Wharton

Financial
Institutions
Center

Customer Learning and Loyalty When Quality is Uncertain

by
Noah Gans

99-11

The Wharton School
University of Pennsylvania

THE WHARTON FINANCIAL INSTITUTIONS CENTER

The Wharton Financial Institutions Center provides a multi-disciplinary research approach to the problems and opportunities facing the financial services industry in its search for competitive excellence. The Center's research focuses on the issues related to managing risk at the firm level as well as ways to improve productivity and performance.

The Center fosters the development of a community of faculty, visiting scholars and Ph.D. candidates whose research interests complement and support the mission of the Center. The Center works closely with industry executives and practitioners to ensure that its research is informed by the operating realities and competitive demands facing industry participants as they pursue competitive excellence.

Copies of the working papers summarized here are available from the Center. If you would like to learn more about the Center or become a member of our research community, please let us know of your interest.

Anthony M. Santomero
Director

The Working Paper Series is made possible by a generous grant from the Alfred P. Sloan Foundation

Customer Learning and Loyalty When Quality is Uncertain

Noah Gans *†

January 1999

Abstract

A consumer has repeated contacts with a set of product or service providers. Each visit to a supplier yields the consumer some randomly distributed utility. The suppliers' utility distributions are unknown to the consumer, and to decide which supplier to visit, she uses a myopic variant of the decision rule used by a classical, utility-maximizing Bayesian. This rule is designed to be roughly consistent with empirical findings regarding individual choice under uncertainty.

For this model, we develop closed-form expressions that characterize both short-term and long-term measures of customer loyalty to a supplier. These results offer a rich picture of how consumer discrimination and prior beliefs interact with the level of quality actually offered by suppliers to determine customer loyalty.

*OPIM Department, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6366.

†Research supported by the Wharton Financial Institutions Center and by NSF Grant SBR-9733739.

1 Introduction

What is the cost of a backorder or stockout? Of making a customer wait longer in queue? Of a plane missing an on-time arrival? Of lowering quality standards just a bit? These costs are embedded in the heart of a wide range of models in operations management. Typically they are described as factors exogenous to the operational models being developed; *given* the stockout cost or the waiting cost, we seek to maximize profit. Often, they are also assumed to be linear: a stockout costs p per unit; waiting costs are c per job per unit of time.

Sometimes these assumptions are reasonable. If stockouts are expedited – the missing items being immediately procured from an alternative source at a premium – then the linear penalty, p per unit, may apply. If the waiting jobs in a queueing system are WIP – with holding costs dominated by the cost of capital invested - then the linear waiting costs, c per unit per unit of time, may be valid.

In systems in which consumers are directly affected by poor service, however, it is not clear that the costs are linear – or that even the basic framework of a stationary stockout cost per unit or holding cost per unit of time – applies. While standard texts that cover inventory management, such as Silver, Pyke, and Peterson [48] and Nahmias [39], mention the costs of lost customer "good will" they note that it is often difficult for managers to assign these costs to stockouts. Traditional operational models, which take these good will costs as exogenously given, are of no help.

In fact, current practice is to directly set service levels. In inventory systems line-item fill rate, the probability that an arriving order can be satisfied from on-hand inventory, is an often-used measure of service quality. In queueing environments such as telephone call centers, the expected delay of an incoming call and the probability that an incoming call will be delayed more than x seconds are analogous service-level targets that managers often set.

It is not surprising, though, that managers are often unsure that they have set the "right" quality level. While they can estimate the marginal change in operating cost that a change in performance will require, they cannot readily quantify the marginal change revenue that the change in quality will obtain. That is, to develop effective operating strategies, a company must first characterize consumer response to aspects of quality, such as product availability or on-time performance, which it only partially controls and are inherently uncertain.

How does one characterize customer response? This question falls within the domain of marketing and appears to be of current interest to marketing researchers. For example, the Marketing Science Institute's 1998-2000 research priorities [35] include the following topics (among others) as "being sufficiently significant and timely that they deserve intensive research attention: ... 1) [understanding the] value of [a] customer – value of loyalty, lifetime value of customer ... ; 2) understanding and driving customer satisfaction; 3) [understanding the] relationship between actual company performance and measures of customer satisfaction."

Thus, the question of how producers should account for consumer response to uncertain quality appears to be open. Nevertheless, over the past 25 years related streams of research in psychology, experimental economics, and marketing have begun to elucidate a picture of human decision-making upon which we can build a reasonable model of this response.

First, 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, and they 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 [28, 53] 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 [24].)

In addition, there are empirical studies of customer satisfaction with product and service quality that offer a picture of a consumer's utility or affect that reflects this categorization. In these models, a consumers' utility is based on a comparison she makes between the quality she receives and some prior expectation. For example, Boulding *et al.* [9] model the consumer as basing her satisfaction with service received on a comparison to two exemplars or standards: the level of service she believes the service provider *will* offer, and the level of service she believes the service provider *should* offer.

Finally, early work studying peoples' heuristic responses to uncertainty, such as that of Kahneman and Tversky [28, 53], has led to experimental work – by Horowitz [27], Grether [11, 21, 22], Camerer [10], Meyer and Shi [38], and Banks *et al.* [7] – that show that, in the context of controlled experiments, people appear to behave in a roughly Bayesian fashion, with some biases, such as a tendency to behave more myopically than is optimal.

These empirically-driven theories of individual perception and decision-making to have prompted us to develop a representation of consumer response to uncertain quality that is a myopic version of a the classical, utility-maximizing Bayesian. In our model, an individual consumer has repeated contacts with a supplier of a product or service. In each contact, she obtains utility that is driven by the quality of the encounter. She categorizes suppliers as being potentially either "bad" or "good," and for each supplier she maintains a primitive prior belief: it may be bad or good with a simple probability. After every encounter with a supplier, the consumer updates her belief concerning the supplier's quality in accordance with her experience. If the posterior probability that a current supplier's quality is good drops below that of a competitor's, then the consumer "defects" to the competitor.

This model of decision-making is similar to the cognitive theory of Gigerenzer and Murray [16]. It also has an intimate connection with the sequential probability ratio test (SPRT), developed by Wald [54] in the 1940's, as well as with the "Cumulative Utility Consumer Theory" of Gilboa and Schmeidler [18] and Gilboa and Pazgal [17]. 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 [54] and [17].

Our results offer closed-form expressions that characterize the duration of customer loyalty in the short term, as well as relative frequency of purchase in the long run. Theorem 2 shows that, for a current customer, the expected duration of loyalty is roughly convex and increasing in the overall level of quality provided by a supplier. Theorem 4 shows that, in the long run, the relative frequency with which a consumer patronizes a supplier is characterized in similar terms; relative "customer share" is convex and increasing in the overall level of quality. The expression for long-term frequency is a clear composition of the competing suppliers' short-term expected loyalties, and its form also bears a striking resemblance that of the multinomial logit (MNL) models of traditional discrete choice theory (for example, see Anderson *et al.* [3]).

These results appear to be quite robust. Our main results do not require the distribution of utility offered by a supplier to follow a particular functional form. Furthermore, we demonstrate that alternative representations of myopic consumers obtain analogous results.

Though stylized, our model and results are also rich enough to admit explicit representation of phenomena such as differences in consumer discrimination and in prior beliefs. In a companion paper [15] we demonstrate how changes in these parameters affect customer loyalty, and we use the expressions developed in this paper as the basis for normative models intended to help suppliers to develop effective marketing and operating strategies.

The remainder of the paper is organized as follows. Section 2 reviews related research, and §3 provides some relevant background on exponential families of distributions. After these preliminaries, we develop our model in §4 and §5. In §6, we characterize the duration of customer loyalty in the short run, and in §7 we characterize relative choice frequencies in the long run. Finally, §8 discusses the robustness of our model, as well as further research and results.

2 Literature Review

Over the years, there have been attempts to explicitly model the effect of poor quality on the distribution of demand. An early example is Schwartz [43, 44, 45] who develops an inventory model in which a stockout causes a short-term reduction in demand, due to customer dissatisfaction, and this effect gradually recedes as customers "forget" the stockout and drift back to their regular order size. This model of dissatisfaction and forgetting does not agree with the learning process that we believe underlies consumer reactions to uncertain quality.

Other inventory models have modeled the effect of stockouts on the aggregate distribution of demand facing an individual supplier. Balcer [4, 5], Fergani [14], and Robinson [42] develop models in which current stockouts negatively affect the next period's demand distribution. In all of these models the analysis begins with the postulation of the behavior of aggregate demand, however. They do not include detailed consideration of individual consumers' behavior.

There is also a stream of game-theoretic analyses of queueing systems that addresses the fact that customers may adapt to system congestion, often by balking or by abandoning the queue before they are provided service. Earlier examples of this work, such as that by Mendelson and Whang [36] and Stidham [51] analyze the case of a monopolist service provider. Later work, such as that of Kalai *et al.* [29], Li and Lee [33], Lederer and Li [32] and Ho and Zheng [26], perform analyses that capture well the competitive externalities that affect a firm's choice of service level. All of these works assume that consumers are well-informed about the (expected) waits at supplier' queues and that, at each arrival to the market, the consumer simply chooses the queue that maximizes her expected utility. In our model, however, there is an important information asymmetry: the consumer is not well-informed about the utility to be obtained. Rather, she must choose among suppliers on the basis of her subjective beliefs regarding the distribution of utility that they offer.

Two more general models do attempt to more explicitly treat the effect of information asymmetries between supplier and consumer. Smallwood and Conlisk [50] develop a model in which consumers react to a product failure by switching to a competitor, and they characterize equilibrium market shares for competitors that are based on both the competitors' breakdown probabilities and the consumers' switching rules. More recently, Hall and Porteus [23] develop a similar model in which inventory and queueing systems compete on the basis of service quality. They show that, when one firm has an advantage of more loyal customer's – with lower probabilities of switching upon a service failure – then it is optimal for the firm to offer a lower level of quality. While both of these models are close in spirit to ours, they differ in that they model customer switching behavior – loyalty – as exogenously driven. We wish to understand what drives loyalty in the first place, and our model attempts to show how loyalty is mediated by customer learning.

There is also research in the economics literature that more directly models the effect of learning on consumer choice. Papers by Hey and McKenna [25], Allen and Faulhaber [1, 2], and Thoman [52], model customer learning as taking place in one trial. As a result, they do not provide insight that we seek into the effect of the distribution of quality on the duration of customer loyalty. In contrast, Bergemann and Välimäki [8] develop an infinite horizon model. Here, however, both suppliers and the consumer share the same prior distributions concerning quality. In our model of consumer learning, in contrast, information asymmetries between a supplier and its customers are the foundation for quality competition.

Finally, there are brand-switching models in the marketing literature that explicitly account for both consumer learning and uncertainty in product quality. The work of Gilboa and Pazgal [17] falls into this category. In addition, papers by Roberts and Urban [41] and Erdem and Keane [12] develop models in which Bayesian learning about product attributes is embedded within MNL representations of consumer choice. While the Bayesian learning described in these papers is analogous to the learning developed in our model, their use of MNL-based choice is not appropriate in our setting. Neither of the two possible sources of (Gumbel-distributed) noise in utility that motivate the MNL model is present in our case (see Anderson *et al.* [3]). First, randomness in realized utilities is already explictly represented in a quite general fashion in our model; for the consumer, there are no utility "externalities" that would generate additional noise. Second, we are not interested in empirically estimating utilities or the choice probabilities of any one selection, so we do not need to account for noise that arises due measurement error.

3 Exponential Families of Distributions

While our fundamental results hold without regard to distributional form, we will sometime extend them by assuming that certain distributions are members of an *exponential family.* A distribution that comes from an exponential family, indexed by $\theta \in \Theta$, may be defined as

$$dF(x|\theta) \quad \stackrel{\triangle}{=} \quad e^{\{\theta x - \psi(\theta)\}} dF(x) \, , \tag{1}$$

where $\psi(\theta) \stackrel{\triangle}{=} \log \int e^{\theta x} dF(x)$ is the *cumulant generating function* of some nondegenerate distribution function, F. Most commonly used distributional forms are exponential families. Examples of one-dimensional families include the following: Bernoulli, exponential, Poisson, normal with mean θ and fixed variance, and normal with fixed mean and variance θ.

A random variable, X, whose distribution comes from a one-dimensional exponential family, indexed by parameter $\theta \in \Theta \subseteq \mathbf{R}$, has mean $\mathsf{E}_\theta[X] = \psi'(\theta)$ and variance $\text{var}_\theta(X) = \psi''(\theta)$. Thus, whenever the underlying family of distributions is not degenerate, $\psi(\cdot)$ is strictly convex, and the mean, $\psi'(\cdot)$, is a monotonically increasing function of θ.

Recall that the *likelihood ratio* of distribution F_1 with respect to distribution F_2, $dF_1(x)/dF_2(x)$, describes the conditional odds, given the observation x, that X is defined by F_1 rather than F_2. Furthermore, the likelihood ratio is *monotone* if it is nonincreasing or nondecreasing in x. It is a fact that, for two distributions that are members of the same, one-dimensional exponential family, the likelihood ratio is monotone:

Lemma 1

If $F(x|\theta^1)$ and $F(x|\theta^2)$ are members of the same, one-dimensional exponential family, then $\theta^1 \leq \theta^2$ implies

$$\frac{dF(x|\theta^1)}{dF(x|\theta^2)} \quad = \quad e^{\{(\theta^1 - \theta^2)x - (\psi(\theta^1) - \psi(\theta^2))\}} \tag{2}$$

is nonincreasing in x.

4 A Model of a Bayesian Consumer

We begin by developing a model of a Bayesian consumer's response to uncertain quality. While we will not directly analyze this Bayesian model, it provides an important first step in the development of our model of a more myopic, "simple" customer. Furthermore, having clearly defined the Bayesian consumer's response, we will be able to offer a rich contrast between it and our model of consumer behavior.

4.1 The Suppliers and Their Choices of Quality Level

A consumer has a periodic need for a product or service, and at regular time intervals $t = 1, 2, \ldots$ she acquires the good from one of m possible suppliers, indexed $i \in \{1, \ldots, m\}$. Because there is inherent uncertainty in the process of delivering the product or service, the utility of the good supplied to the consumer at a time, t, is a random variable, U_t.

While the quality of each item or service encounter is uncertain, a supplier *can* control the overall level of the utility it provides, and for each supplier this choice of a "quality level" manifests itself in the choice of a distribution for U_t. 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.

We assume that each provider's choice of quality level is a strategic decision. That is, supplier i must choose its particular level of quality, $\theta^i \in \Theta$, independently of its competitors, before consumers enter the market. We let $\pi(t)$ denote the supplier chosen by the consumer at time t and assume that, given a fixed θ^i, the utilities obtained from supplier i, $\{U_t : \pi(t) = i\}$, are *i.i.d.*with distribution $F_{\theta^i} \triangleq F(u|\theta^i)$. We also assume that the distributions parameterized by $\theta \in \Theta$ are ordered so that $\theta^i \geq \theta^j$ implies that the $\mathsf{E}[U_t|\pi(t) = i] \geq \mathsf{E}[U_t|\pi(t) = j]$.

4.2 The Consumer's Response to Uncertain Quality

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 anew 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 $\{P_0^1, \ldots, P_0^m\}$ be the "prior" information the consumer has as she enters the market for the first time.

After each contact she uses the new sample of utility, U_t, and Bayes 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

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

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

Let a *policy* $\pi = \{\pi(1), \pi(2)\ \ldots\}$ be a sequence of choices of suppliers, and let Π be the class of policies that is nonanticipating with respect to future rewards. Then the consumer seeks a policy, $\pi \in \Pi$, that will maximize the expected discounted value of the future stream of utilities, $\sup_{\pi\in\Pi} \mathsf{E}_\pi \left[\sum_{t=1}^{\infty} \alpha^t\, U_t\right]$, where $\alpha \in (0,1)$ is the one-period discount rate.

4.3 The Consumer's Problem as a Multi-Armed Bandit

It is well known that for any fixed set of service levels, $\{\theta^1, \ldots, \theta^m\}$, and priors, $\{P_0^1, \ldots, P_0^m\}$, 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, $\pi(t) = i$, at which the arm is played, and, from (3) we see that when the arm is played, its evolution is Markov.

Gittins and Jones [20] 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; 2) at any time t it is optimal for the consumer to use the supplier with the highest Gittins index.

Gittins [19] further characterized the index of supplier i as the result of maximizing expected discounted utility per unit of expected discounted time,

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

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.

Observe that the characterization of the Gittins index is that of an optimal stopping problem for the consumer's use of supplier i. Furthermore, the value of the Gittins index explicitly accounts for the fact that the consumer has the option to stop sampling from i and switch to another supplier if the sample information is unfavorable. Because of this,

$$G(P_t^i) \;\geq\; \frac{\sum_{s=t+1}^{\infty} \alpha^{s-t}\, \mathsf{E}[U(P_t^i)]}{\sum_{s=t+1}^{\infty} \alpha^{s-t}} \;=\; \mathsf{E}[U(P_t^i)]\,, \tag{5}$$

and the inequality is likely to be strict.

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 two.

In practice, the calculation of the Gittins index is a formidable task. In particular, 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, and over the years effort has been invested in developing effective, simple (computable) approximations to the Gittins index (for example, see Lai and Chang [31]).

5 A Model of a "Simple" Consumer

For our purposes, difficulty 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 introduction 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. Keener [30] has previously called a special case of this problem the multi-armed bandit with "simple arms."

The model for simple consumers works as follows. As in the original model, each of the m suppliers chooses a quality level, $\theta^i \in \Theta$. Rather than maintaining a complex set of beliefs concerning suppliers, however, the customer partitions the possible quality levels into two categories – good and bad – with respective utility distributions F_G and F_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 the 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 $\{U_1, U_2, \ldots\}$ are obtained. Then from a direct application of Bayes' rule (3) we find the posterior probability that i is good is

$$p_1^i \;=\; \frac{p_0^i \, dF_G(U_1)}{p_0^i \, dF_G(U_1) + (1-p_0^i)\, dF_B(U_1)} \;=\; \left[1 \;+\; \left(\frac{1-p_0^i}{p_0^i}\right) \; \frac{dF_B(U_1)}{dF_G(U_1)}\right]^{-1} \quad ,$$

where p_0^i denotes the consumer's prior belief concerning the quality of i. In turn, t visits yield

$$p_t^i \;=\; \left[1 \;+\; \left(\frac{1-p_0^i}{p_0^i}\right) \; \prod_{s=1}^{t} \frac{dF_B(U_s)}{dF_G(U_s)}\right]^{-1} \quad . \tag{6}$$

5.1 Relationship Between Bayesian and Simple Consumers

When the suppliers' choices of quality levels are, themselves, simple – $\Theta = \{\theta_B, \theta_G\}$ – 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 (Keener [30] and Banks and Sundaram [6])
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.

5.2 Modeling F_G and F_B

The distributions F_G and F_B need *not* be two members of the same class of distributions. Mathematically, they need only be mutually absolutely continuous; otherwise they may differ arbitrarily from each other. Psychologically, they represent the consumer's "exemplars" of good and bad quality distributions, and these exemplars need not have a direct relationship to Θ. For example, the consumer's expectations for one set of suppliers, who choose among quality levels in Θ, may have been developed through her experience with another set of suppliers, whose quality choices may have been different than Θ. Nevertheless, there are two methods that a consumer might intuitively use Θ to construct F_G and F_B, both of which originated with Wald [54] in the 1940's.

Method 1. An elementary method of constructing the distributions would be to pick two elements of Θ, θ_G and θ_B that represent thresholds for judging whether or not supplier i's performance is acceptable. If the consumer were to know with certainty that $\theta^i \geq \theta_G$ then she would be *satisfied* with i's level of quality and she would remain loyal for all time. Conversely, if she were to know that $\theta^i \leq \theta_B$, then the consumer would be *dissatisfied* and, if given a favorable alternative, would defect immediately.

Method 2. A more elaborate method would use the prior, P_0^i, of the Bayesian model of §4.2 to construct the two marginal distributions, F_G and F_B. Here θ^* might represent a threshold above which the consumer is satisfied and below which she is dissatisfied, so that

$$dF_G \triangleq \frac{\int_{\theta \geq \theta^*} dF(u|\theta) dP_0^i(\theta)}{1 - P_0^i(\theta^*)} \quad \text{and} \quad dF_B \triangleq \frac{\int_{\theta < \theta^*} dF(u|\theta) dP_0^i(\theta)}{P_0^i(\theta^*)} . \tag{7}$$

In either case, if the underlying family of distributions parameterized by θ have monotone likelihood ratios, then F_B and F_G will as well. That is,

Lemma 2

Suppose the consumer constructs F_G and F_B using either Method 1 or Method 2. If $\theta^i < \theta^j$ implies $dF(u|\theta^i)/dF(u|\theta^j)$ is nonincreasing in u for all $\theta^i, \theta^j \in \Theta$, then $dF_B(u)/dF_G(u)$ is nonincreasing in u.

Proof

For *Method 1*, monotonicity holds directly by assumption. For *Method 2*, see Theorem 1.C.11 of Shaked and Shanthikumar [49]. ○

Together, Lemmas 1 and 2 show that, whenever F_B and F_G are members of the same one-dimensional exponential family of distributions, then $dF_B(u)/dF_G(u)$ is monotonically decreasing in u. For a consumer attempting to discern between good and bad quality, this monotonicity is both a natural and appealing property for the likelihood ratio to have. When the ratio, $dF_B(u)/dF_G(u)$ is nonincreasing in u, then a *better* experience at a supplier will consistently lead the consumer to believe that the service provider is *more likely* to be "good."

We note that this likelihood ratio representation of consumer response is consistent with the elements (though not the mathematical structure) of the satisfaction model developed by Boulding *et al.* [9]. More specifically, the probability, p^i_{t-1}, represents the consumer's belief, before the t^{th} trial, concerning the level of quality that supplier i *will* offer. At the t^{th} service encounter the consumer updates her judgment of her satisfaction, p^i_t, directly through the likelihood ratio, $dF_B(U_t)/dF_G(U_t)$, which compares the actual quality received to her conception of what a supplier *should* offer: a realization of $dF_B(U_t)/dF_G(U_t)$ that is greater than one is evidence that the supplier is bad, rather than good, and causes a decrease in her posterior judgment of satisfaction, p^i_t; a realization of $dF_B(U_t)/dF_G(U_t)$ that is less than one is evidence that the supplier is good, rather than bad, and causes an increase in her posterior judgment of satisfaction.

6 The Duration of Customer Loyalty

Keener [30] exploits the fact that Bandit problems are one-sided SPRT's to develop an explicit, closed-form expression for the Gittins index that is monotonically increasing in p^i_t. In the case of a simple consumer, the probability may be used as a satisfaction (or Gittins) index, and it is not difficult to show that consumer's choice problem is essentially a one-sided sequential test.

In this section, we develop the SPRT and its associated random walk. We then use classic results from the theory of sequential testing and of random walks to derive simple bounds on both the expected duration of customer loyalty and on the probability that a customer defects. The form of the bounds implies that the expected duration of customer loyalty is convex and increasing in $E_{\theta^i}[U]$. Finally we present examples – for the exponential, Bernoulli, and normal distributions – in which we refine the bounds and show that the rough characterization provided by the original bounds appears to be robust.

6.1 The Embedded SPRT and Random Walk

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 (6) this is equivalent to

$$\prod_{s=1}^{t} \frac{dF_B(U_s)}{dF_G(U_s)} \quad < \quad \left(\frac{p_0^i}{p_0^j}\right) \times \left(\frac{1-p_0^j}{1-p_0^i}\right) , \tag{8}$$

a one-sided sequential probability ratio test with upper bound $(p_0^i/p_0^j)\left((1-p_0^j)/(1-p_0^i)\right) > 1$.

By taking logs on both sides of (8), we equivalently have a random walk, $S_t = \sum_{s=1}^{t} X_s$, with $i.i.d.$increments

$$X_s \quad \triangleq \quad \log\left(dF_B(U_s)/dF_G(U_s)\right) \tag{9}$$

and stopping boundary

$$b^i \quad \triangleq \quad \log\left(\left(\frac{p_0^i}{p_0^j}\right) \times \left(\frac{1-p_0^j}{1-p_0^i}\right)\right) . \tag{10}$$

Here, the expectation of X, $\mathsf{E}_{\theta^i}[X]$, is evaluated with respect to the utility distribution actually offered by supplier i.

Let

$$\tau \;\triangleq\; \inf\{t:\; p_t^i < p_0^j\} \;=\; \inf\{t:\; S_t > b^i\} \tag{11}$$

be the time at which the customer first perceives that the quality at i is inferior to the quality at j. This, the time at which the consumer "defects" to the competition, is our primary object of study.

There is a vast body of literature, beginning with Wald [54], that characterizes the behavior of τ. We next present some basic results concerning $\mathsf{E}_{\theta^i}[\tau]$ and $P_{\theta^i}\{\tau < \infty\}$, which hold in great generality.

6.2 The Expected Duration of Customer Loyalty

Our characterization of the expected time to defection follows classic results of Wald [54] and Lorden [34]:

Theorem 2

If $0 < \mathsf{E}_{\theta^i}[X] < \infty$*, then*

$$\mathsf{E}_{\theta^i}[\tau] \;=\; \frac{\mathsf{E}_{\theta^i}[S_\tau]}{\mathsf{E}_{\theta^i}[X]} \;=\; \frac{b^i \;+\; \mathsf{E}_{\theta^i}[S_\tau - b^i]}{\mathsf{E}_{\theta^i}[X]}\,. \tag{12}$$

In addition,

$$\frac{b^i}{\mathsf{E}_{\theta^i}[X]} \;\le\; \mathsf{E}_{\theta^i}[\tau] \;\le\; \frac{b^i}{\mathsf{E}_{\theta^i}[X]} \;+\; \frac{\mathsf{E}_{\theta^i}[(X^+)^2]}{\mathsf{E}_{\theta^i}[X]^2} \;\le\; \frac{b^i}{\mathsf{E}_{\theta^i}[X]} \;+\; \frac{\mathsf{E}_{\theta^i}[X^2]}{\mathsf{E}_{\theta^i}[X]^2}\,, \tag{13}$$

the upper bounds holding whenever (respectively) $\mathsf{E}_{\theta^i}[(X^+)^2] < \infty$ *and* $\mathsf{E}_{\theta^i}[X^2] < \infty$*.*

Proof

Since τ is a stopping time with respect to S_t, when $\mathsf{E}_{\theta^i}[X]$ is finite Wald's identity can be applied to derive (12). The lower bound in (13) directly follows from (12). The term $\mathsf{E}_{\theta^i}[S_\tau - b^i]$ on the right hand side of the second equation in (12) is commonly called the *excess over the boundary*, and Theorem 1 of Lorden [34] demonstrates that if $\mathsf{E}_{\theta^i}[X] > 0$ and $\mathsf{E}_{\theta^i}[(X^+)^2] < \infty$, then $\mathsf{E}_{\theta^i}[S_\tau - b^i] \le \mathsf{E}_{\theta^i}[(X^+)^2] \,/\, \mathsf{E}_{\theta^i}[X]$. The first upper bound in (13) then follows, and the fact that $\mathsf{E}[(X^+)^2] \le E[X^2]$ provides the second. ○

When are $\mathsf{E}_{\theta^i}[X]$ and $\mathsf{E}_{\theta^i}[X^2]$ 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 (9) we have

$$X \;=\; \log\left(\frac{e^{\{\theta_B U - \psi(\theta_B)\}}dF}{e^{\{\theta_G U - \psi(\theta_G)\}}dF}\right) \;=\; \psi(\theta_G) - \psi(\theta_B) \;-\; (\theta_G - \theta_B)\,U\,. \tag{14}$$

In turn, we have

$$\mathsf{E}_{\theta^i}[X] \;=\; \psi(\theta_G) - \psi(\theta_B) \;-\; (\theta_G - \theta_B)\,\mathsf{E}_{\theta^i}[U] \tag{15}$$

and

$$\mathsf{E}_{\theta^i}[X^2] \;=\; (\psi(\theta_G)-\psi(\theta_B))^2 \;-\; 2\,(\psi(\theta_G)-\psi(\theta_B))(\theta_G-\theta_B)\,\mathsf{E}_{\theta^i}[U] \;+\; (\theta_G-\theta_B)^2\mathsf{E}_{\theta^i}[U^2]\,. \tag{16}$$

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

If we define

$$\mu^* \;\triangleq\; (\psi(\theta_G)-\psi(\theta_B))\,/\,(\theta_G-\theta_B) \tag{17}$$

then we also see $\mathsf{E}_{\theta^i}[X]\in(0,\infty)$ is equivalent to $\mathsf{E}_{\theta^i}[U]\in(-\infty,\mu^*)$. Furthermore, by substituting (15) into the lower bound of (13) and rearranging terms we obtain the following:

Corollary 1

When F_B and F_G are members of the same exponential family of distributions, then

$$\mathsf{E}_{\theta^i}[\tau] \;\geq\; \frac{b^i\,/\,(\theta_G-\theta_B)}{\mu^* \;-\; \mathsf{E}_{\theta^i}[U]}\,. \tag{18}$$

Differentiating the right hand side of (18) shows that, when F_G and F_B are members of the same exponential family of distributions, the lower bound of Theorem 2 is convex and increasing in $\mathsf{E}_{\theta^i}[U^i]$. In this case, it appears that the bound's behavior offers a good characterization of $\mathsf{E}_{\theta^i}[\tau]$ itself, and in §6.4 and in §6.5 we offer evidence to that effect.

6.3 The Probability of Defection

To determine the probability of defection, $P_{\theta^i}\{\tau<\infty\}$ requires a bit more work. A concise presentation of the analysis, based on Wald [54], can be found in Siegmund [47].

Theorem 3

If $\mathsf{E}_{\theta^i}[X]\neq 0$ and $\mathsf{E}[e^{\varphi X}]<\infty$ for all real φ then there exists a unique $\varphi(\theta^i)\neq 0$ and a probability distribution $F_{\varphi(\theta^i)}(u)$ such that

$$dF_{\varphi(\theta^i)}(u) \;=\; \left(\frac{dF_{\theta_B}(u)}{dF_{\theta_G}(u)}\right)^{\varphi(\theta^i)} dF_{\theta^i}(u)\,; \tag{19}$$

$$P_{\theta^i}\{\tau < \infty\} = e^{-\varphi(\theta^i)b}\, \mathsf{E}_{\varphi(\theta^i)}\left[e^{-\varphi(\theta^i)(S_\tau - b)}\right] \text{ ; and} \tag{20}$$

$$P_{\theta^i}\{\tau < \infty\} \leq e^{-\varphi(\theta^i)b} \tag{21}$$

Proof

For a proof of the existence and uniqueness of (19), see Lemma A.1 and equation (A:17) in Wald [54]. Note that the statement of Wald's lemma in [54] requires $P_{\theta^i}\{X > 0\} > 0$ and $P_{\theta^i}\{X < 0\} > 0$. In our case, $\mathsf{E}_{\theta^i}[X] \neq 0$ implies that $F_B \neq F_G$ and that the inequalities are satisfied.

The form of (20) follows from Wald's likelihood ratio identity, and a special case used for the testing of the simple hypothesis, F_B versus F_G, can be found in (8.3) of Siegmund [47]. A change of measure, based on (19), generalizes the identity to (20) when $F_{\theta^i} \neq F_B$.

The upper bound (21) also follows from Wald's likelihood ratio identity. It is succinctly demonstrated in (2.9) of Siegmund [47]. ○

Again, we would like to know when $\mathsf{E}_{\theta^i}[X] \neq 0$ and when $\mathsf{E}[e^{\varphi X}] < \infty$ for all real φ. For F_B and F_G that are members of the same exponential family of distributions, (15) and (17) show that the first condition holds whenever $\mathsf{E}_{\theta^i}[U] \neq \mu^*$. By substituting (14) into $\mathsf{E}[e^{\varphi X}]$, we see that the second condition holds whenever $\mathsf{E}_{\theta^i}[e^{\varphi U}] < \infty$ for all real φ.

6.4 Examples and Further Approximations

The bounds of Theorems 2 and 3 offer simple limits on the behavior or τ and will be useful in the sensitivity analysis performed in [15]. They need not, however, be particularly tight. Therefore, we would like to verify that the characterization of the lower bound offered in Corollary 1 holds for $\mathsf{E}_{\theta^i}[\tau]$ itself.

The following three examples – for the exponential, Bernoulli, and normal distributions – show that, for an important and varied set of distributional forms, the characterization of Corollary 1 appears to hold. In all three cases, when θ_G and θ_B are fixed and θ^i is bounded above by a constant (for example, by θ_G), then the "excess" term, $S_\tau - b^i$, which is neglected in the corollary can be bounded above by a constant that is independent of b^i. Thus, in these cases, we can construct an upper bound whose form mimics that of the lower bound of (18), and we may infer that the same behavior (roughly) holds for $\mathsf{E}_{\theta^i}[\tau]$, itself. For these distributions, the examples also show how one may compute the quantities presented in Theorems 2 and 3.

6.4.1 Exponential Distribution

One important case in which we may compute the excess exactly is when F_{θ^i} is exponentially distributed. For example, suppose that the supplier operates a queue and that the quality of its service is the delay in queue experienced by customers. Then longer waits have lower utility, and we might model a customer that is risk-neutral with regard to delay in queue as realizing utilities that are exponentially distributed on $(-\infty, 0]$, rather than $[0, \infty)$.

In this case, we have $\Theta = (-\infty, 0]$, $dF_{\theta^i}(u) = -\frac{1}{\theta^i} e^{-\frac{u}{\theta^i}}$ for $u \in (-\infty, 0]$, and $\mathsf{E}_{\theta^i}[U] = \theta^i < 0$. When F_B and F_G are members of the same exponential family of distributions (possibly distinct from F_{θ^i}) with support on $(-\infty, 0]$ we can use (14) to show that for $z > 0$

$$
\begin{aligned}
P_{\theta^i}\{X - y > z | X > y\} &= P_{\theta^i}\{X > y + z | X > y\} && (22)\\
&= P_{\theta^i}\left\{U < \frac{\psi(\theta_G) - \psi(\theta_B) - y - z}{\theta_G - \theta_B} \,|\, U < \frac{\psi(\theta_G) - \psi(\theta_B) - y}{\theta_G - \theta_B}\right\} && (23)\\
&= P_{\theta^i}\left\{U < -\frac{z}{\theta_G - \theta_B}\right\} && (24)\\
&= e^{-\frac{z}{|(\theta_G - \theta_B)\theta^i|}}, && (25)
\end{aligned}
$$

where (24) follows from the memoryless property of exponentially distributed U. Thus for any X_t and y, the excess $X_t - y$, is exponentially distributed, with mean $|(\theta_G - \theta_B)\theta^i|$. This, of course, applies when $t = \tau$ and $y = b^i$, so $\mathsf{E}_{\theta^i}[S_\tau - b^i] = |(\theta_G - \theta_B)\theta^i|$, and (12) becomes

$$
\mathsf{E}_{\theta^i}[\tau] = \frac{b^i + |(\theta_G - \theta_B)\theta^i|}{\psi(\theta_G) - \psi(\theta_B) + (\theta_G - \theta_B)\theta^i}. \qquad (26)
$$

For (20), the definition of X remains (14), but now $U_t \sim F_{\varphi(\theta^i)}$. From (19) we have:

$$dF_{\varphi(\theta^i)}(u) = \left(\frac{dF_B(u)}{dF_G(u)}\right)^{\varphi(\theta^i)} dF_{\theta^i}(u) \tag{27}$$

$$= \left(e^{[\psi(\theta_G)-\psi(\theta_B)-(\theta_G-\theta_B)u]\varphi(\theta^i)}\right)\left(-\frac{1}{\theta^i}e^{-\frac{u}{\theta^i}}\,du\right) \tag{28}$$

$$= -\left(\frac{1}{\theta^i}\,e^{(\psi(\theta_G)-\psi(\theta_B))\varphi(\theta^i)}\right)\left(e^{-[(\theta_G-\theta_B)\varphi(\theta^i)+\frac{1}{\theta^i}]u}\right)du\,. \tag{29}$$

Observe that the term in the first set of parentheses of (29) does not depend on u, while that in the second, $e^{-[(\theta_G-\theta_B)\varphi(\theta^i)+\frac{1}{\theta^i}]u}$, is the (un-normalized) density of an exponentially distributed random variable of mean $[(\theta_G-\theta_B)\varphi(\theta^i)+\frac{1}{\theta^i}]^{-1} < 0$. Note that $\int_{-\infty}^{0} dF_{\varphi(\theta^i)}(u) = 1$, in turn, requires that

$$-\left(\frac{1}{\theta^i}\,e^{(\psi(\theta_G)-\psi(\theta_B))\varphi(\theta^i)}\right) = -\left[(\theta_G-\theta_B)\varphi(\theta^i)+\frac{1}{\theta^i}\right], \tag{30}$$

so that $\varphi(\theta^i)$ is the unique solution to

$$\varphi(\theta^i) = \frac{e^{(\psi(\theta_G)-\psi(\theta_B))\varphi(\theta^i)}-1}{(\theta_G-\theta_B)\theta^i}\,. \tag{31}$$

Furthermore, using the same argument as that for (25) we can show that

$$P_{\varphi(\theta^i)}\{X-y>z|X>y\} = e^{\frac{z\,[(\theta_G-\theta_B)\varphi(\theta^i)+\frac{1}{\theta^i}]}{(\theta_G-\theta_B)}} = e^{-z\left(-[\varphi(\theta^i)+\frac{1}{(\theta_G-\theta_B)\,\theta^i}]\right)}\,. \tag{32}$$

Thus $X_t - y$ is exponentially distributed with mean $-\left[\varphi(\theta^i)+\frac{1}{(\theta_G-\theta_B)\theta^i}\right]^{-1}$, so that

$$\mathsf{E}_{\varphi(\theta^i)}\left[e^{\varphi(\theta^i)(S_\tau-b)}\right] = \frac{-[\varphi(\theta^i)+\frac{1}{(\theta_G-\theta_B)\theta^i}]}{-[\varphi(\theta^i)+\frac{1}{(\theta_G-\theta_B)\theta^i}]-\varphi(\theta^i)} = \frac{(\theta_G-\theta_B)\,\theta^i\,\varphi(\theta^i)\;+\;1}{2\,(\theta_G-\theta_B)\,\theta^i\,\varphi(\theta^i)\;+\;1}\,, \tag{33}$$

and for (20) we have

$$P_{\theta^i}\{\tau<\infty\} = e^{-\varphi(\theta^i)b}\,\frac{(\theta_G-\theta_B)\,\theta^i\,\varphi(\theta^i)\;+\;1}{2\,(\theta_G-\theta_B)\,\theta^i\,\varphi(\theta^i)\;+\;1}\,. \tag{34}$$

6.4.2 Bernoulli Distribution

The Bernoulli distribution can be used to model products and services for which each encounter can be categorically labeled acceptable or unacceptable. This clearly applies to stockouts and product breakdowns, and it may also be used as a rough approximation for more complex quality distributions whenever the outcomes may be comfortably partitioned into these two categories.

Observe that when the utility obtained in each encounter follows a Bernoulli distribution, the distributions of F_B and F_G must, themselves, be Bernoulli. Therefore, let θ_B, θ_G, and θ^i respectively represent the probabilities that an individual encounter is acceptable, given the quality level offered is bad, good, and that actually offered by supplier i. Thus,

$$X_t = \begin{cases} \ln\left(\frac{\theta_B}{\theta_G}\right), & \text{when encounter } t \text{ is acceptable} \\ \ln\left(\frac{1-\theta_B}{1-\theta_G}\right), & \text{otherwise,} \end{cases} \tag{35}$$

and

$$E_{\theta^i}[X] = (1-\theta^i)\ln\left(\frac{1-\theta_B}{1-\theta_G}\right) + \theta^i \ln\left(\frac{\theta_B}{\theta_G}\right). \tag{36}$$

It is immediate from (11) that $X_\tau > 0$ and that this occurs only upon an unacceptable encounter, so it must be the case that $S_\tau - b^i \leq \ln\left(\frac{1-\theta_B}{1-\theta_G}\right)$. We have, therefore, the following simple bounds:

$$\frac{b}{(1-\theta^i)\ln\left(\frac{1-\theta_B}{1-\theta_G}\right) + \theta^i \ln\left(\frac{\theta_B}{\theta_G}\right)} \leq E_{\theta^i}[\tau] \leq \frac{b^i + \ln\left(\frac{1-\theta_B}{1-\theta_G}\right)}{(1-\theta^i)\ln\left(\frac{1-\theta_B}{1-\theta_G}\right) + \theta^i \ln\left(\frac{\theta_B}{\theta_G}\right)}. \tag{37}$$

We can also apply the fact that $S_\tau - b^i \leq \ln\left(\frac{1-\theta_B}{1-\theta_G}\right)$ to the expectation in (20) to develop a lower bound for $P_{\theta^i}\{\tau < \infty\}$, which gives us

$$e^{-\varphi(\theta^i) b^i \ln\left(\frac{1-\theta_B}{1-\theta_G}\right)} \leq P_{\theta^i}\{\tau < \infty\} \leq e^{-\varphi(\theta^i) b}. \tag{38}$$

Wald [54] develops analogous bounds for two-sided versions of Bernoulli SPRT's which are asymptotically equal to those shown here, as one of the stopping boundaries bound grows without bound. In addition, he shows that $\varphi(\theta^i)$ is the unique solution to

$$\theta^i = \frac{1 - \left(\frac{1-\theta_B}{1-\theta_G}\right)^{\varphi(\theta^i)}}{\left(\frac{\theta_B}{\theta_G}\right)^{\varphi(\theta^i)} - \left(\frac{1-\theta_B}{1-\theta_G}\right)^{\varphi(\theta^i)}}. \tag{39}$$

Wald [54] also develops a method, using difference equations, for exactly calculating $\mathsf{E}_{\theta^i}[\tau]$ and $P_{\theta^i}\{\tau < \infty\}$. (See also Chapter XIV of Feller [13].)

6.4.3 Normal Distribution with Unknown Mean and Fixed Variance

Suppose quality is normally distributed, and consider the case in which the mean level of quality varies from one supplier to the next, but the variability of quality is roughly the same for all suppliers. Then, without loss of generality, me may assume that supplier i's quality is normally distributed with mean θ^i and a variance of one, so that

$$\psi(\theta^i) \;=\; (\theta^i)^2/2\,. \tag{40}$$

If F_B and F_G are also normally distributed with unit variance – that is, the consumer discriminates among distributions using Method 1 of §5.2 – then substituting (40) into (15) we have

$$E_{\theta^i}[X] \;=\; (\theta_G - \theta_B)\left(\frac{\theta_B + \theta_G}{2} - \theta^i\right). \tag{41}$$

For $\mathsf{E}_{\theta^i}[X] > 0$ ($\theta^i < \frac{1}{2}(\theta_B + \theta_G)$), we substitute (41) into (12) and can apply the bounds developed by Wald [54] for normally distributed two-sided SPRT's to show

$$\frac{b}{(\theta_G - \theta_B)\,\lambda(\theta^i)} \;\leq\; E_{\theta^i}[\tau] \;\leq\; \frac{b^i + (\theta_G - \theta_B)\left(\frac{\phi(\lambda(\theta^i))}{\Phi(\lambda(\theta^i))} + \lambda(\theta^i)\right)}{(\theta_G - \theta_B)\,\lambda(\theta^i)} \tag{42}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ respectively represent the density and cumulative distribution functions of the standard normal distribution and $\lambda(\theta^i) \triangleq \frac{1}{2}(\theta_B + \theta_G) - \theta^i$.

Note that the upper bound in (42) may be written

$$\frac{b^i + (\theta_G - \theta_B)\left(\frac{\phi(\lambda(\theta^i))}{\Phi(\lambda(\theta^i))}\right)}{(\theta_G - \theta_B)\,\lambda(\theta^i)} \;+\; 1\,. \tag{43}$$

Furthermore, for $\theta^i < \theta_G$, $\lambda(\theta^i)$ is bounded below by $-\frac{1}{2}(\theta_G - \theta_B)$, and $\left(\frac{\phi(\lambda(\theta^i))}{\Phi(\lambda(\theta^i))}\right)$ is, in turn, bounded above by a constant.

Wald [54] also shows that for normally distributed random variables

$$\varphi(\theta^i) \;=\; \frac{\theta_B + \theta_G - 2\,\theta^i}{\theta_G - \theta_B}\,, \tag{44}$$

and we can use Wald's results for two-sided SPRT's of normally distributed random variables to show

$$\frac{\Phi(-|\lambda(\theta^i)|)}{\Phi(|\lambda(\theta^i)|)}\,e^{-\varphi(\theta^i)b} \;\leq\; P_{\theta^i}\{\tau < \infty\} \;\leq\; e^{-\varphi(\theta^i)b}\,. \tag{45}$$

Finally, we note that as $b^i \to \infty$, the distribution of the excess over the boundary, $S_\tau - b^i$, approaches that of the equilibrium distribution of the "ladder height" distribution of the underlying random walk. Siegmund [46, 47] uses the expectation of the equilibrium distribution to develop an approximation for $\mathsf{E}_{\theta^i}[S_\tau - b^i]$, which as $\theta^i \to 0$ is

$$\mathsf{E}_{\theta^i}[S_\tau - b^i] \quad \approx \quad 0.583 \tag{46}$$

for a normal distribution. Without loss of generality, the means of the relevant distributions – θ_B, θ_G, and θ^i – may then be translated so that $\theta^i = 0$, and (46) can be substituted into (12) as an approximation. (The techniques developed in [47] can be applied more generally, and the interested reader may consult the book.)

6.5 Relationship to Other Myopic Models

Our model of the simple consumer is, in large measure, motivated by the assumption that the consumer uses the likelihood ratio, $dF_B(u)/dF_G(u)$, as method of discrimination. Nevertheless, the resulting behavior, which takes the form of a random walk, $\{S_t\}$, is consistent with other myopic models of Bayesian consumers.

More specifically, consider a consumer that is capable of maintaining complex priors and of performing complex Bayesian posterior updates, yet exhibits myopic behavior. At the t^{th} trial she chooses the supplier, i, whose prior distribution, P^i_{t-1}, maximizes the expected utility of her next trial, $E[U(P^i_{t-1})]$ (see the right hand side of (5)). When the form of the prior is conjugate, it is possible to show that the time until customer defection takes the form of a random walk.

For example, suppose the distributions of utilities provided by suppliers are normally distributed with unknown means and fixed standard deviations, ς. If the consumer's initial prior for the mean of supplier i's distribution, P^i_0 is normally distributed with mean μ^i_0 and standard deviation σ^i_0, then her posterior distribution for the mean, P^i_t, is also normally distributed with mean

$$\frac{\mu^i_0/(\sigma^i_0)^2 \;+\; (\sum_{s=1}^{t} U^i_s)/\varsigma^2}{1/(\sigma^i_0)^2 \;+\; t/\varsigma^2}\;. \tag{47}$$

For $\mu_0^i > \mu_0^j$ we define the time at which she will "defect" from i to j to be

$$\xi \triangleq \inf\left\{t \,:\, \frac{\mu_0^i/(\sigma_0^i)^2 \,+\, (\sum_{s=1}^t U_s^i)/\varsigma^2}{1/(\sigma_0^i)^2 \,+\, t/\varsigma^2} \;<\; \mu_0^j\right\} \;=\; \inf\left\{t \,:\, \sum_{s=1}^t (U_s^i \,-\, \mu_0^j/\varsigma^2) \;<\; (\mu_0^j \,-\, \mu_0^i)/(\sigma_0^i)^2\right\}. \quad (48)$$

Then using the approach of Theorem 2, we can derive the following lower bound on the expected stopping time:

$$\mathsf{E}[\xi] \;\;\leq\;\; \frac{(\mu_0^i - \mu_0^j)/(\sigma_0^i)^2}{\mu_0^j/\varsigma^2 \;-\; \mathsf{E}[U^i]}\,, \qquad (49)$$

whenever $\mathsf{E}[U^i] \;<\; \mu_0^j/\varsigma^2$. This bound has same form as that of Corollary 1.

Our representation of consumer decision making is also intimately connected with a special case of the Cumulative Utility Consumer Theory (CUCT) proposed by Gilboa and Schmeidler [18]. As we shall demonstrate in the following section, our model begins with a more detailed characterization of consumer response to quality variation and may be seen as an enrichment of the theory in the context of quality uncertainty.

7 Long Run Choice Frequencies

Theorems 2 and 3 offer a short-term characterization of customer loyalty. That is, they describe the duration of loyalty of a *current* customer. In some competitive environments, this characterization may be sufficient; a customer that is lost is gone forever, and there is little chance that she will return. In other markets, however, the short-term characterization in clearly incomplete. In the long run, consumers that defect to the competition may return.

In this section we develop a characterization for the long-run case that is analogous to the results developed for the short run. The results show that long-run frequency with which a consumer patronizes supplier i is described by a simple ratio among the various supplier's $\mathsf{E}_{\theta^i}[X]$'s. The results also closely parallel those developed by Gilboa and Pazgal [17]. We make this connection explicit and apply their results to develop our own.

7.1 Relationship to Cumulative Utility Consumer Theory

We consider the cardinal version of the Cumulative Utility model presented in Gilboa and Pazgal [17]. Here, the consumer maintains cumulative balances of the utilities obtained from each of the suppliers. When she enters the market, she may have an initial utility balance for supplier i, b_0^i, that represents her prior impression of i. Each encounter with supplier i yields a randomly distributed instantaneous utility, Y^i, which is added to i's balance. Letting $S_t^i = \sum_{s=1}^t Y_s^i$, we see that t trips to i yield a balance $b_0^i + S_t^i$.

The consumer behaves myopically; in each period she visits the supplier with the highest cumulative balance. Thus, for $b_0^i > b_0^j$, the stopping time at which the consumer defects from i to j is

$$\tau \triangleq \inf\left\{t : b_0^i + S_t^i < b_0^j\right\} = \inf\left\{t : -S_t^i > b_0^i - b_0^j\right\} . \tag{50}$$

If we let $b^i = b_0^i - b_0^j$ and let $-Y^i = X^i$, where X^i is defined as in (9), then we recover our original model for the simple consumer.

Note that the intention of the Cumulative Utility models somewhat different from ours. In Gilboa and Pazgal's [17] formulation, a consumer makes repeated choices among a number of competing brands, and switching behavior is induced is due to "variety seeking," rather than inherent defect of the product: the more negative $\mathsf{E}[Y^i]$, the more easily bored the consumer becomes on average with choice i; randomness in the realizations of Y^i correspond to unpredictable differences in the consumers response to i, rather than any (measurable) uncertainty in the actual performance of i.

7.2 Basic Results

Consider a competitive market in which there are m suppliers competing for the patronage of a consumer. Then

Theorem 4 (Gilboa and Pazgal [17])

Let $f_i \triangleq \lim_{t\to\infty} \frac{1}{t}\sum_{s=1}^{t} \mathbf{1}\{\pi(s)=i\}$ *denote the limiting relative frequency with which the consumer chooses supplier* i*. If* $\mathsf{E}_{\theta^i}[Y] < 0$ *and* $\mathrm{var}_{\theta^i}(Y) < \infty$ *for* $i = 1,\ldots,m$*, then with probability one* f_i *exists for* $i = 1,\ldots,m$*, and*

$$f_i \;=\; \frac{\prod_{j\neq i} \mathsf{E}_{\theta^j}[Y]}{\sum_{k=1}^{m} \prod_{j\neq k} \mathsf{E}_{\theta^j}[Y]}\,. \tag{51}$$

If we multiply both numerator and denominator by $-\prod_{j=1}^{m} \mathsf{E}_{\theta^j}[Y]$ and note that $X \equiv -Y$, then we recover

Corollary 2

If $\mathsf{E}_{\theta^i}[X] > 0$ *and* $\mathrm{var}_{\theta^i}(X) < \infty$ *for* $i = 1,\ldots,m$*, then with probability one* f_i *exists for* $i = 1,\ldots,m$*, and*

$$f_i \;=\; \frac{\mathsf{E}_{\theta^i}[X]^{-1}}{\sum_{j=1}^{m} \mathsf{E}_{\theta^i}[X]^{-1}}\,. \tag{52}$$

Furthermore, if F_B *and* F_G *are members of the same family of exponential distributions, then* $E_{\theta^i}[X] > 0 \iff \mathsf{E}_{\theta^i}[U] < \mu^*$*,* $\mathrm{var}_{\theta^i}(X) < \infty \iff \mathrm{var}_{\theta^i}(U) < \infty$*, and*

$$f_i \;=\; \frac{1/(\mu^* - \mathsf{E}_{\theta^i}[U])}{\sum_{j=1}^{m} 1/(\mu^* - \mathsf{E}_{\theta^j}[U])}\,. \tag{53}$$

Observe that (53) describes the long-run frequencies solely in terms of the expected satisfaction with an individual service encounter, $\mu^* - \mathsf{E}_{\theta^i}[U]$. It is not surprising that the consumer's priors beliefs, $\{b^1,\ldots,b^m\}$, do not have a long-term effect. What is more interesting, however, is the fact that the excess, $S_\tau - b^i$, does not contribute to the long-run frequencies.

It is also interesting to note that, if we let $z_i = 1/\mathsf{E}_{\theta^i}[X]$, then

$$f_i \;=\; \frac{z_i}{\sum_{j=1}^{m} z_j} \tag{54}$$

bears a striking resemblance to the choice probability derived from classic MNL models of discrete choice. While the form of the expression is the same, the model of consumer behavior which generates it is quite different (see Andersen *et al.* [3]).

8 Conclusion

What is the cost of a stockout or of making a customer wait a bit long in queue? To answer this question, we have formulated a model of consumer response and developed expressions that characterize short and long term loyalty.

In [15] we use the results developed here to provide a fuller answer. We offer a rich characterization of how the expected duration of customer loyalty changes with the overall level of service, with the customer's discrimination function, and with the consumer's prior beliefs. For the long-term frequency of purchase, we present analogous results for the overall level of service and for the customer's discrimination function. We also use the expressions as the basis of simple optimization models and games that are of normative value for suppliers.

Of course, the value of all of these results rests on the underlying adequacy of our model of consumer behavior. In particular, even if one accepts the assumption that people behave in roughly Bayesian fashion, but that they are "more" myopic than is optimal, she one might wonder if our model is "too" myopic, that people are more foresighted than our representation permits.

A natural analysis to perform would be to compare the behavior of τ in our model to that predicted for a rational Bayesian. Then, if the behavior of the fully Bayesian bound on behavior is similar to that shown in our model, then we might conclude that we have adequately characterized the behavior itself. Difficulty in calculating the Gittins index requires us to make the comparison using numerical methods, and this work is beyond the scope of this paper. Nevertheless, for our characterization of the expected time to defection, we note two pieces of evidence that point to just such a result.

First, Erdem and Keane [12] develop two models in which Bayesian learning is embedded in MNL representations of choice: in the first, the consumer is a myopic Bayesian, and in the second, more farsighted. In empirical tests in which they fit both models to the same panel data, the authors report small (but statistically significant) differences in the parameter estimates generated by the two models. Thus, for a family of models which are similar to ours – but for which the investigators' estimates of switching behavior are mediated by MNL-inducing noise – the bias induced by using a too-myopic model of consumer behavior appears to be small.

Second, Lai and Chang [31] develop an approximation to the Gittins index which, for quality distributions that are members of the same exponential family, they demonstrate is asymptotically Bayes as the discount rate approaches one. Furthermore, for a customer using the approximate index, as $\alpha \to 1$ the expected time to defection converges to

$$\lim_{\alpha \to 1} \mathsf{E}_{\theta^i}[\tau] \quad \approx \quad \frac{\log\left((\theta_{CE} - \theta^i)^2\right) \; + \log\left((1-\alpha)^{-1}\right)}{\psi(\theta_{CE}) - \psi(\theta^i) \; - \; (\theta_{CE} - \theta^i)\psi'(\theta^i)}\,, \tag{55}$$

where θ_{CE} denotes the quality level of the certainty equivalent to the (approximation to) the Gittins index of the best alternative to supplier i, and the convergence is uniform in $|\theta_{CE} - \theta^i|$ over some relevant range. (See (4.2) in [31].)

The statement of the result does not include the rate (in α) at which the convergence takes place. Nevertheless, as a "back of the envelope" analysis, we might divide both sides of (55) by $\log\left((1-\alpha)^{-1}\right)$ to derive an alternative form of the limit

$$\lim_{\alpha \to 1} \frac{\mathsf{E}_{\theta^i}[\tau]}{\log\left((1-\alpha)^{-1}\right)} \quad \stackrel{?}{=} \quad \frac{1}{\psi(\theta_{CE}) - \psi(\theta^i) \; - \; (\theta_{CE} - \theta^i)\psi'(\theta^i)}\,, \tag{56}$$

which is precisely the form of the denominator of the lower bound of (13) when $U \sim F_{\theta^i}$ and $X = \log(dF_{\theta^i}(U)/dF_{CE}(U))$.

Thus, when F_{θ^i} and F_{CE} are members of the same exponential family of distributions, then the behavior of (56) with respect to θ^i is the same as that of the simple consumer. As $\alpha \to 1$, the rate at which the expected duration of customer loyalty is growing is convex and increasing in θ^i.

One may also ask how our characterization of long-term relative purchase frequencies compares with that of the fully Bayesian consumer. If $E_{\theta^i}[U] < \mu^*$ 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, $dF_B(u)/dF_G(u)$, nor the dispersion of her posterior distributions, p_t^i, is changing with experience. With experience, however, the posterior distributions of a true Bayesian would become less dispersed, and her choices would converge. In the case of long run frequencies, then, we may expect find consumers' actual behavior fall somewhere in between the pictures offered by the simple and Bayesian models.

Indeed, empirical findings, such as those of Camerer [10], do imply just such a difference between short and long-term behavior. While people may act myopically in the short run, their long-run patterns of choice appear to become more fully Bayesian. Thus, our model of consumer behavior may not be wholly adequate for the long run; it may be profitably modified or extended.

More broadly, one may note that the model presented in this paper considers the myopic behavior associated with categorical thinking but does not address other well-documented biases associated with individual decision-making under uncertainty. (For example, see Payne *et al.* [40] or Meyer and Hutchinson [37].) Nevertheless, we believe that the learning associated with repeated trials is a fundamental determinant of customer loyalty in the context of quality uncertainty. In this sense the model of a simple consumer provides an advance over more traditional models in which loyalty is exogenously defined. How much of an advance may be determined through empirical validation.

Acknowledgments

Thanks to David Croson, Rachel Croson, Pat Harker, Teck-Hua Ho, Eric Johnson, Ziv Katalan, Howard Kunreuther, Robert Meyer, Garrett van Ryzin, and Yu-Sheng Zheng, as well as to seminar participants at Northwestern, for valuable comments and references.

References

[1] Allen, F., and G. R. Faulhaber (1988), "Optimism Invites Deception," *The Quarterly Journal of Economics* **103**, 397 – 407.

[2] Allen, F., and G. R. Faulhaber (1991), "Quality Control in the Service Firm and Consumer Learning," in *Service Quality* (S. W. Brown, E. Gummelson, B. Edvardsson, and B. Gustavsson, eds.), Lexington, Massachusetts: Lexington Books, 289 – 302.

[3] Anderson, S. P., A. de Palma, and J-F. Thisse (1992), *Discrete Choice Theory of Product Differentiation*, Cambridge: MIT Press.

[4] Balcer, Y. (1980), "Partially Controlled Demand and Inventory Control: an additive model," *Naval Research Logistics Quarterly* **27**, 273 – 288.

[5] Balcer, Y. (1983), "Optimal Advertising and Inventory Control of Perishable Goods," *Naval Research Logistics Quarterly* **30**, 609 – 625.

[6] Banks, J. S., and R. K. Sundaram (1992), "A Class of Bandit Problems Yielding Myopic Optimal Strategies," *Journal of Applied Probability*, **29**, 625 – 632.

[7] Banks, J. S., M. Olson, and D. Porter (1997), "An Experimental Analysis of the Bandit Problem," *Economic Theory* **10**, 55 – 77.

[8] Bergemann, D. and J. Välimäki (1996), "Learning and Strategic Pricing," *Econometrica* **64**, 1125 – 1149.

[9] Boulding, W., A. Kalra, R. Staelin, and V. A. Zeithaml (1993), "A Dynamic Process Model of Service Quality: from expectations to behavioral intentions," *Journal of Marketing Research* **30**, 7 – 27.

[10] Camerer, C. F. (1987), "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence," *American Economic Review* **77**, 981 – 997.

[11] El-Gamal, M. A., and D. Grether (1995), "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association* **90**, 1137 – 1145.

[12] Erdem, T., and M. P. Keane, "Decision-making Under Uncertainty: Capturing Dynamic Brand Choice in Turbulent Consumer Goods Markets," *Marketing Science* **15**, 1 – 20.

[13] Feller, W. (1968), *An Introduction to Probability Theory and Its Applications*, Vol. 1, 3^{rd} ed., New York: John Wiley & Sons.

[14] Fergani, Y. (1976), *A Market Oriented Stochastic Inventory Model*, unpublished Ph.D. dissertation, Stanford University.

[15] Gans, N. (1999), "Customer Loyalty and Supplier Strategies for Quality Competition," Working Paper, OPIM Department, The Wharton School, University of Pennsylvania.

[16] Gegerenzer, G. and D. J. Murray (1987), *Cognition as Intuitive Statistics*, Hillsdale, New York: Lawrence Erlbaum Associates.

[17] Gilboa, I., and A. Pazgal (1995), "History Dependent Brand Switching: theory and evidence," Working Paper, KGSM-MEDS, Northwestern University.

[18] Gilboa, I. and D. Schmeidler (1997), "Cumulative Utility Consumer Theory," *International Economic Review* **38**, 737 – 761.

[19] Gittins, J. C. (1979), "Bandit Processes and Dynamic Allocation Indices," *Journal of the Royal Statistical Society B* **41**, 148 – 177.

[20] Gittins, J. C., and D. M. Jones (1974), "A Dynamic Allocation Index for the Sequential Design of Experiments," In *Progress in Statistics*, ed. J. Gani et al., Amsterdam: North Holland, 241 – 266.

[21] Grether, D. (1980), "Bayes Rule as a Descriptive Model: the representativeness heuristic," *The Quarterly Journal of Economics* **95**, 537 – 557.

[22] Grether, D. (1992), "Testing Bayes Rule and the Representativeness Heuristic: some experimental evidence," *Journal of Economic Behavior and Organization* **17**, 31 – 57.

[23] Hall, J. M., and E. Porteus (1998), "Dynamic Customer Service Competition," INFORMS Fall 1998 Conference, Seattle.

[24] Henderson, P. W., and R. A. Peterson (1992), "Mental Accounting and Categorization," *Organizational Behavior and Human Decision Processes* **51**, 92 – 117.

[25] Hey, J. D., and C. J. McKenna (1981), "Customer Search with Uncertain Product Quality," *Journal of Political Economy* **89**, 54 – 66.

[26] Ho, T-H., and Y-S. Zheng (1996), "Setting Customer Expectation in Service Delivery," Working Paper, Marketing Department, The Wharton School, University of Pennsylvania.

[27] Horowitz, A .D. (1973), *Experimental Study of the Two-Armed Bandit Problem*, unpublished Ph.D. dissertation, University of North Carolina at Chapel Hill,

[28] Kahneman, D., and A. Tversky, (1973), "On the Psychology of Prediction," *Psychological Review* **80**, 237 – 251.

[29] Kalai, E., M. I. Kamien, and M. Rubinovitch (1992), "Optimal Service Speeds in a Competitive Environment," *Management Science* **38**, 1154 – 1163.

[30] Keener, R. (1986), "Multi-Armed Bandits with Simple Arms," *Advances in Applied Mathematics* **7**, 199 – 204.

[31] Lai, T. L. and F. Chang (1987), "Optimal Stopping and Dynamic Allocation," *Advances in Applied Probability* **19**, 829 – 853.

[32] Lederer, P. J., and L. Li (1997), "Pricing, Production, Scheduling, and Delivery-Time Competition," *Operations Research* **45**, 407 – 420.

[33] Li, L., and Y. S. Lee (1994), "Pricing and Delivery-Time Performance in a Competitive Environment," *Management Science* **40**, 633 – 646.

[34] Lorden, G. (1970), "On Excess over the Boundary," *The Annals of Mathematical Statistics* **41**, 520 – 527.

[35] Marketing Science Institute (1998), *Research Priorities: a guide to MSI research programs and procedures*, Cambridge, Massachusetts: Marketing Science Institute.

[36] Mendelson, H. and S. Whang (1990), "Optimal Incentive-Compatible Priority Pricing for the M/M/1 Queue," *Operations Research* **38**, 870 – 883.

[37] Meyer, R. J., and J. W. Hutchinson (1994), "Intuitive Dynamic Decision Making: an analysis of normative theory as a descriptive model," Working Paper 94-011, Wharton Marketing Department, University of Pennsylvania.

[38] Meyer, R. J., and Y. Shi (1995), "Sequential Choice Under Ambiguity: intuitive solutions to the armed-bandit problem," *Management Science* **41**, 817 – 834.

[39] Nahmias, S. (1997), *Production and Operations Analysis*, 3^{rd} ed., Chicago: Irwin.

[40] Payne, J. W., J. R. Bettman, and E. J. Johnson (1993), *The Adaptive Decision Maker*, Cambridge: Cambridge University Press.

[41] Roberts, J. H., and G. L. Urban (1988), "Modeling Multiattribute Utility, Risk, and Belief Dynamics for New Consumer Durable Brand Choice," *Management Science* **34**, 167 – 185.

[42] Robinson, L. W. (1990), "Appropriate Inventory Policies When Service Affects Future Demands," Working Paper 88-08, Johnson Graduate School of Management, Cornell University.

[43] Schwartz, B. L. (1965), *Inventory Models in which Stockouts Influence Subsequent Demand*, unpublished Ph.D. dissertation, Stanford University.

[44] Schwartz, B. L. (1966), "A New Approach to Stockout Penalties," *Management Science* **12**, B-538 – B-544.

[45] Schwartz, B. L. (1970), "Optimal Inventory Policies in Perturbed Demand Models," *Management Science* **16**, B-509 – B-518.

[46] Siegmund, D. (1979), "Corrected Diffusion Approximations in Certain Random Walk Problems," *Advances in Applied Probability* **11**, 701 – 719.

[47] Siegmund, D. (1985), *Sequential Analysis: tests and confidence intervals*, New York, Springer Verlag.

[48] Silver, E. A., D. F. Pyke, and R. Peterson (1998), *Inventory Management and Production Planning and Scheduling*, 3^{rd} ed., New York: John Wiley & Sons.

[49] Shaked, M., and J. G. Shanthikumar (1994), *Stochastic Orders and Their Applications*, Boston: Academic Press.

[50] Smallwood, D. E., and J. Conlisk (1979), "Product Quality in Markets where Consumers are Imperfectly Informed," *Quarterly Journal of Economics* **XCIII**, 1 – 23.

[51] Stidham, S. Jr. (1992), "Pricing and Capacity for a Service Facility: stability and multiple local optima," *Management Science* **38**, 1121 – 1139.

[52] Thoman, L. (1994), "Repeat Purchases Under Quality Uncertainty," *Economics Letters* **46**, 33 – 40.

[53] Tversky, A., and D. Kahneman, (1974), "Judgment under Uncertainty: heuristics and biases," *Science* **185**, 1124 – 1131.

[54] Wald, A. (1947), *Sequential Analysis*, New York: John Wiley & Sons.