

Customer Loyalty and Supplier Quality Competition

Noah Gans

*OPIM Department, The Wharton School, University of Pennsylvania,
Philadelphia, Pennsylvania 19104
gans@wharton.upenn.edu*

We develop a model of customer choice in response to random variation in quality. The choice model yields closed-form expressions which reflect the effect of competing suppliers' service quality on the long-run fraction of purchases a customer makes at the various competitors. We then use the expressions as the basis of simple normative models for suppliers seeking to maximize their long-run average profits. The results provide insight into the effect of switching behavior on the service levels offered by competing suppliers. (*Customer Loyalty; Quality Competition; Bayesian Bandit*)

1. Introduction

What should the fill rate be? How should we set the expected delay in queue? In traditional operations models, these types of service levels are exogenously defined.

In practice, companies commonly set quality levels by simply matching the competition. Benchmarking exercises and data drawn from trade associations allow managers to compare the service levels of their operations to those of their competitors. Still, managers may wonder whether they have converged on the right service level, whether there is something to be gained from deviating from the industry standard.

The "right" level of service depends on how customers respond to variations in suppliers' quality levels. In many cases, instances of poor service may lead customers to change suppliers, and this switching behavior drives supplier market share and profitability. Without an understanding of customer choice behavior and competitive dynamics, the question of how best to set quality levels cannot be properly addressed.

In this paper, we seek to improve our understanding by directly modeling these phenomena. We develop a model of consumers' long-term responses to quality failures that is intended to capture the effect

of variation in service quality on customer switching behavior. We then use the results of the consumer model as the basis of a model of quality competition in an oligopoly.

In our consumer model, a customer repeatedly chooses among a set of suppliers, and the outcome of each visit to a supplier is some (instantaneous) utility. The utility offered by each supplier is, in fact, a random variable that reflects the quality of that supplier's offering: whether or not the flight arrives on time, how fresh the fish is that day. The customer is not well informed about supplier quality levels, and she uses a crude form of Bayesian revision to keep track of which of the suppliers she prefers. Each time she enters the market, she myopically chooses the supplier that she thinks is most likely to be best.

The consumer model yields succinct closed-form expressions for the long-run fraction of purchases a customer makes at the various suppliers. Furthermore, when the appropriate random variables are members of exponential families of distributions, each supplier's fraction can be shown to be increasing and convex in its own quality level, and decreasing and concave in its competitors' quality levels. Thus, a supplier's "share of customer" responds increasingly strongly to changes in quality levels.

We then use these expressions for customer share to model the effect of oligopoly competition on the supplier's choice of quality level. We formulate a non-cooperative game in which unit prices are fixed and suppliers compete on the basis of the quality levels that they set. When aggregated, the switching behavior of individual customers obtains market shares that evolve dynamically and stochastically over time.

Given this stochastic view of market shares, we formulate the supplier's objective as the maximization of long-run average profit. For suppliers whose costs 1) are convex and increasing in the overall level of quality offered, and 2) do not display economies of scale, we find the following:

- When competitors are symmetric, there is a unique pure-strategy Nash equilibrium, and it is symmetric.
- When competitors have asymmetric costs, equilibria are always asymmetric. In a duopoly, the competitor with lower costs will invest its advantage to increase quality, capture market share, and earn higher long-run average profits.
- As the number of competitors m increases, competitive pressure drives quality levels to increase. Still, consumers' lack of quality information about providers allows suppliers to earn positive profits from current customers even as $m \rightarrow \infty$.

Thus the results demonstrate that, even though customers are not well informed about producers' quality levels, their ability to switch suppliers in response to poor service provides the discipline necessary to drive competitors to adopt an industry quality standard. Conversely, although the standard increases with increasing competition, consumers' lack of information is significant enough to permit suppliers to earn positive margins, even as $m \rightarrow \infty$.

Finally, we note that the types of optimization models used to derive these results can be used (in principle) by a supplier to verify whether or not the current standard in its industry is consistent with the long-run competitive equilibrium. To the extent that it is not, a supplier can use the models to better understand what its own "off-equilibrium" quality strategy should be.

The remainder of the paper is organized as follows. Section 2 reviews related literature. In §3, we describe

the supplier's decision problem, and in §4 we develop our model of customer reaction to random variation in service quality. Then §5 uses the customer model to characterize a supplier's share of a customer as a function of its and its competitors' quality levels. In turn, §6 integrates the customer share results into the supplier's problem to analyze quality competition among suppliers. Finally, §7 discusses the results, as well as directions for future work.

2. Literature Review

Our model of quality competition is related to two recent streams of papers in the operations management literature. The first group looks at quality competition more generally and (like this paper) represents the supplier's cost of service in a more highly stylized fashion. Recent examples of these papers include Karmarkar and Pitbladdo (1997), Cohen and Whang (1997), Tsay and Agrawal (2000), and the references therein. The second group analyzes queueing systems, and in them delay in queue or system sojourn time is typically defined to be the measure of quality, the shorter the better. Examples that analyze the case of a monopolist service provider include papers by Mendelson (1985), Mendelson and Whang (1990), Dewan and Mendelson (1990), Stidham (1992), and van Mieghem (2000). Mandelbaum and Shimkin (2000) characterize equilibrium behavior of customer abandonments from a monopolist. Other work, such as that of Li (1992), Li and Lee (1994), Lederer and Li (1997), Ho and Zheng (1996), Cachon and Harker (2001), and Chayet and Hopp (1999) perform competitive analyses in which a supplier competes within an oligopoly. Li (1992) analyzes a queueing-inventory system and obtains newsvendor-like expressions for which the "underage" cost reflects a balance between holding costs and the discount rate.

The analyses in both groups assume that in equilibrium consumers are well informed about (or have beliefs that are consistent with) the service level offered by a supplier: Each customer knows sufficient statistics concerning the service quality at the various suppliers. In Li (1992) and Li and Lee (1994) customers continually monitor queue lengths and jockey among suppliers; they exhibit no loyalty towards suppliers. In the remaining work, customers know the

expected quality level offered by all suppliers in equilibrium, and they frequent only the supplier that maximizes expected utility. They do not modify their choice based on the history of the level of service they actually receive.

Our underlying model of consumer behavior differs from the above work in that we assume that consumers are not well informed about quality levels, and the resulting behavior falls somewhere in between these two extremes. Consumers do not continually jockey among suppliers. In the short run they remain “loyal” and stay with one supplier. But they do respond to the history of the service they actually receive, and the resulting equilibria are ones in which consumers continue to switch among suppliers.

Our results concerning customer behavior are consistent with a stream of articles and books on customer loyalty generated by Sasser and colleagues (for example, see Heskett et al. 1994, Jones and Sasser 1995). This literature reports a strong relationship between the level of quality offered by a supplier and the resulting loyalty displayed by customers. It does not propose a specific model of consumer behavior that generates this response, however.

Our representation of consumer response to uncertain quality builds on several streams of work in the behavioral literature, and it attempts to reflect two sets of findings from empirically driven work on individual perception and decision making. First, in the context of controlled experiments—such as those of Horowitz (1973), Meyer and Shi (1995), and Banks et al. (1997)—people appear to behave in a roughly Bayesian fashion, with some biases, such as a tendency to behave more myopically than is optimal. Second, it appears that people systematically categorize as they make sense of their experiences. That is, people maintain mental examples of how entities in the world behave. They then interpret an experience with an entity by comparing their perceptions with the typical or exemplary characteristics of their mental picture of how that category behaves. This structure shows up in Kahneman and Tversky’s (Kahneman and Tversky 1973, Tversky, and Kahneman 1974) well-known “representativeness heuristic,” and it forms the basis of category and exemplar theory in cognitive psychology.

(For a review and interpretation, see Henderson and Peterson 1992.)

Our resulting model of decision making is similar to the cognitive theory of Gigerenzer and Murray (1987). It also has an intimate connection with the sequential probability ratio test (SPRT) developed by Wald (1947) in the 1940s, as well as with the “Cumulative Utility Consumer Theory” of Gilboa and Schmeidler (1997) and Gilboa and Pazgal (2001). Indeed, our representation of consumer choice may be thought of as a composition of these models, and our results are analogues of those developed in Gilboa and Pazgal (2001).

The paper closest in spirit to ours is Hall and Porteus (2000), which develops a similar model of quality competition and applies it to inventory and queueing systems for which there are no economies of scale. (For an earlier related paper see also Smallwood and Conlisk 1979.) They show that when one firm has an advantage of more loyal customers—with lower probabilities of switching upon a service failure—then it is optimal for the firm to offer a lower level of quality than its competitor.

Hall and Porteus (2000) also explicitly models customers switching among suppliers, but it differs from the current paper in two respects. First, it considers dynamic policies in which suppliers may change service quality in response to short-term changes in market share. We consider only static quality policies in which each supplier decides on a quality level and maintains it. Our stationary quality policies are simpler to analyze and implement, and they are intended to be consistent with industry practice, which typically sets stationary targets for service levels such as fill rates and queueing delays. Nevertheless, they clearly may not perform as well as more dynamic policies. Second, our underlying model of customer switching behavior differs from theirs. In Hall and Porteus (2000), a customer switches suppliers as an immediate response to a service failure at the current supplier, rather than as a result of the history of service she receives at the various suppliers. In addition, in Hall and Porteus (2000) switching is mediated by an exogenously defined loyalty factor. In our model of consumer behavior, loyalty is completely (endogenously) driven by supplier quality.

3. Supplier Model

There are m suppliers of a product or service. At discrete times $t = 1, 2, \dots$ each of n customers chooses one of the suppliers and acquires a unit of the good. Let n_t^i denote the number of customers patronizing supplier i at time t , so that $n = \sum_{i=1}^m n_t^i$.

The quality of the good is inherently uncertain, and the utility obtained when supplier i provides the good to customer j at time t is a random variable, $U_t^{i,j}$. We make two fundamental assumptions that limit the nature of the variation:

ASSUMPTION 1. *All customers derive the same utility from the same physical experience.*

ASSUMPTION 2. *The sequence of utilities that any customer obtains from a supplier is i.i.d.*

The first assumption allows us to translate any physical measure of the distribution of quality directly into a distribution of utility. The second implies that each supplier's choice of a quality distribution is a strategic decision. Supplier i must choose its particular quality distribution once, before consumers enter the market, and then live with the consequences.

Together, the assumptions allow us to describe a supplier's quality choice as a single distribution of utility, U^i . We may think of physical measures of system performance and of the resulting customer utility in similar terms, although for consumers that are risk averse with respect to the physical measure of service quality the two need not be identical.

For example, suppose that a physical measure of system quality is normally distributed but that customers are risk averse. Then two service providers that offer the same mean level of the physical quantity, but different levels of variability, would have different—and not necessarily normally distributed— U^i 's. Furthermore, the supplier with the lower variance would provide a higher expected utility.

Let $\Theta \subseteq \mathbf{R}$ be the set of distributions from which each of the suppliers chooses. The set may be finite or countable, or it may represent a continuum of choices. Then supplier i 's choice of quality distribution is $U^i \sim F(u|\theta^i)$, and we define $\mu_i \triangleq \mathbf{E}[U^i] = \mathbf{E}_{\theta^i}[U]$. If Θ represents an exponential family of distributions, then $\mu_i = \psi'(\theta^i)$.

REMARK. A distribution that comes from an *exponential family*, indexed by $\theta \in \Theta$, may be defined as

$$dF(x|\theta) \triangleq e^{\{\theta x - \psi(\theta)\}} dF(x), \quad (1)$$

where $\psi(\theta) \triangleq \ln \int e^{\theta x} dF(x)$ is the *cumulant generating function* of some nondegenerate (nonpoint-mass) distribution function, F . A random variable X , whose distribution comes from a one-dimensional exponential family indexed by parameter $\theta \in \Theta \subseteq \mathbf{R}$, has mean $\mathbf{E}_{\theta}[X] = \psi'(\theta)$ and variance $\text{var}_{\theta}(X) = \psi''(\theta)$. Thus, $\psi(\cdot)$ is strictly convex, and the mean $\psi'(\cdot)$ is a monotonically increasing function of θ . Most commonly used distributional forms—such as Bernoulli, exponential, Poisson, normal with mean θ and fixed variance, and normal with variance θ and fixed mean—are exponential families.

We assume that prices are fixed and that each customer visit yields a supplier \$ r of revenue. Therefore, to determine supplier i 's total revenues in period t we need only derive the number of its patrons, n_t^i .

We define supplier i 's cost of delivering the product or service to an arriving customer as a function $c_i(\mu_i)$ that is determined at least in part by the average level of utility offered, μ_i . Note that for risk-averse consumers an increase in μ_i may reflect a decrease in the variability, as well as an increase in the mean level of the physical measure of quality offered by the supplier. We assume that $c_i(\mu_i)$ is increasing and convex in μ_i and that there are no economies of scale in production. Thus, $c_i(\cdot)$ does not vary with n_t^i .

Supplier i seeks to maximize its long-run average profits:

$$\begin{aligned} \max_{U^i} \left\{ \Pi_i^m(U^i) \triangleq \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t (r n_s^i - c_i(\mu_i) n_s^i) \right. \\ \left. = (r - c_i(\mu_i)) \times \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t n_s^i \right\}. \quad (2) \end{aligned}$$

Here, Π_i^m denotes supplier i 's profit when m suppliers compete in an oligopoly, and $\{n_t^i : t = 1, 2, \dots\}$ is determined by customer switching behavior.

Given increasing and convex unit costs, the unit margin $(r - c_i(\mu_i))$ is a decreasing and concave function of average quality. To precisely define long-run average profits, however, we must specify how the

competitors' quality distributions affect supplier i 's long-run average market share, $\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t n_s^i$. The following two sections develop these results.

4. Customer Model

In this section we develop a model of consumer behavior that reflects the effect of the suppliers' quality distributions on a customer's choice of provider. We begin by describing the consumer's response to uncertain quality as the solution to a Bayesian multi-armed bandit problem. We then use this representation as the basis of a model of a more myopic, "simple" consumer whose heuristic solution to the bandit problem reflects well-known biases described in the behavioral literature.

4.1. Bandit Model of Customer Choice

The consumer is aware of the nature of the supply process described above, though she does not know what θ s the suppliers have chosen. Each time she returns to the marketplace, she uses the information she has acquired through past samples of the providers' performance to decide her choice of firm.

More formally, the consumer views each θ^i as a random variable, and for each provider she maintains a probability distribution of θ^i , $P_t^i(\theta)$, that represents her understanding at time t of the utility distribution under which she believes the supplier to be operating. Let $\{P_0^1, \dots, P_0^m\}$ be the "prior" information the consumer has as she enters the market for the first time.

After each contact with a supplier she uses the new sample of utility and Bayes's rule to update this belief. Thus, if the consumer uses provider i at time t and receives utility u , then her new belief distribution will be

$$dP_t^i(\theta|u) = \frac{dP_{t-1}^i(\theta) dF(u|\theta)}{\int_{\theta} dP_{t-1}^i(\theta) dF(u|\theta)}. \tag{3}$$

If $j \neq i$, then $dP_t^j(\theta) = dP_{t-1}^j(\theta)$.

Let a policy $\gamma = \{\gamma(1), \gamma(2) \dots\}$ be a sequence of choices of suppliers by a consumer, and let Γ be the class of policies that is nonanticipating with respect to future utility. Then the consumer seeks a policy, $\gamma \in \Gamma$, that will maximize the expected discounted value

of the future stream of utilities, $\sup_{\gamma \in \Gamma} E_{\gamma}[\sum_{t=1}^{\infty} \alpha^t U_t]$, where $\alpha \in (0, 1)$ is the one-period discount rate.

It is well known that for any fixed set of service distributions, $\{\theta^1, \dots, \theta^m\}$, and priors, $\{P_0^1, \dots, P_0^m\}$, the consumer's problem can be represented as a multi-armed bandit. Here, each supplier represents an arm and P_t^i the arm's state at time t . Arm i 's state evolves only at epochs, $\gamma(t) = i$, at which the arm is played, and from (3) we see that when the arm is played its evolution is Markov.

Gittins and Jones (1974) showed that (1) for each supplier i the consumer may construct an index, commonly called the *Gittins index*, which is calculated independently of the information concerning the other suppliers; and (2) at any time t it is optimal for the consumer to use the supplier with the highest Gittins index. Gittins (1979) further characterized the index of supplier i as the result of maximizing expected discounted utility per unit of expected discounted time,

$$G(P_t^i) = \sup_{\tau_i > t} \left\{ \frac{E[E[\sum_{s=t+1}^{\tau_i-1} \alpha^{s-t} U(P_{s-1}^i) | P_t^i]]}{E[E[\sum_{s=t+1}^{\tau_i-1} \alpha^{s-t} | P_t^i]]} \right\}, \tag{4}$$

where τ_i is a stopping time with respect to the history of the process through time $t - 1$, and the notation, $U(P_{s-1}^i)$, emphasizes the fact that the marginal (subjective) distribution of utility at $s - 1$ is a function of the distribution on the consumer's belief at the time.

To maximize her expected discounted utility, the consumer can in theory follow this simple algorithm. First, calculate the Gittins indices of all m suppliers. Second, choose the supplier, i , with the largest $G(P_t^i)$ and sample from i until the stopping time, τ_i , is reached. Then, using the new sample information, recalculate the Gittins index for i and go back to Step 2.

In practice, however, the calculation of the Gittins index is a formidable task. The state space of a general Bayesian bandit, as it evolves, covers the set of all possible sequences of posterior distributions generated by sample paths of the reward process. From a prescriptive standpoint, this poses a computational problem for rational decision makers who are faced with bandit problems.

4.2. A Model of a “Simple” Consumer

For our purposes, difficulty in calculating the Gittins index also poses a *descriptive* problem. That is, if the Gittins index is so difficult to calculate, can we believe that consumers behave “as if” they are calculating them?

In fact, the empirical work described in the literature review suggests that people display systematic biases away from purely Bayesian behavior. Our approach to modeling the problem is to simplify the consumer’s decision-making process in a way that results in an analytically tractable structure that is consistent with these biases. We refer to this model as one of a *simple consumer* because, as we shall see, it is analogous to the case of a one-sided test of a simple hypothesis in sequential analysis.

The model for simple consumers works as follows. As in the original model, each of the m suppliers chooses a quality distribution $\theta^i \in \Theta$. Rather than maintaining a complex set of beliefs concerning suppliers, however, the customer partitions the suppliers into two categories—good and bad—with respective utility distributions θ_G and θ_B . For each supplier, the consumer maintains, in turn, a belief distribution that is solely the probability that the supplier is good or bad. Thus, rather than judging *how* good or bad a supplier is, the consumer’s problem is more simply to decide whether a supplier is good or bad. At any time, t , the consumer myopically chooses the supplier that has the highest probability of being good.

Let p_t^i be the consumer’s subjective probability at time t that supplier i is good, and consider a sequence of visits to supplier i in which utilities $\{U_1, U_2, \dots\}$ are obtained. Then from a direct application of Bayes’s rule (3) we find the posterior probability that i is good after period t is

$$p_t^i = \left[1 + \left(\frac{1 - p_0^i}{p_0^i} \right) \prod_{s=1}^t \frac{dF_B(U_s)}{dF_G(U_s)} \right]^{-1}, \tag{5}$$

where $F_B = F(\cdot|\theta_B)$, $F_G = F(\cdot|\theta_G)$, $U_s \sim F_i = F(\cdot|\theta^i)$, and p_0^i denotes the consumer’s prior belief concerning the quality of i .

Observe that when the suppliers’ choices of quality distributions are themselves simple—so that $\Theta = \{\theta_B, \theta_G\}$ —then the binary nature of alternatives also

leads to a fundamental simplification of the calculation of the Gittins index, and the simple consumer’s myopic policy *is* optimal:

THEOREM 1 (BANKS AND SUNDARAM 1992). *Suppose $\Theta = \{\theta_B, \theta_G\}$. Then the Gittins index of supplier i , G_t^i , is monotonically increasing with the probability that i is good, p_t^i .*

Even when Θ is more complex than $\{\theta_B, \theta_G\}$, if the consumer only believes $\Theta = \{\theta_B, \theta_G\}$ or can only perceptually discriminate on the basis of $\Theta = \{\theta_B, \theta_G\}$, then it is optimal for her to act myopically. Thus, for a simple consumer, myopic choice behavior represents the rational response to an inability to distinguish among many different quality distributions, rather than an inherent inability to weigh the future consequences of current choices.

5. Long-Run Customer and Market Share

Consider a market in which there are m suppliers competing for the patronage of a consumer, and let $f_i \triangleq \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \mathbf{1}\{\gamma(s) = i\}$ denote the limiting relative frequency with which the consumer chooses supplier i . In this section we analyze the model of a simple consumer to characterize f_i as a function of the quality distributions set by the suppliers.

First, we let $j = \arg \max\{k : p_0^k \leq p_0^i\}$. Then a simple consumer will buy from i as long as $p_t^i \geq p_t^j$. By (5), this is equivalent to

$$\prod_{s=1}^t \frac{dF_G(U_s)}{dF_B(U_s)} \geq \left(\frac{p_0^j}{p_0^i} \right) \times \left(\frac{1 - p_0^j}{1 - p_0^i} \right), \tag{6}$$

a one-sided test of a simple hypothesis against a simple alternative (see Wald 1947).

By taking logs on both sides of (6), we equivalently have a random walk, $S_t = \sum_{s=1}^t X_s$, with i.i.d. increments

$$X_s \triangleq \ln(dF_G(U_s)/dF_B(U_s)) \tag{7}$$

and stopping boundary

$$b^i \triangleq \ln(p_0^j/p_0^i \times (1 - p_0^j)/(1 - p_0^i)). \tag{8}$$

Here X is a function of the random variable $U_s \sim U^i$, and its expectation, $E_{\theta^i}[X]$, is evaluated with respect to the utility distribution actually offered by supplier i .

Let $\tau \triangleq \inf\{t : p_t^i < p_0^j\} = \inf\{t : S_t < b^i\}$ be the time at which the customer first believes supplier i is less likely to be good than supplier j . This is the time at which the consumer “defects” to the competition, and one can characterize this time to defection as the stopping time of a one-sided SPRT (see Gans 2000).

Here, we concentrate on the customer’s long-run behavior, however. We wish to characterize how long-run shares $\{f_1, \dots, f_m\}$ are determined by the suppliers’ quality distributions $\{\theta^1, \dots, \theta^m\}$, and we use results from the cumulative utility consumer choice model developed in Gilboa and Pazgal (2001) to do so:

THEOREM 2 (GILBOA AND PAZGAL 2001). *If $E_{\theta^i}[X] < 0$ and $\text{var}_{\theta^i}(X) < \infty$ for $i = 1, \dots, m$, then with probability one f_i exists for $i = 1, \dots, m$, and*

$$f_i = \frac{\prod_{j \neq i} E_{\theta^j}[X]}{\sum_{k=1}^m \prod_{j \neq k} E_{\theta^j}[X]}. \tag{9}$$

Note that, while the mathematical structure of the choice behavior is the same, the intention of the cumulative utility models is somewhat different from ours. In Gilboa and Pazgal’s (2001) formulation, a consumer makes repeated choices among a number of competing brands, and switching behavior is induced due to “variety-seeking” behavior: The more negative $E_{\theta^i}[X]$, the more easily bored the consumer becomes on average with choice i . Randomness in the realizations of X correspond to unpredictable differences in the consumer’s response to i , rather than to any (measurable) uncertainty in the actual performance of i .

When are $E_{\theta^i}[X]$ and $\text{var}_{\theta^i}(X)$ finite? For F_B and F_G that are members of the same exponential family of distributions, the answer is straightforward to determine. From (1) and (7) we have

$$\begin{aligned} X &= \ln\left(\frac{e^{(\theta_G U - \psi(\theta_G))} dF}{e^{(\theta_B U - \psi(\theta_B))} dF}\right) \\ &= (\theta_G - \theta_B)U - (\psi(\theta_G) - \psi(\theta_B)). \end{aligned} \tag{10}$$

In turn, we have

$$E_{\theta^i}[X] = (\theta_G - \theta_B)\mu_i - (\psi(\theta_G) - \psi(\theta_B)) \tag{11}$$

and

$$\begin{aligned} E_{\theta^i}[X^2] &= (\theta_G - \theta_B)^2 E_{\theta^i}[U^2] \\ &\quad - 2(\theta_G - \theta_B)(\psi(\theta_G) - \psi(\theta_B))\mu_i \\ &\quad + (\psi(\theta_G) - \psi(\theta_B))^2. \end{aligned} \tag{12}$$

Thus, if $\psi(\theta_G)$ and $\psi(\theta_B)$ are well defined, then $E_{\theta^i}[X]$ and $E_{\theta^i}[X^2]$ are finite whenever μ_i and $E_{\theta^i}[U^2]$ are.

If we define

$$\mu^* \triangleq (\psi(\theta_G) - \psi(\theta_B))/(\theta_G - \theta_B), \tag{13}$$

then we also see $E_{\theta^i}[X] < 0$ if and only if $\mu_i < \mu^*$. Recalling that $X = \ln(dF_G(U)/dF_B(U))$, we see from (13) that μ^* is the level of quality at which the consumer cannot, in expectation, determine whether i ’s distribution is good or bad. One may think of μ^* as the customer’s “aspiration” level, the level of average quality at which she is satisfied, and of $(\mu^* - \mu_i)$ as describing how far short supplier i falls, on average, in each service encounter.

If we divide both numerator and denominator by $\prod_{j=1}^m E_{\theta^j}[X]$, then from (9) we recover

$$f_i = \frac{E_{\theta^i}[X]^{-1}}{\sum_{j=1}^m E_{\theta^j}[X]^{-1}}, \tag{14}$$

and from (13) we have:

COROLLARY 1. *Suppose F_B and F_G are members of the same exponential family of distributions. If $\mu_i < \mu^*$ and $\text{var}_{\theta^i}(U) < \infty$ for $i = 1, \dots, m$, then f_i exists for $i = 1, \dots, m$ with probability one and*

$$f_i = \frac{1/(\mu^* - \mu_i)}{\sum_{j=1}^m 1/(\mu^* - \mu_j)} = \frac{1}{1 + (\mu^* - \mu_i)\Delta}, \tag{15}$$

where $\Delta \triangleq \sum_{j \neq i} (\mu^* - \mu_j)^{-1}$.

Observe that when F_G and F_B are members of the same exponential family of distributions, f_i is determined solely by the expected utilities offered by the suppliers and is insensitive to the distributional forms of the various U^i ’s. When we wish to emphasize this functional relationship, we will write $f(\mu_i, \mu_{-i})$, where $\mu_{-i} \triangleq \{\mu_j : j \neq i\}$.

Furthermore, the condition $\mu_i < \mu^*$ for all i requires that the customer not be satisfied with any of the

suppliers. Indeed, by offering μ^* a supplier locks in customer loyalty, and if a supplier can profitably offer μ^* (that is, $r - c_i(\mu^*) > 0$), then it can expect to profitably serve the customer for all time. In this case, the expression for the suppliers' long-run shares of the customer would be degenerate.

In broader terms, if suppliers can afford to offer quality level μ^* , then they satisfy the customer's quality requirements, and quality is no longer the basis of competition. In this case all suppliers may offer μ^* , and the focus of competition would shift to other arenas, such as efforts to attract new customers or product and service differentiation.

Note also that, because f_i is a long-run average, the consumer's prior beliefs $\{b^1, \dots, b^m\}$ do not have a long-term effect. Equation (15) describes the supplier's long-run share of a customer solely in terms of her expected satisfaction with a single service, $(\mu^* - \mu_i)$.

Finally, Δ can be thought of as an aggregate measure of the competitive intensity supplier i faces. Δ increases with each μ_j and, given a fixed set of μ_j 's, it increases with m . As Δ increases, f_i decreases.

By differentiating f_i with respect to the suppliers' quality levels we see that customer response to changes in supplier quality is sharply increasing:

COROLLARY 2. Suppose F_B and F_C are members of the same exponential family of distributions. If $\mu_i < \mu^*$ and $\text{var}_{\theta_i}(U) < \infty$ for $i = 1, \dots, m$, then:


$$(i) \quad \frac{\partial f_i}{\partial \mu_i} = \frac{f_i(1-f_i)}{\mu^* - \mu_i} > 0 \quad \text{and} \quad \frac{\partial^2 f_i}{\partial \mu_i^2} = \frac{2f_i(1-f_i)^2}{(\mu^* - \mu_i)^2} > 0, \quad (16)$$

$$(ii) \quad \frac{\partial f_i}{\partial \mu_j} = -\frac{f_j^2}{\mu^* - \mu_i} < 0 \quad \text{and} \quad \frac{\partial^2 f_i}{\partial \mu_j^2} = -\frac{2f_j^2(1-f_j)}{(\mu^* - \mu_i)(\mu^* - \mu_j)} < 0, \quad (17)$$

$$(iii) \quad \frac{\partial^2 f_i}{\partial \mu_i \partial \mu_j} = \left(\frac{f_j}{\mu^* - \mu_i} \right) \times \left(\frac{2f_i^2}{\mu^* - \mu_j} - \frac{f_j}{\mu^* - \mu_i} \right). \quad (18)$$

Furthermore, (18) implies that $\frac{\partial^2 f_i}{\partial \mu_i \partial \mu_j} > 0 \iff f_i > 1/2$.

Figure 1 Customer Share as a Function of μ_i and Δ



PROOF. Omitted. \square

Parts (i) and (ii) show that relative purchase frequencies at i are strictly convex and increasing in μ_i and strictly concave and decreasing in μ_j for all $j \neq i$. Thus, long-run customer share responds increasingly strongly as changes in quality grow larger.

These relationships can be seen in Figure 1, which shows customer share as a function of both μ_i and $\Delta = \sum_{j \neq i} (\mu^* - \mu_j)^{-1}$, the aggregate measure of competitive intensity faced by supplier i . Here, larger Δ drives the response to μ_i to be first flatter and then more sharply convex. That is, the greater Δ , the further down and to the right the "elbow" of the curve lies. It is worth noting that the relationship pictured in Figure 1 roughly matches the empirically based curves of the more competitive industries shown in Jones and Sasser (1995).

Part (iii) of the corollary shows that the crossderivative of f_i with respect to μ_i and μ_j is increasing if and only if the relative frequency at i is already greater than one half. Given a competitor's marginal increase in quality, a dominant high-quality supplier would see increasing marginal (revenue) returns to raising its own quality, while a smaller lower-quality competitor would see decreasing benefits. Similarly, in a market with $m > 2$ competitors of equal quality levels, each supplier would find that increases in competitors' quality levels would decrease the customer-share "return" it would obtain by marginally increasing its quality.

Finally, given identical customers it is not difficult to show that long-run average market share

is a multiple of the long-run share of an arbitrary customer:

COROLLARY 3. *Given Assumptions 1 and 2, suppose $-\infty < \mu_i < \mu^*$ and $\text{var}(U^i) < \infty$ for $i = 1, \dots, m$. Then*

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t n_s^i &= \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \sum_{j=1}^n \mathbf{1}\{\gamma^j(s) = i\} \\ &= \sum_{j=1}^n \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \mathbf{1}\{\gamma^j(s) = i\} \\ &= nf(\mu_i, \mu_{-i}). \end{aligned} \tag{19}$$

Note that the interchange of limit and summation is justified by the fact that $0 \leq \sum_{j=1}^n \mathbf{1}\{\gamma^j(s) = i\} \leq n$ is bounded for every s .

Thus, while customer switching implies that the population of customers patronizing supplier i , $\{n_t^i : t = 1, \dots\}$, is dynamic and stochastic, we can use f_i to characterize supplier i 's long-run average market share. This characterization is the last essential element required for us to analyze quality competition.

REMARK. One may also be interested in the limiting distribution, $\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \mathbf{1}\{n_s^i = x\}$. For example, if we were to assume that in any period t the distribution of utility obtained by one customer is independent of that of the next, then it should not be difficult to show that the limiting random variable is binomially distributed with parameters n and $f_i(\mu_i, \mu_{-i})$.

6. Quality Competition

In this section we derive the fundamentals of a supplier's quality strategy. To do so, we model the problem as that of a noncooperative game in which each of m suppliers sets its quality level μ_i to maximize its long-run average profits. Each supplier's quality level determines its own unit cost, and together the suppliers' quality levels jointly determine each supplier's long-run average market share.

As before, we assume that customers are identical, that suppliers' quality distributions are stationary, and that F_G and F_B are from the same exponential family of distributions, so that (15) can be used to describe the long-run average share of each customer. We also assume that each supplier knows its competitors' cost functions, as well as the form of (15).

Using the result of Corollary 3, we can rewrite supplier i 's problem as

$$\max_{\mu_i} \{ \Pi_i^m(\mu_i) \triangleq (r - c_i(\mu_i)) \times n f_i(\mu_i, \mu_{-i}) \}. \tag{20}$$

Here, the supplier maximizes over its average quality level μ_i , rather than over its entire quality distribution $U^i \sim F(\cdot|\theta^i)$, and competitors' average quality levels, $\mu_{-i} = \{\mu_j | j \neq i\}$, are assumed to be fixed.

Given fixed μ_{-i} and $c_i(\mu_i)$ that is convex in μ_i , it is not difficult to show that $\Pi_i^m(\mu_i)$ is pseudoconcave in μ_i . The unit profit $(r - c_i(\mu_i))$ is concave and, if we let $g_i = 1/f_i = 1 + (\mu^* - \mu_i)\Delta$, then g_i is linear and decreasing in μ_i . In turn, for $\mu_i < \mu^*$ the profit, $(r - c_i(\mu_i))n/g_i$, is pseudoconcave in μ_i (see Mangasarian 1969). Therefore,

PROPOSITION 1. *Suppose that the $c_i(\cdot)$ are convex and increasing and that there exist $-\infty < \underline{\mu}_i < \bar{\mu}_i < \mu^*$ such that $\mu_i \in [\underline{\mu}_i, \bar{\mu}_i]$, for $i = 1, \dots, m$. Then there exists a pure-strategy Nash equilibrium to the quality competition game.*

PROOF. The strategy spaces of the suppliers are nonempty, compact, convex subsets of the real line, and each supplier's response function is quasi-concave in its quality level. Therefore, from Debreu (1952) the result follows. \square

The assumption that there exists a $\underline{\mu}_i > -\infty$ is mild. A reasonable assumption for $\bar{\mu}_i$ would be that $(r - c_i(\mu^*)) < 0$, so that $\bar{\mu}_i = \{\mu : r - c_i(\mu) = 0\}$. This is consistent with the assumptions that $c(\cdot)$ is increasing and that no supplier can profitably satisfy customers in expectation.

The first-order conditions are satisfied when the marginal increase in profits, due to an increase in market share, equals the marginal increase in cost over that share. Using (16) we have:

$$\Pi_i^{m'}(\mu_i) = (r - c_i(\mu_i))n \frac{f_i(1 - f_i)}{\mu^* - \mu_i} - c'_i(\mu_i)nf_i = 0, \tag{21}$$

for all suppliers, $i = 1, \dots, m$.

Note that the set of Nash equilibria includes

$$\{(\mu_1, \dots, \mu_m) : \Pi_i^{m'}(\mu_i) = 0 \quad \forall i = 1, \dots, m\}, \tag{22}$$

the set of quality vectors for which the first-order conditions (21) hold for all i . We call these equilibria

interior. It also may include vectors with some elements on the boundary of the action space: $\{\mu_i = \underline{\mu}_i \cap \Pi_i^m(\underline{\mu}_i) \leq 0\}$ or $\{\mu_i = \bar{\mu}_i \cap \Pi_i^m(\bar{\mu}_i) \geq 0\}$.

One general problem with using the search for Nash equilibria as a method of strategic analysis is that the set of possible equilibrium points may be large. In this case one must justify how competitors may converge on one of many possible equilibria.

Given our problem structure, we are able to show that the set of Nash equilibria is well structured, however. In particular, when the suppliers' technologies are identical, we can use the first-order conditions (21) to prove the following.

PROPOSITION 2. For $i = 1, \dots, m$, suppose that $c_i(\cdot) \equiv c(\cdot)$ is convex and increasing, that there exist $-\infty < \underline{\mu} < \bar{\mu} < \mu^*$ such that $\mu_i \in [\underline{\mu}, \bar{\mu}]$, and that $r - c(\bar{\mu}) = 0$. Then there exists a unique pure strategy Nash equilibrium, and it is symmetric.

PROOF. Please see the Appendix. \square

Thus, although simple consumers do not use all of the information available in the bandit setting, their ability to switch suppliers in response to poor service provides the discipline necessary to drive competitors to adopt an industry quality standard.

Just as symmetric costs drive symmetric equilibria, asymmetric costs drive equilibria to be asymmetric. As the following proposition shows, when one supplier has a (percentage) cost advantage over another, the two never choose the same quality level in equilibrium.

PROPOSITION 3. Let $c_i(\cdot) \equiv \alpha_i c(\cdot)$ for $i = 1, \dots, m$, where $\alpha_i \in (0, \infty)$ and $c(\cdot)$ is positive, convex, and increasing. For each i suppose there exist $\underline{\mu}_i < \bar{\mu}_i < \mu^*$ such that $\mu_i \in [\underline{\mu}_i, \bar{\mu}_i]$ and that $r - \alpha_i c(\bar{\mu}_i) = 0$.

(i) For any m , consider arbitrary suppliers j and k . If $\alpha_j < \alpha_k$, then there is no interior pure strategy equilibrium with $\mu_j = \mu_k$.

(ii) When $m = 2$, if $\alpha_1 < \alpha_2$, then any interior pure strategy equilibrium has $\mu_1 > \mu_2$. Furthermore, $\Pi_1^m(\mu_1) > \Pi_2^m(\mu_2)$.

PROOF. Please see the Appendix.

Part (i) of the proposition shows that asymmetric equilibria will generally result from asymmetric costs. Without more structure, however, it is difficult to say

more about the nature of the equilibria. Part (ii) adds the restriction that $m = 2$. In this case we can also say that for any positive, convex, and increasing cost function, the supplier with a cost advantage will offer higher quality and earn higher profits.

REMARK. Proposition 3 is an analogue to results in Lederer and Li (1997) and a complement to results in Hall and Porteus (2000). In our case a lower cost structure allows a supplier to increase quality, customer loyalty, and average profits. In Hall and Porteus (2000) the authors show that, given customers are a priori more loyal, a supplier may be able to decrease quality and improve profits.

Finally, we show that as the number of suppliers increases, industry quality levels rise as well, though never so much as to completely erode unit margins:

PROPOSITION 4. Suppose that for any m : (a) $c_i(\cdot)$ is convex and increasing for $i = 1, \dots, m$; and (b) there exist $-\infty < \underline{\mu} < \bar{\mu} < \mu^*$ such that $\mu_i \in [\underline{\mu}, \bar{\mu}]$ and that $r - c_i(\bar{\mu}) \leq 0$. Let μ_i^m be the equilibrium quality level offered by supplier i in such an m -player oligopoly equilibrium.

(i) Suppose that for any m , $c_i(\cdot) \equiv c(\cdot)$ for $i = 1, \dots, m$. Then the quality of service offered in equilibrium is increasing in m : $\mu_i^{m+1} \geq \mu_i^m$. Furthermore, if the equilibrium is interior, then quality is strictly increasing in m : $\mu_i^{m+1} > \mu_i^m$.

(ii) Consider a fixed supplier i with cost function $c_i(\cdot)$. There exists a quality level $\mu_i^\infty < \bar{\mu}$ such that for any fixed m , $\mu_i^m \leq \mu_i^\infty$. Furthermore, given any sequence of $m = 2, 3, \dots$ competitors, $\lim_{m \rightarrow \infty} \mu_i^m = \mu_i^\infty$.

PROOF. Please see the Appendix.

Part (i) of the proposition shows that for symmetric competitors the equilibrium level of quality is increasing with the number of competitors. Part (ii) shows that, furthermore, even in the limit increased competition never drives up quality so far that $r - c_i(\mu_i^\infty) = 0$. Rather, the customer's lack of information concerning i 's quality allows $r - c_i(\mu_i^\infty) > 0$, and each customer sojourn with supplier i earns positive profits. For a fixed number of customers n , as $m \rightarrow \infty$ long-run average supplier profits do vanish however, because i 's market share drops to zero. In the limit, the expected time between customer visits to i becomes infinite.

The proposition shows that consumers' lack of information is significant enough to permit suppliers to earn positive margins, even as $m \rightarrow \infty$. This limiting behavior is analogous to that found by Lippman and McCardle (1997) in a single-period newsvendor setting. In the newsvendor model, the deterministic allocation of customers to competing suppliers plays the same role as our model's information asymmetry; it "locks in" customers and allows the industry to earn positive profits in the limit.

The result also shows that for symmetric competitors there is a clear incentive for a set of suppliers to consolidate. As m decreases quality decreases, and unit, firm, and industry profits increase. Of course, this effect also reflects the fact that neither revenue nor the size of the customer base increases with the overall level of service quality offered in the industry.

7. Discussion

In this paper we have sought to better understand how suppliers should set quality levels in markets in which quality has unavoidable random variation. To that end we have developed a model of customer response to quality variation and have used it to demonstrate the effect of switching behavior on suppliers' quality strategies. Both elements of the analysis yield insight into the interplay between service quality and switching behavior.

Our model of consumer behavior provides simple closed-form expressions in which consumer response grows increasingly strongly with marginal improvements in quality. This behavior is consistent with empirical observations made by Sasser and colleagues (Heskett et al. 1994, Jones et al. 1995), and it is identical to that of Cumulative Utility Consumer Theory (Gilboa and Pazgal 2001).

It is also interesting to note that if we let $z_i = 1/E_{\theta_i}[X]$ in (14), then $f_i = z_i/(\sum_{j=1}^m z_j)$ bears a strong resemblance to the choice probability derived from classic multinomial logit (MNL) models of discrete choice, as well as to expressions for market share used in "attraction" models in the marketing literature (see Andersen et al. 1992 and Karnani 1985). There are two important differences, however.

First, while the expression, $f_i = z_i/(\sum_{j=1}^m z_j)$, is the same, the form of z_i differs among the models, and

this may lead to differences in market behavior. A systematic analysis of the differences may sharpen our insight into how the quality competition modeled in this paper varies from other forms of competition.

Second, the MNL model's expressions describe the probability of choosing supplier i in any single period and are not history dependent. In our model f_i is a long-run average, however. In any period the probability of a simple consumer choosing supplier i is highly correlated with its previous choices.

Conversely, the long-run behavior of simple consumers differs significantly from that of the bandit model described in §4.1. If $\mu_i < \mu^*$ for all of the suppliers, then the simple consumer will switch among all suppliers indefinitely and will never converge upon a preferred supplier (or a subset of suppliers). This is due to the fact that neither the form of her discrimination function, $dF_G(u)/dF_B(u)$, nor the dispersion of her posterior distributions, p_i^t , is changing with experience. In contrast, with experience the posterior distributions of a Bayesian with more complex priors would become less dispersed, and her choices would converge to a single supplier.

Given a market made up of identical consumers, this difference in individual behavior translates directly into significant differences in long-run market behavior. If customers are "complex" Bayesians, then as $t \rightarrow \infty$ each customer will settle on a supplier, market shares will be realized, and there will be no customer switching. In an analogous market with n simple consumers and m suppliers with $\mu_i < \mu^*$, however, customers will continue to circulate among the suppliers even as $t \rightarrow \infty$.

It is likely that neither of these limiting views (as $t \rightarrow \infty$) wholly captures the nature of customer switching. In many real markets, in each period new customers enter the market, existing customers exit, and the number of customers patronizing each supplier is likely to change. Thus, a fuller representation would model an "open" system, rather than a closed system in which n customers move about. The aggregation of the transient behavior of each customer—from the time she enters the market until the time she exits—into that of a whole population would provide a clearer characterization of the market response to changes in supplier quality.

Thus, although it may be stylized, in one important respect the aggregate behavior of a market full of simple consumers is appealing: It naturally leads to a representation of market share that is dynamic and stochastic, rather than fixed. For example, if the behavior of each of n customers is independent of that of the others, then we might model the steady-state distribution of the number of customers patronizing supplier i as binomially distributed with mean nf_i . Furthermore, the limited empirical evidence presented in Jones and Sasser (1995) is consistent with the long-term relationship described in (15), so this simplified representation may not be unduly biased.

Still, the underlying choice model may be further developed and extended. For example, switching costs are likely to systematically change customer loyalty behavior. Similarly, advertising and word of mouth also influence customer satisfaction and are likely to influence switching behavior. Furthermore, the underlying assumption that suppliers' quality levels are stationary may be revisited. In particular, in industries that enjoy rapid technological advances, supplier quality may be systematically improving, and expectations of improvement may induce customers to test alternative suppliers more frequently.

The paper's competitive analysis demonstrates that the simple consumer's switching behavior provides enough discipline to force suppliers into maintaining an industry norm, and that this industry standard increases with the number of competitors. Alternatively, when one competitor has a cost advantage over another it can invest part of its savings to improve quality, increase market share, and earn greater total profits.

Given an estimate of μ^* and a specific functional form for costs, the analysis allows a supplier to calculate what the equilibrium level of quality should be. Furthermore, if a supplier finds that the market is not in equilibrium and believes that competitors will not respond (quickly) to its change in quality, then it can use the reaction function (20) to help determine an "off-equilibrium" strategy as well.

The information required by i is as follows: an estimate of customer preferences, μ^* ; its own cost function, $c_i(\cdot)$; its current quality level μ_i and market share f_i ; and details about the form of the optimization

problem (20). In particular—given a knowledge μ^* , μ_i , and f_i —supplier i can use (9) to back out Δ , the required measure of competitive intensity, without having to know the actual quality levels of individual competitors.

Just as the consumer model may be further developed, elements of the competitive analysis may be extended. For example, the current analysis is limited to systems for which economies of scale have been exhausted. In many operational environments, however, economies of scale are important. Because unit costs in these cases do not vary linearly with demand, the steady-state distribution of demand—rather than just the long-term average—is required for us to characterize long-run average costs and analyze competitive equilibria. Other important factors that may be considered are customer heterogeneity and the ability of suppliers to trade off quality and price.

Thus, the analysis of this paper leads to two directions for further research. In the short run, these results may be extended to systems with economies of scale. In the longer term, further development and testing of these types of choice models will lead to a better understanding of both the determinants of switching behavior and how to manage them.

Acknowledgments

The author thanks Gérard Cachon and Rachel Croson, as well as the referees and editors, for helpful comments. Research was supported by the Wharton Financial Institutions Center and by NSF Grant SBR-9733739.

Appendix

PROOF OF PROPOSITION 2. First we show that, given symmetric suppliers, there exists no asymmetric pure strategy equilibrium. By contradiction, consider the quality levels of two suppliers, j and k , in an asymmetric equilibrium and, without loss of generality, suppose $\underline{\mu} \leq \mu_k < \mu_j$.

Letting $c_i(\cdot) \equiv c(\cdot)$ for arbitrary supplier i , we note that $r - c(\bar{\mu}) = 0$ implies that $\mu_j < \bar{\mu}$, since μ_j is an equilibrium solution, and inspection of (21) shows that $\Pi_j^{m'}(\bar{\mu}) < 0$. Similarly, the fact that $\underline{\mu} \leq \mu_k < \mu_j$ implies that $\Pi_j^{m'}(\mu_j) = 0$, since $\mu_j \in (\underline{\mu}, \bar{\mu})$ is an interior equilibrium point for j . Thus, we have $\Pi_k^{m'}(\mu_k) \leq \Pi_j^{m'}(\mu_j) = 0$.

We can write the first-order conditions (21) as:

$$\frac{nf_i(1-f_i)}{\mu^* - \mu_i} \left[(r - c(\mu_i)) - \frac{\mu^* - \mu_i}{1-f_i} c'(\mu_i) \right] = 0; \quad (A1)$$

$$\frac{nf_i(1-f_i)}{\mu^* - \mu_i} \left[(r - c(\mu_i)) - \frac{(\mu^* - \mu_i) \sum_{l=1}^m (\mu^* - \mu_l)^{-1}}{\sum_{l \neq i} (\mu^* - \mu_l)^{-1}} c'(\mu_i) \right] = 0. \quad (A2)$$

In turn, letting $\delta \triangleq \sum_{\{i:i \neq j,k\}} (\mu^* - \mu_i)^{-1}$, we can rewrite the term in square brackets within (A2) to define

$$g(\mu_1, \mu_2) \triangleq (r - c(\mu_1)) - \left(\frac{\mu^* - \mu_2}{\delta(\mu^* - \mu_2) + 1} + (\mu^* - \mu_1) \right) c'(\mu_1) \quad (A3)$$

for arbitrary suppliers 1 and 2. Note that $\Pi_k^{m'}(\mu_k) \leq \Pi_j^{m'}(\mu_j) = 0$ implies that $g(\mu_k, \mu_j) \leq g(\mu_j, \mu_k) = 0$ as well, since the terms outside of the square brackets in (A2) are positive.

At the same time, note that for any $\mu_1 < \mu^*$, $\mu_2 < \mu^*$, and convex, increasing $c(\cdot)$, we have

$$\begin{aligned} \frac{\partial g}{\partial \mu_1} &= -(\mu^* - \mu_1)c''(\mu_1) \leq 0 \quad \text{and} \\ \frac{\partial g}{\partial \mu_2} &= \frac{c'(\mu_1)}{(\delta(\mu^* - \mu_2) + 1)^2} > 0. \end{aligned} \quad (A4)$$

This implies that $g(\mu_j, \mu_k) \leq g(\mu_k, \mu_k) < g(\mu_k, \mu_j)$, which contradicts $g(\mu_k, \mu_j) \leq g(\mu_j, \mu_k) = 0$. Thus, given symmetric competitors there cannot be an asymmetric equilibrium.

Second, we show that, given all equilibria are symmetric, there exists at most one equilibrium that satisfies the first-order conditions. Since the equilibrium is symmetric, we let $f_i = 1/m$ and $\mu_i = \mu$ for all i , and we rewrite first-order conditions (21) as follows:

$$\begin{aligned} \Pi^{m'}(\mu) &= \left(\frac{m-1}{m^2} \right) \left(\frac{n}{\mu^* - \mu} \right) \\ &\times \left[(r - c(\mu)) - \frac{m}{m-1} (\mu^* - \mu)c'(\mu) \right] = 0. \end{aligned} \quad (A5)$$

Thus, the first-order conditions are satisfied if and only if the terms in the square brackets sum to zero, or equivalently if $g(\mu) \triangleq c(\mu) + \left(\frac{m}{m-1} \right) (\mu^* - \mu)c'(\mu) = r$. Differentiating $g(\mu)$ and collecting terms, we see that

$$g'(\mu) = \frac{1}{m} c'(\mu) + \left(\frac{m}{m-1} \right) (\mu^* - \mu)c''(\mu) \quad (A6)$$

is positive for all $\mu < \mu^*$, since $c(\cdot)$ is increasing and convex. Thus, there can be at most one μ for which $\Pi^{m'}(\mu) = 0$.

Finally, recall that $r - c(\underline{\mu}) = 0$ implies that $\Pi^{m'}(\underline{\mu}) < 0$. If there is a $\hat{\mu} \in (\underline{\mu}, \bar{\mu})$ such that $\Pi^{m'}(\hat{\mu}) = 0$, then $\Pi^{m'}(\mu) < 0$ for all $\mu \in (\underline{\mu}, \hat{\mu})$, $\Pi^{m'}(\mu) > 0$ for all $\mu \in (\hat{\mu}, \bar{\mu})$, and there is exactly one equilibrium. Otherwise, $\Pi^{m'}(\mu) < 0$ for all $\mu \in (\underline{\mu}, \bar{\mu})$, and $\underline{\mu}$ is the unique equilibrium. \square

PROOF OF PROPOSITION 3. When $c_i(\cdot) = \alpha_i c(\cdot)$, we can write the first-order conditions (21) as

$$(r - \alpha_i c(\mu_i)) n \frac{f_i(1-f_i)}{\mu^* - \mu_i} - \alpha_i c'(\mu_i) n f_i = 0. \quad (A7)$$

Part (i). Proof by contradiction. Assume that there exists an equilibrium with suppliers j and k in which $\alpha_j < \alpha_k$ but $\mu_j = \mu_k$ and $\Pi_j^{m'}(\mu_j) = \Pi_k^{m'}(\mu_k) = 0$. Then the only difference between the first-order conditions of j and k are the α terms. Furthermore, since $c(\cdot)$ is positive and increasing, all terms in (A7) are positive, so it must be the case that $\Pi_j^{m'}(\mu_j) > \Pi_k^{m'}(\mu_k)$, a contradiction.

Part (ii). For $m = 2$, we assume $\alpha_1 < \alpha_2$ and rewrite the first-order conditions (A7) for arbitrary supplier i as follows:

$$\begin{aligned} n f_i \alpha_i \left[(r/\alpha_i - c(\mu_i)) \frac{1-f_i}{\mu^* - \mu_i} - c'(\mu_i) \right] \\ = n f_i \alpha_i \left[\frac{r/\alpha_i - c(\mu_i)}{(\mu^* - \mu_i) + (\mu^* - \mu_j)} - c'(\mu_i) \right] \end{aligned} \quad (A8)$$

since $1 - f_i = f_j = \frac{(\mu^* - \mu_j)}{(\mu^* - \mu_i) + (\mu^* - \mu_j)}$. Furthermore, $\Pi_i^{m'}(\mu_i) = 0$ if and only if the term in square brackets equals zero, or equivalently when

$$r = \alpha_i c(\mu_i) + \alpha_i c'(\mu_i) (\mu^* - \mu_i + \mu^* - \mu_j). \quad (A9)$$

Since $c(\cdot)$ is positive and increasing, $\alpha_1 < \alpha_2$ implies that $\mu_1 > \mu_2$ is necessary to satisfy $\Pi_1^{m'}(\mu_1) = \Pi_2^{m'}(\mu_2) = 0$. Thus, lower costs prompt supplier 1 to offer higher average quality.

Finally we show that equilibrium profits must be higher for supplier 1. By contradiction, suppose that $\alpha_1 < \alpha_2$ and $\mu_1 > \mu_2$ but that $\Pi_1^m(\mu_1) \leq \Pi_2^m(\mu_2)$. Noting that $\mu_1 > \mu_2$ implies $f_2 < 1/2 < f_1$, we have $\Pi_2^m(\mu_2) = (r - \alpha_2 c(\mu_2)) n f_2 < (r - \alpha_2 c(\mu_2)) n \frac{1}{2} < (r - \alpha_1 c(\mu_2)) n \frac{1}{2}$. Thus, supplier 1 could have earned higher profits than supplier 2 simply by lowering its quality level from μ_1 to μ_2 , and μ_1 must not have been profit maximizing in the first place. \square

PROOF OF PROPOSITION 4. We rewrite (21) as follows:

$$\Pi_i^{m'}(\mu_i) = n f_i (1 - f_i) \left[\frac{r - c_i(\mu_i)}{\mu^* - \mu_i} - c'_i(\mu_i) - \frac{f_i}{1 - f_i} c'_i(\mu_i) \right]. \quad (A10)$$

Part (i). From Proposition 2 we know that symmetric suppliers will find a unique symmetric equilibrium, and we can rewrite the term within the square brackets of (A10) as follows:

$$\frac{r - c(\mu)}{\mu^* - \mu} - c'(\mu) - \frac{1}{m-1} c'(\mu). \quad (A11)$$

Observe that $\Pi_i^{m'}(\mu^m) = 0$ if and only if (A11) evaluates to zero and suppose that, for some m , this is the case. Then given $m + 1$ suppliers and the same μ^m , the far right term decreases from $c'(\mu^m)/(m-1)$ to $c'(\mu^m)/m$, and the entire expression becomes positive. Thus, $\Pi_i^{m'}(\mu^m) = 0 \implies \Pi^{m+1'}(\mu^m) > 0 \implies \mu^{m+1} > \mu^m$.

Alternatively, if $\Pi^m(\mu^m) < 0$, then $\mu^m = \underline{\mu}$. In this case, it must be that $\mu^m \leq \mu^{m+1}$.

Part (ii). Recalling from (15) that $f_i = 1/(1 + (\mu^* - \mu_i)\Delta)$, we write the term in square brackets of (A10) as

$$g(\mu_i, \Delta) \triangleq \frac{r - c_i(\mu_i)}{\mu^* - \mu_i} - c'_i(\mu_i) - \frac{1}{(\mu^* - \mu_i)\Delta} c'_i(\mu_i). \quad (A12)$$

Note that the first-order conditions are satisfied if and only if $g(\mu_i, \Delta) = 0$.

Now define

$$h_i(\mu_i) \triangleq \frac{r - c_i(\mu_i)}{\mu^* - \mu_i} - c'_i(\mu_i) = g(\mu_i, \Delta) - \frac{1}{(\mu^* - \mu_i)\Delta} c'_i(\mu_i), \quad (A13)$$

and let $\mu_i^\infty = \sup\{\mu \in [\underline{\mu}, \bar{\mu}] | h_i(\mu) \geq 0\}$. If $\sup\{\mu \in [\underline{\mu}, \bar{\mu}] | h_i(\mu) \geq 0\}$ is empty, let $\mu_i^\infty = \underline{\mu}$. Because $r - c_i(\bar{\mu}) = 0$ and $c'(\mu) > 0$ for all $\mu \in [\underline{\mu}, \bar{\mu}]$, it must be the case that $\mu_i^\infty < \bar{\mu}$.

To see that $\mu_i^m \leq \mu_i^\infty$ we note that for any fixed $\Delta > 0$, $g(\mu, \Delta) < h(\mu)$ for all $\mu \in [\underline{\mu}, \bar{\mu}]$, since $c(\cdot)$ is increasing. Thus, when $\underline{\mu} < \mu_i^\infty$, we have $g(\mu_i^\infty) < h(\mu_i^\infty) = 0$. The fact that Π^m is pseudoconcave then implies that $\mu_i^m \in [\underline{\mu}, \mu_i^\infty)$. Similarly, when $\mu_i^\infty = \bar{\mu}$, then $g(\bar{\mu}) < h(\bar{\mu}) \leq 0$, and $\bar{\mu}$ is optimal in the oligopoly equilibrium.

Next, we show that for a fixed competitor, i , the sequence $\{\mu_i^m, \mu_i^{m+1}, \dots\}$ converges to μ_i^∞ . The preceding argument implies that if $\mu_i^\infty = \underline{\mu}$, then $\mu_i^m = \underline{\mu}$ for all m . If, however, $\mu_i^\infty \in (\underline{\mu}, \bar{\mu})$, we can use the Implicit Function Theorem to prove the result. Note that the partial derivative

$$\frac{\partial g}{\partial \mu_i} = \frac{r - c_i(\mu_i) - c_i'(\mu_i)(\mu^* - \mu_i)}{(\mu^* - \mu_i)^2} - c_i''(\mu_i) - \frac{c_i'(\mu_i)}{(\mu^* - \mu_i)^2 \Delta} - \frac{c_i''(\mu_i)}{(\mu^* - \mu_i) \Delta} \quad (A14)$$

$$= \frac{1}{\mu^* - \mu_i} g(\mu_i, \Delta) - c_i''(\mu_i) \left(1 - \frac{1}{(\mu^* - \mu_i) \Delta} \right) \quad (A15)$$

and that $\Pi_i^{m'}(\mu_i, \mu_{-i}) \leq 0$ implies $g(\mu_i, \Delta) \leq 0$. This, in turn, implies $\frac{\partial g(\mu_i, \Delta)}{\partial \mu_i} < 0$. Therefore, the inverse of the partial derivative exists at μ_i^m , and by the Implicit Function Theorem there is mapping from $\hat{\Delta} \triangleq 1/\Delta$ to μ_i^m that is continuously differentiable at $\hat{\Delta}$. Call this function $\mu(\hat{\Delta})$.

Now in three steps we show that the sequence $\{\mu_i^m, \mu_i^{m+1}, \dots\}$ converges to μ_i^∞ . First, since $\mu(\hat{\Delta})$ is continuous at $(\mu, \hat{\Delta})$, then for any $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $|\hat{\Delta}' - \hat{\Delta}| < \delta$, $|\mu(\hat{\Delta}') - \mu(\hat{\Delta})| < \varepsilon$. Second, call the equilibrium value of $\hat{\Delta}$ for some arbitrary set of m suppliers $\hat{\Delta}^m$ and note that

$$\frac{(\mu^* - \mu)}{(m-1)} \geq \hat{\Delta}^m = \frac{1}{\sum_{j \neq i} (\mu^* - \mu_j^m)^{-1}} \geq \frac{(\mu^* - \bar{\mu})}{(m-1)}. \quad (A16)$$

Therefore, for any $\delta > 0$ there exists an $m^* > 0$ such that for all $m, n \geq m^*$, $|\hat{\Delta}^m - \hat{\Delta}^n| < \delta$. Together, these two facts imply that for any $\varepsilon > 0$ there exists an m^* such that for all $m, n \geq m^*$, $|\mu(\hat{\Delta}^m) - \mu(\hat{\Delta}^n)| < \varepsilon$. Thus, the sequence $\{\mu_i^m = \mu(\hat{\Delta}^m), \mu_i^{m+1} = \mu(\hat{\Delta}^{m+1}), \dots\}$ converges. Third, the fact that the sequence converges to μ_i^∞ can be seen by comparing $g_i(\cdot, \cdot)$ to $h_i(\cdot)$ and noting that: (1) $\frac{c(\mu)}{\mu^* - \mu} \in (0, r/(\mu^* - \bar{\mu})]$ for all $\mu \in [\underline{\mu}, \bar{\mu}]$; and (2) $\lim_{m \rightarrow \infty} \hat{\Delta}^m = 0$. \square

References

Anderson, S. P., A. de Palma, J.-F. Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.
 Banks, J. S., R. K. Sundaram. 1992. A class of bandit problems yielding myopic optimal strategies. *J. Appl. Probab.* **29** 625–632.
 ———, M. Olson, D. Porter. 1997. An experimental analysis of the bandit problem. *Econom. Theory* **10** 55–77.
 Cachon, G. P., P. T. Harker. 2001. Competition and outsourcing with scaled economics. Working paper, OPIM Department, The Wharton School, University of Pennsylvania, Philadelphia, PA.
 Chayet, S., W. J. Hopp. 1999. Lead time competition under uncertainty. Working paper, Department of IE&MS, Northwestern University, Evanston, IL.

Cohen, M. A., S. Whang. 1997. Competing in product and service: A product life-cycle model. *Management Sci.* **43** 535–545.
 Debreu, D. 1952. A social equilibrium existence theorem. *Proc. National Acad. Sci.* **38** 886–893.
 Dewan, S., H. Mendelson. 1990. User delay costs and internal pricing for a service facility. *Management Sci.* **36** 1502–1517.
 Gans, N. 2000. Customer loyalty and supplier strategies for quality competition I: Customer loyalty. Working paper, OPIM Department, The Wharton School, University of Pennsylvania, Philadelphia, PA.
 Gigerenzer, G., D. J. Murray. 1987. *Cognition as Intuitive Statistics*. Lawrence Erlbaum Associates, Hillsdale, NY.
 Gilboa, I., A. Pazgal. 2001. Cumulative discrete choice. *Marketing Lett.* **12** 119–130.
 ———, D. Schmeidler. 1997. Cumulative utility consumer theory. *Internat. Econom. Rev.* **38** 737–761.
 Gittins, J. C. 1979. Bandit processes and dynamic allocation indices. *J. Roy. Statist. Soc. B* **41** 148–177.
 ———, D. M. Jones. 1974. A dynamic allocation index for the sequential design of experiments. J. Gani, K. Sakadi, I. Vinczo, eds. *Progress in Statistics*. North-Holland, Amsterdam, The Netherlands, 241–266.
 Hall, J. M., E. L. Porteus. 2000. Customer service competition in capacitated systems. *Manufacturing Service Oper. Management* **2** 144–165.
 Henderson, P. W., R. A. Peterson. 1992. Mental accounting and categorization. *Organ. Behavior Human Decision Processes* **51** 92–117.
 Heskett, J. L., T. O. Jones, G. W. Loveman, W. E. Sasser Jr., L. A. Schlesinger. 1994. Putting the service profit chain to work. *Harvard Bus. Rev.* **72** (March–April) 164–174.
 Ho, T.-H., Y.-S. Zheng. 1996. Setting customer expectation in service delivery. Working paper, Marketing Department, The Wharton School, University of Pennsylvania, Philadelphia, PA.
 Horowitz, A. D. 1973. Experimental study of the two-armed bandit problem. Unpublished Ph.D. dissertation, University of North Carolina at Chapel Hill, Chapel Hill, NC.
 Jones, T. O., W. E. Sasser Jr. 1995. Why satisfied customers defect. *Harvard Bus. Rev.* **73** (November–December) 88–99.
 Kahneman, D., A. Tversky. 1973. On the psychology of prediction. *Psych. Rev.* **80** 237–251.
 Karmarkar, U. S., R. C. Pitbladdo. 1997. Quality, class and competition. *Management Sci.* **43** 27–39.
 Karnani, A. 1985. Strategic implications of market share attraction models. *Management Sci.* **31** 536–547.
 Lederer, P. J., L. Li. 1997. Pricing, production, scheduling, and delivery-time competition. *Oper. Res.* **45** 407–420.
 Li, L. 1992. The role of inventory in delivery-time competition. *Management Sci.* **38** 182–197.
 ———, Y. S. Lee. 1994. Pricing and delivery-time performance in a competitive environment. *Management Sci.* **40** 633–646.
 Lippman, S. A., K. F. McCardle. 1997. The competitive newsboy. *Oper. Res.* **45** 54–65.
 Mandelbaum, A., N. Shimkin. 2000. A model of rational abandonments from invisible queues. *QUESTA* **36** 141–173.

- Mangasarian, O. 1969. *Nonlinear Programming*, McGraw-Hill, New York.
- Mendelson, H. 1985. Pricing computer services: Queueing effects. *Comm. ACM* **28** 312–321.
- , S. Whang. 1990. Optimal incentive-compatible priority pricing for the $M/M/1$ queue. *Oper. Res.* **38** 870–883.
- Meyer, R. J., Y. Shi. 1995. Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem. *Management Sci.* **41** 817–834.
- Smallwood, D. E., J. Conlisk. 1979. Product quality in markets where consumers are imperfectly informed. *Quart J. Econom.* **93** 1–23.
- Stidham, S. Jr. 1992. Pricing and capacity for a service facility: Stability and multiple local optima. *Management Sci.* **38** 1121–1139.
- Tsay, A., N. Agrawal. 2000. Supply chain dynamics under price and service competition. *Manufacturing Service Oper. Management* **2** 372–391.
- Tversky, A., D. Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *Science* **185** 1124–1131.
- van Mieghem, J. A. 2000. Price and service discrimination in queuing systems: Compatibility of $Gc\mu$ scheduling. *Management Sci.* **46** 1249–1267.
- Wald, A. 1947. *Sequential Analysis*. John Wiley & Sons, New York.

Accepted by Wallace J. Hopp; received August 7, 2000. This paper was with the author 3 months for 2 revisions.