice of wages offers little insight so the whole question of labor cost has been left as an avenue of future research. Thus the approach used here is more like [8] than [5], or [7], although the underlying structure of all these approaches is the same.

The limited capabilities of a firms employees are explicitly modelled by describing each individual as a processors and the organization as a programmed network (which is equivalent to a distributed memory P-RAM).

A processor has an *inbox* and a *register* (which together comprise its *memory*). In one period of time, a processor can perform all or some of the following operations in the following order:

1. Read a number from the sample or its inbox and store it in its register, overwriting the previous contents of the register.

2. Take one number, y, from its inbox, calculate a linear function y and the contents of its register and store the result in its register;

3. Write the contents of its register to the inbox of any processor to which it is connected.

Together processors and connections form a network. A program specifies which operations each processor performs in each period.

Although specified here as a linear function, so that the network can calculate a sample mean, the generalized form of this model only assumes that processors can perform an associative binary operation. The adding of numbers should be considered as an analogy of the processing of information and the producing of reports, recommendations and decisions that really go on inside the management structure of an organization. It might seem a trivial task to collate sample information, however the information which flows within this programmed network should also be thought of as the decisions and recommendations that are based on raw information.

4

3 The Single Sample Approach

Potentially a firm might choose any number of samples and any decision rule based on these in order to determine where to locate its product.

The simplest possible approach is to take no sample at all and locate at the point which the entrepreneur believes will maximize profits. If the process by which consumer preferences change is stationary, then there exists a long run mean to preferences and this would be the sensible place at which to locate if no other information is used. This approach is likely to do well if the shifts in the market are small compared to the appeal of the product. This approach makes a useful bench mark for comparison in stationary markets.

If, as assumed here, the market is non-stationary then by definition there does not exist a mean for the market transformation process which is independent of time. Thus there is no sensible reason for picking any point without sample information, since the expected profit for a randomly drawn location is zero.

The next simplest case is to take a single sample and to make all decisions based upon this. Although perhaps not the optimal sampling approach if sampling is costless and unconstrained, this approach is intuitively plausible.

Sampling is not in general free and may have sizable fixed or marginal costs limiting the size and number of surveys taken. In addition there are institutional factors that may well constrain the number of surveys. The surveying process has not been modelled here, it just provides a number of data points which are to be processed to make a decision. The most common paradigm for market research is some kind of questionnaire based survey. It is however important to consider whether the survey questions are general or based on responses to a specific experience.

Returning to the feature film example, a general question such as "How much violence do you like in a film?" is going to be of limited use to film producers trying to decide if their film has too much violence. Instead they would like to know, in response to seeing their film, how many people thought it contained too much violence, and which particular scenes were disturbing. In order to gain this information pre-release showings have to be arranged. It is not just the direct cost of re-editing and printing the film and organizing a showing that make these showings rare. It is important for products, like films, to hit the market with an up beat bang. Too much pretesting can indicate to consumers a product with a problem, the product looses its novelty and it gives more opportunity for competitors to copy or duplicate the product.

For these reasons a firm may well only use one sample, and then just has to decide how best to process it, and simultaneously with the choice of processing, how large the sample should be. This situation is close to the sampling problem of classical statistics, but in this model the cost associated with a sample is now endogenously determined by its impact of profits, rather than the cost of actually collecting the data.

We begin by assuming that firm k takes a sample of size $N(k)$ which is the only source of empirical information used in the product decision rule. This sample is used to calculate a sample mean and the firm relocates on the basis of this sample mean. The calculation of the sample mean is performed by the people who work for the organization, by forming some programed network.

3.1 An example: the regular binary hierarchy solution

Potentially, for an organization with P employees, it might take a $P \times P$ matrix (called the adjacency matrix) to represent all the communication links between employees. This is the situation where for each individual it is necessary to state whether they pass information to each of the other $P - 1$ individuals. In a management situation the number of connections is typically less than this. [3, p18] site a real world example of excessive information transmission in which 92 connections are shown between 17 committees. Working with the management principal "simpler organizational structure is better" theoretical work by economists on management structures has concentrated almost exclusively on hierarchies.

In order to reduce the number of parameters from $P(P - 1)$ to a more manageable level some simplifying assumptions on the shape of the network are generally made. It is generally assumed that the organization has a hierarchical form (it can be easily proved that given certain informational assumptions this will be the case, see for example [2]. Given that the organization is a hierarchy two more assumptions are generally employed to reduce the parameters to a manageable number, see for example [6], [9] or [1]. First it is assumed that the hierarchy has L levels, (with level 1 at the bottom and the boss at level L) and that individuals at all levels above 1 have subordinates, the number of subordinates is know as an individuals span of control. The second assumption is that either all individuals at the same level have the same span of control (hence reducing the hierarchy to L parameters) or that the span of control at each level is the same, reducing the hierarchy to two parameters (this is known as a regular hierarchy).

The binary hierarchy example is useful because its hierarchies are easy to visualize (unlike the optimal single sample hierarchies). The functions relating delay and sample size are also smooth, well behaved and can be parameterized by the single variable L.

In a binary hierarchy there is one person at level L, 2 people at level $L - 1$, and so on down so that there are 2^{L-1} employees at the bottom (level 1). Summing over the number of levels gives $\sum_{i=1}^{L} 2^{i-1}$ total individuals in the hierarchy.

It takes an individual two periods to add up two pieces of information. Regardless of whether they are subtotals received from another person in the hierarchy or raw data from the sample. Assume that each individual at level 1 receives 2 pieces of raw sample data then the hierarchy receives 2^L sample points and takes $2L$ periods to process them into a decision.

Figure 1 gives an example of a three level binary hierarchy processing the sample mean of the data set: $\{c_1, ..., c_8\}$. Individuals are represented by the shaped boxes. Each box shows the linear computation performed by that individual. The data points are shown being read in at the bottom of the hierarchy.

These assumptions reduce a firms organizational and sample decision to a single parameter, L. By choosing the number of levels in the hierarchy a firm also fixes the amount of information it can process and the amount of time it takes to do this processing and hence which period to take the sample in (since for equal sized samples, the most recent one will always be the most useful).

Next we need to consider the expected market share and hence expected profits associated with each sized hierarchy, and how these depend on the environment in which the firm operates.

Assume that the survey was taken in period t. The sample $C_t = \{c_1, \ldots c_N\}$ that a firm
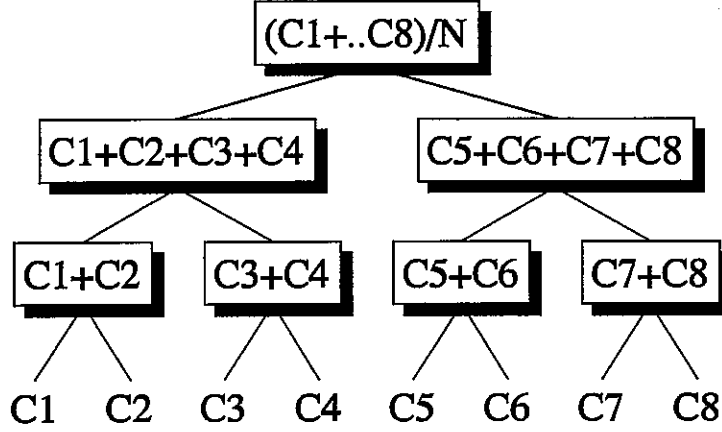
Figure 1: A three level hierarchy calculating the sample mean of eight numbers.

receives is the set of ideal points for the consumers who were surveyed. By calculating the sample mean $\overline{C}$ the firm has the best unbiased estimate of $M(t)$, which was the mean of the distribution from which the sample came. Since this sample is the only information on which to condition expectations it follows that

$$E[M(t)|C_t] = E[M(t+\delta)|C_t] = \overline{C}, \qquad \forall \delta > 0. \tag{5}$$

Expected market share is the expected overlap between the firms interval, which will be $[\overline{C} - r, \overline{C} + r]$ and the market interval $[M(t+\delta) - \frac{1}{2}, M(t+\delta) + \frac{1}{2}]$, where δ is the number of periods of delay caused by calculating $\overline{C}$. Since the market shifts according to a random walk with Normally distributed shocks, the market interval can be rewritten as

$$[M(t+\delta) - \frac{1}{2}, M(t+\delta) + \frac{1}{2}] = [M(t) + \sum_{i=1}^{\delta} \varepsilon_{t+i} - \frac{1}{2}, M(t) + \sum_{i=1}^{\delta} \varepsilon_{t+i} + \frac{1}{2}] \tag{6}$$

$$= [M(t) + T - \frac{1}{2}, M(t) + T + \frac{1}{2}]. \tag{7}$$

Where $T = \sum_{i=1}^{\delta} \varepsilon_{t+i}$, is the sum of the shock terms. The shock terms are independent and identically distributed, $\varepsilon_t \sim N(0, \sigma^2)$, hence $T \sim N(0, \delta\sigma^2)$.

The length of the overlap is a function of $d = \left| M(t) + T - \overline{C} \right|$, the distance between the centres of the two intervals. There are two cases depending on the size of w (recall $w = 2r$) the firm's interval width. The following analysis assumes $w \leq 1$, this is the economically more interesting case of a product whose width of appeal is no bigger than the market.

$$\text{Overlap}(d) = \begin{cases} w & if & 0 < d < \frac{1-w}{2} \\ \frac{1+w}{2} - d & if & \frac{1-w}{2} < d < \frac{1+w}{2} \\ 0 & if & \frac{1+w}{2} < d \end{cases} \tag{8}$$

7

To find the expected overlap the distribution of $d = \left| M(t) + T - \overline{C} \right|$ is needed. This is calculated easily from the distribution of $z = (M(t) + T - \overline{C})$. Let $y(z), q(d)$ be the respective probability densities of z and d. Then $q(d) = 0$ if $d < 0$ and $q(d) = 2f(z)$ otherwise.

Now $E[z] = E[M(t) + T - \overline{C}] = 0$ and $Var[z] = Var[M(t) + T - \overline{C}] = Var[T] + Var[\overline{C}]$. The variance of T is the sum of variances of the shock terms of which there are $2L$ and the variance of $\overline{C}$ is simply the variance of the sample mean of 2^L randomly chosen points from a Uniform distribution on $[M(t) - \frac{1}{2}, M(t) + \frac{1}{2}]$. Hence

$$Var[T] + Var[\overline{C}] = 2L\sigma^2 + \frac{1}{12 \times 2^L}. \tag{9}$$

Although the distribution of z is not known, it can approximated by a second order Hermite polynomial. This gives the approximation: $z \sim N(0, 2L\sigma^2 + \frac{1}{12 \times 2^L})$, where the density function $f(z)$ of the Normal distribution with mean zero and variance $2L\sigma^2 + \frac{1}{12 \times 2^L}$ is given by

$$f(z) \;=\; \frac{1}{\sqrt{2\pi(\delta\sigma^2 + \frac{1}{12N(k)})}} \exp\left(\frac{-z^2}{2(\delta\sigma^2 + \frac{1}{12N(k)})}\right) \tag{10}$$

$$=\; \frac{1}{\sqrt{2\pi(2L\sigma^2 + \frac{1}{12 \times 2^L})}} \exp\left(\frac{-z^2}{2(2L\sigma^2 + \frac{1}{12 \times 2^L})}\right) \tag{11}$$

Thus expected market share for $w \leq 1$ is

$$s(\overline{C}|L, w, \sigma) = \int_0^{\frac{1-w}{2}} 2w f(z) dz + \int_{\frac{1-w}{2}}^{\frac{1+w}{2}} 2\left(\frac{1+w}{2} - z\right) f(z) dz \tag{12}$$

This integral equation cannot be solved analytically. However the relationships between the variables can be examined graphically by solving the equation numerically which is the approach adopted below.

A difficulty arises in that market share is a function of three parameters, L, w and σ. Thus at least one of these parameters must be assigned a specific value in order to produce a graph. We begin by comparing two hierarchies with fixed numbers of levels, L. The cases $L = 1$ and $L = 6$ are shown in Figure 2 and Figure 3 respectively. These cases are of interest because they correspond to a single individual processing the sample ($L = 1$) and a rather deep hierarchy ($L = 6$). Many large companies have around six layers to their hierarchies, so by real world standards this type of hierarchy is on the big side in terms of levels (although the number of employees here is relatively small because only one task is being considered). The graphs show the market share these hierarchies expect to gain for ranges of w (the width of appeal of the firms product) and σ (the standard deviation of ε_t).

Both Figures show the same general shape. In both cases a larger market share is obtained when the standard deviation of the shocks to preferences are small. This occurs because when shocks are small there will be little difference between preferences in the period sampled and the period of the launch. If shocks are small it is not so much speed as accuracy which is important, hence the six level hierarchy which takes a larger sample does better when shocks are small. As shocks become large the firms expected market share drops to zero - there is just too much noise for it to expect its forecast to be accurate.
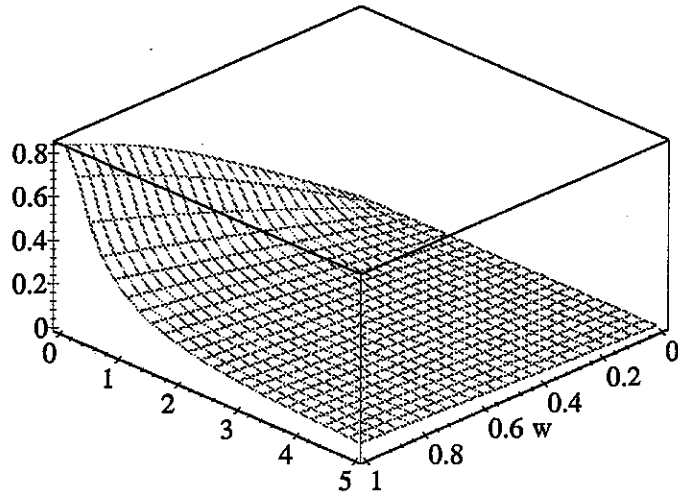
Figure 2: Expected market share for ($L = 1$). (Expected market share on vertical axis, σ on left axis, w on right axis).
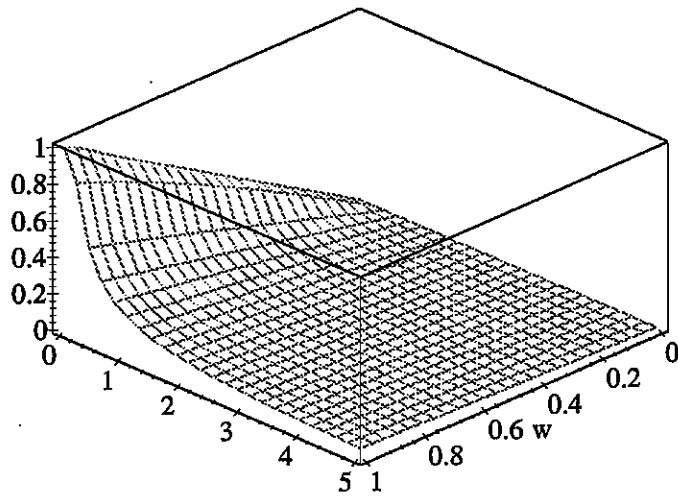

Figure 3: Expected market share for ($L = 6$). (Expected market share on vertical axis, σ on left axis, w on right axis).

As the width of appeal of the product increases so does the market share, as long as the shocks are not too large. Again this is intuitive since the market share a firm gains can be no larger than the width of appeal of its product.

The shapes of the two graphs do differ suggesting that there are potentially interesting interaction effects amongst the parameters. These are not clear from the three dimensional plots, hence below, two dimensional graphs for two different values of a third parameter are compared on the same plot, to make the effects clearer. We shall focus in particular on two types of products. A niche product ($w = 0.1$), which can gain at most 10% of the market and a mass appeal product ($w = 1$) which could gain 100% of the market.

Aiming for a niche is a favorite principle of modern management, see for example [3, p182-186]. The model used here allows rigorous analysis of the predicted performance of niche products relative to mass appeal products. Intuitively one might believe that only having a small market might leave a firm more exposed to shifts in preferences, we see that this is not necessarily the case.

First we compare the performance of a small hierarchy ($L = 1$) and a large hierarchy ($L = 6$), for both niche and mass products. The differences in expected market share for these two hierarchies are most shown in Figure 4. Each line shows the expected market share of a one level hierarchy less the expected market share of a six level hierarchy, for a specific type of product and for a range of shocks.

In both cases the large hierarchy does best when the shocks are small (negative values on the vertical axis) and the smaller (and hence faster hierarchy) does better for larger shocks. Note that counter intuitively the niche product is less effected in both absolute and relative terms by the size of the shocks. The niche product width is smaller than the market interval, hence by aiming for the middle of the market the niche firm gives its self plenty of room to miss. The mass appeal product has a wider interval so it always has a larger market share than the niche firm, but is consequently relatively more effected when its forecast is inaccurate

Next in Figure 5 we examine the optimal choice of hierarchy for a niche product when shocks are small ($\sigma = 0.05$) or a little larger ($\sigma = 0.1$). For small shocks the optimal hierarchy has four levels. As the shocks become larger the optimal size of hierarchy decreases (for the $\sigma = 0.1$ case, it has dropped to a two level hierarchy). When shocks are smaller the cost of having too large a hierarchy (or too small) are relatively low, shown by the flatness of the $\sigma = 0.05$ line. As the world becomes more uncertain it becomes more important to be the right size. This sits well with the intuition that firms in a very static environment can survive despite inefficient bureaucracies.

The situation for a mass appeal product, shown in Figure 6, is much the same. The mass appeal product is less effected by the increase in the size of the shocks, since the slopes of the two curves are more similar than in Figure 5. In part this is because the mass appeal product was less able to fulfill its potential market share even when shocks were small.

Instead of absolute performances we now consider relative performance.

$$\text{Expected percentage of potential market share} = \frac{\text{Expect market share} \times 100}{\text{Product width of appeal}}. \tag{13}$$

Recall, a product's width of appeal defines the maximin share of the market it could gain. Thus by this measure, if a firm which has $w = 0.1$ (the potential to gain 10% of the market)
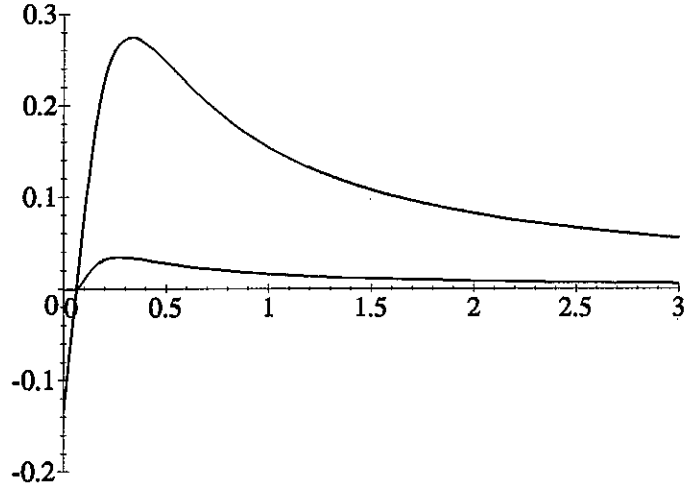
Figure 4: Expected market share of $(L = 1)$ less expected market share of $(L = 6)$, for niche product $(w = 0.1)$ and mass appeal product $(w = 1)$.(Difference in expected market share on vertical axis, levels in hierarchy on horizontal axis).
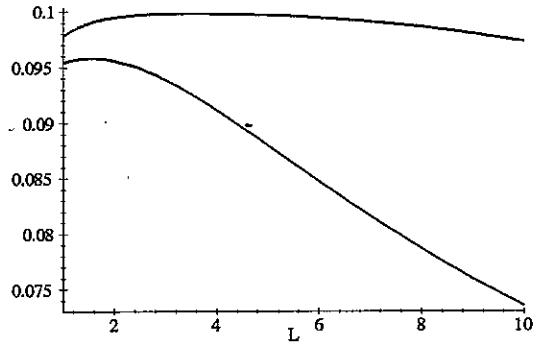

Figure 5: Expected market share by levels for niche product $(w = 0.1)$ for $\sigma = 0.05$ and $\sigma = 0.1$. (Expected market share on vertical axis, levels in hierarchy on horizontal axis).
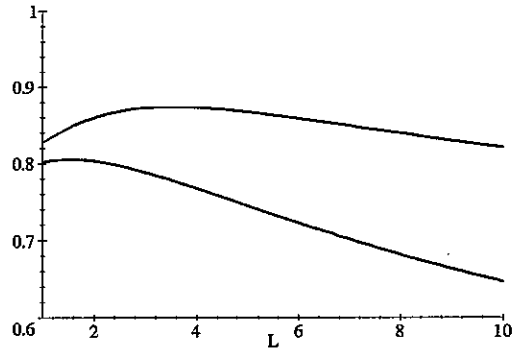
Figure 6: Expected market share for mass appeal product ($w = 1$) for $\sigma = 0.05$ and $\sigma = 0.1$. (Expected market share on vertical axis, levels in hierarchy on horizontal axis).

expects to get a market share of 0.1 then it has fulfilled 100% of its potential. This is an important real world consideration because it is the rate of return on investment, not the magnitude of profits, which is really important to investors. A niche product should go hand in hand with plans for a lower level of production and hence less investment.

Figure 7 shows that for a small shock, the niche product is expected to fulfill more of its potential than the mass appeal product. Also the niche products ability to fulfill its potential is reasonably insensitive to a suboptimal choice of hierarchy. However as the magnitude of the shocks increases the relative performance of the two products converge, see Figure 8.
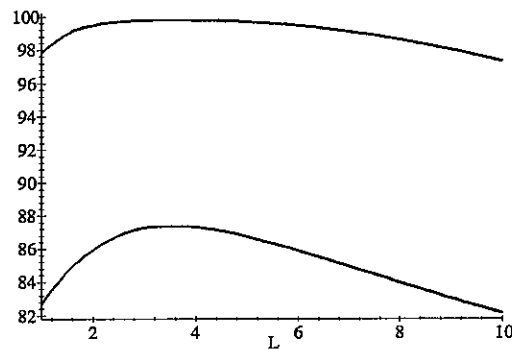

Figure 7: Expected percentage of potential market share for niche product ($w = 0.1$) and mass appeal product ($w = 1$) when $\sigma = 0.05$. (Expected percentage of potential market share on vertical axis, levels in hierarchy on horizontal axis).
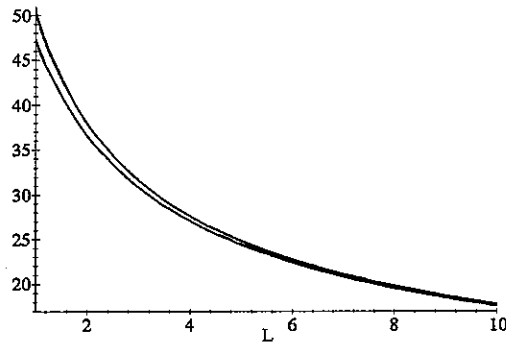
12

Figure 8: Expected percentage of potential market share for niche product ($w = 0.1$) and mass appeal product ($w = 1$) for $\sigma = 0.5$. (Expected percentage of potential market share on vertical axis, levels in hierarchy on horizontal axis).

4 Optimal Single Sample Hierarchies

The set of optimal hierarchies for a firm can be described as the set of hierarchies that give the minimum delay for a given sample size. [5] provides an algorithm for converting regular binary hierarchies into hierarchies with minimum delay for a given sample size and a given number of employees.

The details of the algorithm will not be repeated here, but the intuition for the method is as follows. In a regular binary hierarchy the supervisors of the level which is currently processing are idle. Time could be saved if some of those in the current level were eliminated and their work passed to their supervisors. Time is saved in this way because there is an implicit cost in rereading information every time a report is passed upwards. This can be somewhat alleviated by eliminating some unnecessary reports (and the individuals who authored them). Repeating this procedure until no more improvements can be made gives an optimal solution.

These optimal solutions give the minimum delay for a sample of size N by P employees. The analysis in this paper, however, has ignored the labor cost induced by the number of employees in the various hierarchical solutions. Thus for our purposes we can focus on the number of employees that gives the overall minimum delay for a sample of size N. The minimum delay can always be achieved with $P > \frac{N}{2}$ employees, and is given by

$$\text{Minimum delay}(N) = \delta(N) = 1 + \lceil \log_2 N \rceil . \qquad (14)$$

Thus in these fastest hierarchies the number of employees can be approximated as half the sample size N, and the choice of hierarchy can again be reduced to one parameter, N. These hierarchies do not have a regular shape and are as readily interpreted in terms of levels as the regular hierarchies. However conclusions about the size of the management structure can still be drawn.

The labor cost of a hierarchy is of importance in deciding on an optimal hierarchy. Meagher(1995) outlines how to calculate labor cost for these hierarchies under various em-

ployment regimes and shows the impact of different employment regimes on the choice of optimal hierarchy.

Whether labor cost will have a significant impact on the choice of hierarchy depends on the wage bill relative to the costs associated with the accuracy of the forecast. Wages have not been endogenized in this model, neither have costs been calibrated with empirical data. Hence the impact of considering labor costs would depend on an arbitrary choice of wages. Thus it is assumed that wages are small relative to the potential loss in profits from not choosing the correct sampling and hierarchical solution, and hence the cost of labor can safely be ignored as a constraint in choosing the optimal hierarchy.

We can analysis the performance of these optimal hierarchies in different environments by substituting equation 14 into equation 10 to create the approximating Normal distribution

$$f(z) = \frac{1}{\sqrt{2\pi((1 + \lceil \log_2 N \rceil)\sigma^2 + \frac{1}{12N})}} \exp(\frac{-z^2}{2((1 + \lceil \log_2 N \rceil)\sigma^2 + \frac{1}{12N})}). \qquad (15)$$

This in turn is substituted into equation 12 to give an approximate equation for the expected market share of an optimal hierarchy processing a sample of size N. This can be analysed by numerical simulation as before.

5 Conclusion

The first part of this paper takes the standard characteristic space model of demand for differentiated products, and moves it a to a dynamic stochastic setting. Consumer preferences change over time and are not directly observable, hence the decision of which good to produce becomes significant and difficult.

Firms solve the problem of which good to produce by using sample data from the market. However the individuals using the sample data are of bounded rationality, hence it becomes important for the firm to choose some form of hierarchical organization so that coordination can be used to overcome the limits of the individuals rationality.

The model endogenizes the cost of delay in processing information. It followed from this that large hierarchies perform best in reasonably static environments when delay in making a decision was outweighed by the importance of accuracy.

It was also shown that the choice of hierarchy was less critical for a niche product firm than a mass appeal product firm when consumer preferences experience only small shocks. Also counter intuitively niche product firms were expected to be more successful at gaining their potential market share than mass appeal product firms. However, the expected success in gaining market share for the two types of products, converged as shocks to consumer preferences became large.

This theoretical model of market research has many obvious extensions (some of which are the subject of on going research), including multiperiod sampling, repeated play and multiple firms interacting strategically.

References

Beckmann,"Some Aspects of Returns to Scale in Business Administration", Quarterly Journal of Economics.

Meagher, Kieron J.,"Efficient Hierarchies: Equivalence under Differing Employment Regimes", unpublished working paper, 1995.

Peters and Waterman,"In Search of Excellence", Harper & Row Publishers Inc, New York, 1982.

Radner, "Hierarchy: The Economics of Managing", Journal of Economic Literature, 30 (1992), 1382-1414.

Radner, "The Organization of Decentralized Information Processing", Econometrica, Vol. 61, No. 5 (September, 1993), 1109-1146).

Rosen, "Authourity, control and the distribution of earnings", Bell Journal of Economics, Vol. 13 (1982), 311-323.

Van Zandt, "The structure and returns to scale of organizations that process information with endogenous agents", unpublished working paper, Princeton University, 1994.

Van Zandt and Radner, "Information Processing and Returns to Scale of a Statistical Decision Problem", unpublished working paper, 1995.

Yingyi Qian, "Incentives and Loss of Control in an Optimal Hierarchy", Review of Economic Studies, (1994), 61, 527-544.