Unstable Consumer Learning Models:

Structural Models and Experimental Investigation

by

Mitchell James Lovett

Department of Business Administration
Duke University

Date:_______________________
Approved:

____
William Boulding, Co-Supervisor

____
Richard Staelin, Co-Supervisor

____
Ron Shachar

____
Mike West

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Business Administration in the Graduate School
of Duke University

2008

ABSTRACT

Unstable Consumer Learning Models:

Structural Models and Experimental Investigation

by

Mitchell James Lovett

Department of Business Administration
Duke University

Date:_______________________
Approved:

____
William Boulding, Co-Supervisor

____
Richard Staelin, Co-Supervisor

____
Ron Shachar

____
Mike West

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Business Administration in the Graduate School
of Duke University

2008

Copyright by
Mitchell James Lovett
2008

Abstract

This dissertation explores how consumers learn from repeated experiences with a product offering. It develops a new Bayesian consumer learning model, the unstable learning model. This model expands on existing models that explore learning when quality is stable, by considering when quality is changing. Further, the dissertation examines situations in which consumers may act as if quality is changing when it is stable or vice versa. This examination proceeds in two essays.

The first essay uses two experiments to examine how consumers learn when product quality is stable or changing. By collecting repeated measures of expectation data and experiences, more information enables estimation to discriminate between stable and unstable learning. The key conclusions are that (1) most consumers act as if quality is unstable, even when it is stable, and (2) consumers respond to the environment they face, adjusting their learning in the correct direction. These conclusions have important implications for the formation and value of brand equity.

Based on the conclusions of this first essay, the second essay develops a choice model of consumer learning when consumers believe quality is changing, even though it is not. A Monte Carlo experiment tests the efficacy of this model versus the standard model. The key conclusion is that both models perform similarly well when the model assumptions match the way consumers actually learn, but with a mismatch the existing

model is biased, while the new model continues to perform well. These biases could lead to suboptimal branding decisions.

Dedication

This work is dedicated to my wife, Cali, without whose love and support it would not have been possible, and to my parents and the memory of my Aunt Frances, who taught me to value education.

Contents

List of Tables

List of Figures

Acknowledgements

To begin I would like to thank my esteemed committee --Bill Boulding, Rick Staelin, Ron Shachar and Mike West--without whose help and guidance completing this work would not have been possible. In particular, Rick and Bill have patiently advised me throughout the process, providing guidance not only on research and teaching, but also on life as an academic. They also graciously provided financial support for the data collection and computation and always made sure I had the necessary resources.

Numerous other faculty members have given me advice and help during the PhD process. Before I mention a few specific individuals, I would like to acknowledge the helpful comments and support the marketing faculty has given me, not to mention my course work and training. I thank Christine Moorman for pushing me to be a broader, more open researcher, and for freeing me to do things my way, Carl Mela for many helpful conversations that made my dissertation and general approach to research better, and Jim Bettman for patiently handling administrative details of completing the dissertation. I also want to thank Abhijit Guha and other students in the PhD program who helped with data collection and helped me think through problems.

Finally, I want to thank my wife for putting up with me during this process. While I had this one "baby" of a dissertation, she gave us two babies. Twice as productive as I am, and I needed every bit of the love and support she gave me.

1. Introduction

In many product settings, consumers have repeat experiences with a product offering and use these experiences to infer the quality of the product. Past research on how consumers learn from repeat product experience and update their beliefs has focused on situations where (1) the product offering has some true (mean) product quality that remains constant over time so that consumers' experiences with this product offering vary around this true quality level (henceforth, we refer to this as a **stable** attribute of quality) and (2) consumers know that the true (mean) quality is constant (e.g., Boulding, Kalra, and Staelin 1999; Erdem and Keane 1996). This dissertation expands on these existing models by considering models that allow both for true quality to change from one experience to the next (henceforth, we refer to this as an **unstable** attribute or quality) and for consumers to act as if it is changing, even when it is not, presenting potentially irrational behavior.

We follow the tradition of Bayesian learning models in marketing (Boulding et al. 1999; Rust et al. 1999; Erdem and Keane 1996) to understand approximately how consumers learn. The first essay addresses the question of whether consumers behave as if the true attribute is changing, when in fact it may or may not be. To answer this question, we collect experimental data that allows direct estimation of the consumer updating process from observed expectations and experiences. Through this estimation process we get a much more nuanced perspective of consumer learning than in previous

Bayesian learning studies. One key result from this first essay is that most consumers act as if the underlying (unobserved) quality level is changing, even when it is stable. This finding has potential implications for a large and growing literature in choice models that incorporate Bayesian learning.

The second essay explores these implications and presents choice models that allow consumers to behave as if the world is changing when it is not. Further, using simulated data, we explore the biases that arise when researchers use the incorrect model to estimate consumer learning. We find that if the researcher assumes consumers believe the world is stable when in fact consumers do not, estimates are heavily biased in ways that could lead to bad marketing investments. We then show that our extension of the standard model performs well regardless of the actual consumer behavior.

To motivate this dissertation, we both relate our work to the existing research in Bayesian learning models and present an application of our ideas to branding. Our work is closely related to the Bayesian learning literature contained in both marketing and economics. This literature includes early contributions by Meyer (1981), Meyer and Sathi (1985), and Roberts and Urban (1988) using experimental and clinical data, and Eckstein, Horsky, and Raban (1988), and Horsky and Raban (1988) using diary data. Later models developed in this same tradition largely apply the Bayesian learning formulation to scanner data or scanner data augmented by additional sources of information. In these models, choices are typically the only observed quantity; hence, the consumers' experiences, expectations, and uncertainties are all unobserved. All of these unobserved quantities are assumed to follow parametric forms dictated by the particular assumptions of the Bayesian learning model. Researchers have then demonstrated improvements over existing state-dependence formulations in terms of prediction and fit (e.g., Erdem and Keane 1996; Mehta, Rajiv, and Srinivasan 2004). These papers have, for example, incorporated forward-looking consumers (e.g., Erdem and Keane 1996), advertising and experience effects (e.g., Ackerberg 2003), learning about variances (Iyengar, Ansari, and Gupta 2007), and transfers of information (e.g., Szymanowski and Gijsbrechts 2008). Moreover, this literature continues to spawn a large number of working papers (e.g., Shin et al. 2007; Chan et al. 2007; Chintagunta et al. 2007; Zhao et al. 2007; Osborne 2006; Narayanan and Manchanda 2007; Szymanowski and Gijsbrechts 2008).

All of these papers use data where expectations and experiences are unobserved. In marketing, a number of papers have used very similar assumptions, but focused on other types of data. These papers cover the topics of service quality (Boulding et al. 1999; Rust et al. 1999) and wait times (Kumar et al. 1997). In all of these papers, the measures of interest were expectations rather than, or in addition to, choices. Such data can allow more refined models of consumer learning to gain additional insight.

We build upon this literature and extend the core model. In particular, we develop an approach that allows non-optimal behavior, a growing area in this literature (e.g., Boulding et al. 1999; Mehta et al. 2004; Camacho, Donkers, Stremersch 2008). Nearly all of these models make a key dual assumption that we relax--that product quality is stable once a product is introduced and that consumers believe this to be the case. We do so by allowing consumers to behave as if quality is unstable, a) when it is in fact stable and b) when the quality level is actually unstable. We apply these settings to both expectations and experience data in the first essay and to choice data in the second essay.

Because this dissertation focuses on consumer learning when quality could be changing, our second motivation arises from the psychology of identification of change. A body of research examines the ability of individuals to detect relationships from repeated experiences with stimuli (cues) (e.g., Brehmer 1980). For example, this work suggests individuals rely too much on deterministic versus probabilistic rules for learning relationships (Brehmer 1980), that how individuals learn is influenced by how noisy the signals are (Muchinsky and Dudycha 1975), and that individuals are slow to adjust to new informative cues (Edgell 1983).

This psychology of identification of change.work has spawned two approaches to studying how individuals respond to change in regime that are relevant to this thesis. The first set of studies examines changes in individuals' predicted relationships in response to changes in the true underlying relationship. The findings in this work include that individuals are slower to learn after a change than prior to the change and that individuals respond quite differently to change (Peterson, Hammond, and Summers 1965), that changes in relationship and signal invoke slower responses than either individually (Summers 1969), that changes from high to low variance environments results in worse relative performance than low to high, and that an increase in error variance is easier to detect than a decrease in error variance (Lindeberg and Brehmer 1976). The second set of studies focuses on predicting the timing of a regime change in the underlying process (Brown and Bane 1975; Rapoport 1979; Barry and Pitz 1979; Massey and Wu 2005). This work provides evidence that individuals consistently overweight the indications of change and evaluate the particular case at hand without an appreciation for the class in which it belongs.

Together, these two sets of studies suggest for this dissertation the important role of **perceptions** of change whether or not change is actually present. While these perceptions of change are likely to correspond to some degree with actual changes, the correspondence may be of less interest, per se, than the influence of these perceptions of change on the way consumers learn and try products. For example, some evidence from choice experiments suggests that even after repeated trials subjects continue to try suboptimal options, even when the outcome of that option is fixed (Erev and Haruvy 2008) . We suggest one explanation may be a perception that the value of that option may be changing. In the context of Bayesian learning models, we will focus on how consumers learn when they perceive the true attribute as changing, regardless of whether or not it is changing. In this way, we open the avenue for investigating how, through repeat experiences, consumer perceptions of change may influence choices and expectations.

The third motivation arises from an incongruity between the theory of branding and popular press accounts of consumer loyalty behaviors. Marketers have long sought the benefits of brand loyalty and being able to keep customers coming back even if they have a less than perfect experience with the product or service. This willingness to repeat purchase is often referred to as brand loyalty. The foundation of brand loyalty is a brand's equity and requires consumers to have readily accessible, positive and unique beliefs about the product offering (Keller 1993). However, popular press accounts suggest that consumers do not exhibit brand loyalty, but are instead fickle and promiscuous, moving from one brand to another (Surowiecki 2004). We focus on one explanation that could produce both brand equity and fickle behavior—the way consumers learn from experiences and update their beliefs about product attributes. By considering how consumers learn about the underlying quality level of a product when they believe the true (mean) product quality is changing, this dissertation presents a way to reconcile both fickle and brand loyal behavior based on learning behaviors.

The first essay focuses on understanding how consumers learn and how marketers may be influenced by this behavior. Most existing Bayesian learning models assume underlying quality is stable.[1] We refer to models with this assumption as stable learning models. In this essay, we extend these stable learning models to incorporate changing quality and develop a rational Bayesian belief-updating model.

We show that in the unstable quality setting—one in which the true quality is changing—even after multiple experiences, consumers should continue to place substantial weight on the most recent experience when updating their beliefs. In contrast if the quality level is stable they should place substantial weight on the prior belief. Further, we demonstrate how this unstable learning model subsumes the stable learning model as a special case and can generate learning patterns that have the rational consumer place more weight on his or her prior belief over time as well as learning patterns that continue to react strongly to each new experience. Our unstable learning model can explain why a consumer with a string of good experiences followed by one bad experience can switch to another offering and still behave normatively. At the same time, our model can generate beliefs and behavior consistent with brand loyalty.

We also explore in this first essay the possibility that consumers apply the wrong learning model to a given situation. For example, we look at the situation where a

---

[1] Erdem et al. (2005) allow for changing price expectations and levels with constant relative quality levels. Another set of research allows quality to change, but consumers know true quality (Ackura et al. 2004; Sriram et al. 2006). These latter two papers model dynamic quality attributes, but treat quality perceptions as certain and changing rather than develop optimal updating rules for consumers that acknowledge consumer uncertainty.

consumer acts as if he or she is using an unstable learning model even though the true underlying quality is stable over time as well as the situation where the consumer acts as if he or she is using a stable learning model when the true quality level is changing over time. We show that if consumers act as if the quality level is stable when in fact it is changing, they will exhibit excessively brand loyal behavior, whereas if they act as if the quality level is changing when it is in fact stable, they will exhibit promiscuous behavior. Thus, we envision a world where consumers update their quality beliefs accurately or inaccurately to produce accurate, promiscuous, and excessively brand loyal behavior. Further, we explain how the inaccurate behaviors will also impact predictive uncertainty, predictive error, and choice.

Our investigation of inaccurate learning models leads us to consider whether such "inaccurate" behavior could also be a result of a rational response. We illustrate this possibility with a model that incorporates consumer uncertainty about whether quality is changing. In this model, the consumer uses experiences to help infer whether quality is changing or not. The resulting way in which consumers learn is a mixture between the stable and unstable learning models. By analyzing this illustrative model, we develop several propositions, including that it can be rational for consumers to act as if quality is changing, even when it is stable.

We next conduct two experiments employing two different product-offering settings to test the implications of our model. In both settings, participants are exposed to a series of service encounters. After each experience, they are asked what they think will happen during their next experience and how uncertain they are about what will happen next. These subjective beliefs are used to understand the learning patterns that consumers exhibit over time.

To analyze these patterns, we apply the stable and unstable models of learning to consumer behavior. More specifically, we develop a Bayesian dynamic mixture model that allows consumers to act according to the stable or unstable learning models and estimates their primitive beliefs. This approach allows us to both test which model fits consumer behavior best and identify when consumers behave inaccurately. We identify considerable heterogeneity in the learning patterns consumers use in response to the same product quality experiences. Capturing these heterogeneous responses allows us to generate insight beyond what has been studied in the past.

In the first experiment we find that even when the consumer faces a stable product quality environment, the unstable learning model captures consumers' actual learning patterns better than the stable learning model. That is, consumers tend to act too promiscuously when encountering stable quality experiences. This tendency to use an unstable learning model was even more pronounced when consumers experienced changing quality levels. Thus, not only are consumers less likely to exhibit brand loyal behavior when a firm's mean quality level changes over time, consumers are sensitive to their environment and modify their learning model to reflect this environment. Still, some consumers' model parameters do not accurately represent the underlying learning environment. This implies that some consumers update their quality beliefs too much (too little) based on a single experience. We also find that these consumers who are less sensitive to their environment have more predictive error and predictive uncertainty.

In the second experiment we replicate the essential findings from the first study in a new, much more complicated setting. Participants in this study experienced two different product options over time and chose between them at five different points of time. We show that when making choices, consumers who exhibit unstable learning patterns weigh the most recent experiences more relative to their prior belief compared to consumers who exhibit stable learning patterns. Hence, the observed heterogeneity in learning patterns directly influences choice.

Overall, our research supports a much more complicated picture of the quality learning behavior of consumers. Thus, we find consumers exhibit a variety of accurate and inaccurate learning behaviors. We also document that modeling consumer heterogeneity at a segment level is valuable in capturing how consumers learn, even given the same experiences. By dealing with heterogeneity and introducing unstable learning, we shed light on not only how consumers behave in an unstable world, but also how consumers behave in traditional stable settings. Specifically, some consumers behave accurately in both stable and unstable product settings, while some consumers are excessively loyal in unstable settings and still others are promiscuous in stable settings.

In the second essay, we explore the implications of these findings for models that rely on choice data and assume consumers learn about the underlying quality level of each product in the product category. Such choice models incorporate Bayesian learning processes, a learning process that has a long tradition in marketing. Most of these models assume that the firm's quality level is, in fact, stable over time and that the consumers know it is stable and use an optimal updating policy (i.e., using what we term a stable learning model). However, based on the evidence provided in essay one that consumer learning patterns often are better approximated by an unstable learning model, even when true quality is stable, in the second essay we examine the potential influence of this inaccurate consumer behavior on the estimation of choice models as well as investigate the estimation of an unstable learning model.

We do this as follows. First, we expand the analysis presented in essay one by developing estimation models based primarily on choice data rather than expectations data. Specifically, we develop structural estimation models for when consumers act as if the products are stable or unstable, when products are actually stable. We then develop a MCMC estimation procedure for the structural estimation models.

Second, we demonstrate how these estimation models can be derived from commonly used, statistical models in time-series (West and Harrison 1997).[2] Specifically, we map the Bayesian learning models onto the dynamic linear model (DLM), presenting the observation, transition, and variance matrices. This provides a clear depiction of how the Bayesian learning model shares a common framework with the DLM, but also demonstrates the specific constraints that produce the structural parameters. Thus, a comparison between the two previously distinct models becomes straightforward.

Third, we use simulated data to explore the performance of these samplers in small samples. We discuss the samplers' convergence behavior and ability to recover the true parameters. We find that both structural estimation models do well in recovering the true parameters when the data is generated according to the assumptions of the model. Further, when the estimation model assumes consumers use an unstable learning model when in fact they use a stable learning model, model performance is similar to when the model matches the actual environment. However, when the estimation model assumes consumers use a stable learning model when in fact they use an unstable learning model, important biases are introduced. This investigation indicates the nature of those biases as well as the value of applying the unstable learning model even to cases when the underlying attribute is stable.

---

[2] We thank Carl Mela for several thoughtful conversations and intellectual prodding that helped clarify and refine this mapping.

2. Essay 1: The World Is Changing Isn't It? Implications for Consumer Learning, Accuracy, and Choice

2.1 Introduction

The world is changing at a rapid pace, perhaps increasingly so. Each month, week, and day of the year, firms introduce products and services that usher in new standards of quality within markets. For existing products, firms regularly introduce new versions and processes, all of which change the quality of these products. How should and do consumers learn from repeated experiences when product quality is changing?[1]

We illustrate this kind of learning with a simple example. Imagine a consumer has used a particular car lube service three times in the past and the service took 26, 20 and 29 minutes to complete. Based on these past experiences the individual develops a belief that the underlying (unobserved) quality level of the service process has a 25 minute average service time. Now, imagine that on the next visit the service time is 45 minutes. How long does the consumer expect the next visit will last? Based on the new information, how does the consumer adjust the 25 minute belief about the average service time?

[1]We use product to refer to products, services, or a combination of the two.

A sizable literature employs a model of consumer learning that answers these questions when the product quality is stable, in the sense that product experiences vary around a constant mean (true) product quality (e.g., Roberts and Urban 1988; Erdem and Keane 1996; Boulding, et al. 1999; Rust, et al. 1999; Erdem 1998; Crawford and Shum 2003; Ackerberg 2003; Israel 2005). In these models, consumers are uncertain about a stable mean (true) product quality and have experiences that differ from one encounter to the next. Consumers update in a Bayesian fashion their uncertain beliefs about this stable product quality. As a result, consumers learn over repeated experiences, and their beliefs about quality converge to the fixed, or stable, true quality level. Because both these models assume consumers know that product quality is stable, we term models for how consumers learn based on this assumption **stable learning models**.

However, if the world is changing, consumers using a **stable learning model** will have their beliefs converge towards something that does not exist. Recognizing and capturing the dynamic nature of the environment is the first differentiating point of our research.[2] We develop a consumer learning model that we term the **unstable learning model** in which consumers take into account the fact that they believe the mean (true) quality level of the product changes over time. More technically, consumers acknowledge that the variability that they notice in product experiences is due not only

---

[2] See Erdem, et. al., 2005 for a model that allows price expectations to change while holding fixed relative quality. Further, see Akcura, et. al. 2004 for a model in which the quality is changing, but not due to learning about uncertainty. In their model, consumers know true quality.

to the inherent variability around the true underlying quality level, but also because the true quality level is changing over time. However, since a consumer never observes the true underlying quality level, it is possible for this consumer to believe the true quality level is changing (and thus use an unstable learning model), even when the true quality is stable. Thus in this essay we do not assume that the consumer uses the correct learning model. Instead we investigate the implications that follow from having the consumer's learning model match (not match) the actual state of the world.

Taking such an approach allows us to a) demonstrate how a consumer **should** update his or her belief about the true quality level after having a product experience based on whether the individual believes the world (i.e., the underlying quality process that generates wait times) has changed or not and b) determine the impact of consumers using the correct (incorrect) learning model on the choice behavior of these consumers. In this way we are able to provide a reason why we observe some consumers exhibiting excessively brand loyal behavior (i.e. staying with a brand even after experiencing a number of "bad" outcomes) while other consumers exhibit promiscuous behavior (i.e., switch brands after experiencing just one "bad" experience).

This essay adds to a growing body of research that attempts to build into the Bayesian learning framework non-normative behaviors (e.g., Mehta et al. 2004; Boulding, et al. 1999). To this end our focus is on identifying to what extent consumer learning behaviors can be approximated by a consumer learning model that assumes the world is changing versus one that assumes the world is stable. If consumers act as if the world is changing, when in fact it is stable (or vice versa) these beliefs are inaccurate and are inconsistent with normative behavior, at least as assumed in the prior literature.

Beyond the unstable-stable differentiation, this research is different from the extant learning models that all are estimated using choice data. Thus, instead of inferring learning from choices, we use expectation data to directly estimate the consumer learning process. Such an approach allows us to gain a much finer-grained view of learning. In particular, we allow consumers to make errors in their updating in addition to allowing consumers to use different learning models that imply different assumptions about whether quality is changing or not. Further, we are able to identify heterogeneity in the learning model as well as beliefs (see Shin, Misra, and Horsky (2007) for an alternative approach to learning process heterogeneity). Thus, this work should be viewed as exploring the foundations of the latent learning process and can support extensions to the choice literature. However, it is not in direct competition with these models since the phenomena to be explained differs, i.e our attention is on expectations while the extant literature centers on choices.

The essay proceeds as follows. In section 2.2, we first extend existing consumer learning models to situations in which the true, latent, quality is changing from one experience to the next. We then generalize this model to situations in which consumers are uncertain about whether or not the underlying quality environment is changing.

From this model, we generate a number of testable propositions. In sections 3 and 4, we test our model using two experiments in which we manipulate the type of underlying quality environment and thus the actual product experiences. Consumers are asked to give predictions about what they believe will happen on their next product experience. Using the experience and expectation data, we estimate different structural learning models, and identify which model best fits the measured beliefs. To foreshadow our results, we find that consumers act in a manner consistent with our learning model. In addition, even though we find that consumers are sensitive to the environment they encounter, a meaningful portion of consumers act as if the environment is unstable when it is stable, and vice versa, i.e., they choose the wrong, or, more precisely, less accurate learning model. In the final section, we synthesize the results, discuss managerial implications of our findings, and identify areas for future research.

2.2 Model development

We start with the observation that for many product offerings, the consumer's experiences can vary over time. We assume this variation can be attributed to two sources. First, the underlying quality level can change over time resulting in different experiences with each new product encounter. We call this unstable variation because it leads the true quality to be unstable (i.e., changing). Second, even if the underlying quality level is stable, the product experiences are fallible representations of this level. We call this latter variation stable variation, since it is variance around a stable mean.

For services, this stable variability in product experiences could arise due to the particular service employee(s) encountered and their behaviors, both of which can differ across encounters. For physical products purchased repeatedly, stable variability could be caused by imperfectly reliable quality processes. For example, a bag of potato chips could differ from experience to experience because the potatoes change, the cooking oil heat differs, salt sticks more or less to the chips, and so on. Finally, even for consumer durables (e.g., an automobile), the experiences can differ over time due to factors such as usage context effects.

We also assume that consumers never directly observe the true quality level of the product offering. Instead, they infer quality via product experiences.[3] Using these assumptions, we present a normative learning model, where the consumer knows the two sources of variation.

We then investigate what happens if consumers use an inaccurate learning model given the true state of the world (i.e., stable or unstable). These inaccuracies can lead to learning behaviors that update beliefs too much or too little based on each new experience. Moreover, we argue that uncertainty about whether the world is unstable or stable could lead consumers to rationally use an "inaccurate" learning model that mixes between the stable and unstable learning models. Here, we put inaccurate in quotes

---

[3] Although in our empirical context we focus on learning from repeated experiences with a product, our model is agnostic about the source of information, e.g., consumers might learn from sources other than their own experiences.

because in this rational model, the consumer could do no better given what they know. However, if they knew whether the world was stable or unstable, they would be acting incorrectly. We illustrate these ideas with a simple selection model of how consumers might learn whether to use a stable or unstable learning model.

2.2.1 Firm quality delivery process

We assume the firm's offering has a true, underlying quality, Q_t, at time t and this quality level affects each consumer's experience with the offering. We follow the lead of Boulding et al. (1999) and assume that the delivered quality experiences for consumer i, DE_{it}, can differ from this true quality. Thus, each delivered experience can be viewed as a fallible measure of the true quality of the firm. Formally, the delivered experience is a realization from a stochastic quality delivery process. Following the literature, we assume that the error associated with this delivery process is normally distributed with mean Q_t and variance S, i.e.,

$$DE_{it} = Q_t + s_{it}\text{, where } s_{it} \sim N[0, S] \quad (1)$$

We next allow for the underlying quality process to vary over time. We acknowledge that there are numerous alternatives for how this quality level might change. We elect to use an alternative that imposes little structure. Specifically, we assume the underlying quality process is equally likely to increase or decrease at any point in time. Although firms presumably intend for changes to yield quality increases, there are numerous reasons why the quality of a product might decrease as well as increase. First, quality might decrease when firms try to cut costs. Second, the quality process can degrade due to the loss of key personnel or access to suppliers. Third, since quality is often judged relative to competition (Boulding et al. 1999), increases in the competitive quality level can cause a decrease in the perceptions of the target firm's quality. Fourth, since the outcomes of a quality process are ultimately determined by subjective consumer perceptions and a complex delivery system, managers may make mistakes when attempting to improve the underlying quality process. We quantify these ideas by assuming the change in quality, d_t, is distributed normally with variance D, i.e.

$$Q_t = Q_{t-1} + d_t, \text{ where } d_t \sim N[0, D] \tag{2}$$

2.2.2 Consumer beliefs about quality

The consumer's objective is to form accurate beliefs (i.e., learn) about the underlying quality level, Q_t. This learning is summarized by a belief structure which we represent in terms of a probability distribution. More specifically, we assume consumer i's belief about the current quality level after an experience at time t is $b_{it}(Q_t)$, where $b(\bullet)$ represents the distribution of values that the consumer believes are possible. Following the literature, we assume these beliefs are normally distributed with mean, $WE_{it,}$ and variance (uncertainty), U_{it}. Similarly, we define the belief $b_{it}(DE_{it+1})$ to be the belief about the next product experience. This is the belief that consumers use when deciding on their next product choice. It is also the belief that consumers use to determine if their learning model is accurate, since it provides a mechanism for comparing a stated belief with an actual occurrence. This belief is a predictive distribution that has the same mean, WE_{it},[4] but greater variance (or predictive uncertainty), P_{it}, where $P_{it} = U_{it}+S+D$.

2.2.3 Consumer belief updating about quality

Our interest centers on developing what we call the **unstable learning model**, a model of normative belief updating when consumers know there is both stable (S) and dynamic (D) variance. West and Harrison (1997), using standard Bayesian statistical analysis, present a number of statistical models, one of which captures the quality process and delivery experience modeled in Equations 1 and 2. Their models can be used to show that a consumer would (normatively) update his or her quality beliefs as follows:

$$WE_{it} = \alpha_{it} * DE_{it} + (1- \alpha_{it}) * WE_{it-1} \quad (3)$$
$$U_{it} = \alpha_{it} * S \quad (4)$$
$$\text{where } \alpha_{it} = [(U_{it-1} + D)/(U_{it-1} + S + D)] \quad (5)$$

Equations 3-5 have a number of interesting properties. First, they fully describe the consumer's updating process since $b_{it}(Q_t)$ is uniquely defined once the mean, WE_{it}, and variance, U_{it}, are specified. Second, these equations imply that normative behavior restricts how much consumers should learn from each new experience. More specifically, the weights are restricted between zero and one, i.e., $0 < \alpha_{it} < 1$, and sum to one. Third, assuming U_{i0} is greater than S, the person's uncertainty (U_{it}) about the

---

[4] As shown by Boulding et al. (1999), WE_{it}, is both the consumer's best guess of what will happen during the next product experience and the current belief about the quality level. These researchers refer to this best guess as the consumer's Will Expectation, hence, the notation WE_{it} .

current quality level decreases with each new experience and reduces to less than the stable variance after the initial experience. Fourth, equations 3 and 5 are similar in form to adaptive exponential smoothing and are a special case of the Kalman filter. Finally, note that the prior models of Boulding et al. (1999) and Rust et al. (1999) can be characterized by these same equations, where the dynamic variance, D, is set to zero. This special case, which we term the **stable learning model**, has particular significance because it underlies many Bayesian consumer learning models in marketing (e.g., Erdem and Keane 1996).

Equations 4 and 5 are also useful in specifying how a person's weights (and thus learning patterns) should change over time. Solving for α_{it} in terms of α_{it-1}, we note the following:

$$\alpha_{it} = (\alpha_{it-1} + D/S) / (\alpha_{it-1} + D/S + 1) \tag{6}$$

This equation implies that α_{it} should decrease with each new experience and specifies the rate of this decrease. It also highlights the need for consumers to accurately partition the observed variation in the delivered experiences between the stable variance, S, and the dynamic variance, D. If the consumer acts according to the **stable learning model**, i.e., assumes D=0, the learning model implies that α_{it} quickly approaches zero (Rust et al. 1999). However, if D/S > 0, the normative implications are that α_{it} will not asymptote to zero. In fact, it is possible for α_{it} to stay very close to one if the dynamic variance is much larger than the stable variance. Finally, as shown in

Figure 1, given the same starting value α_{i0}, the value of α_{it} in the **unstable learning model** can never be less than the value of α_{it} in the **stable learning model**).

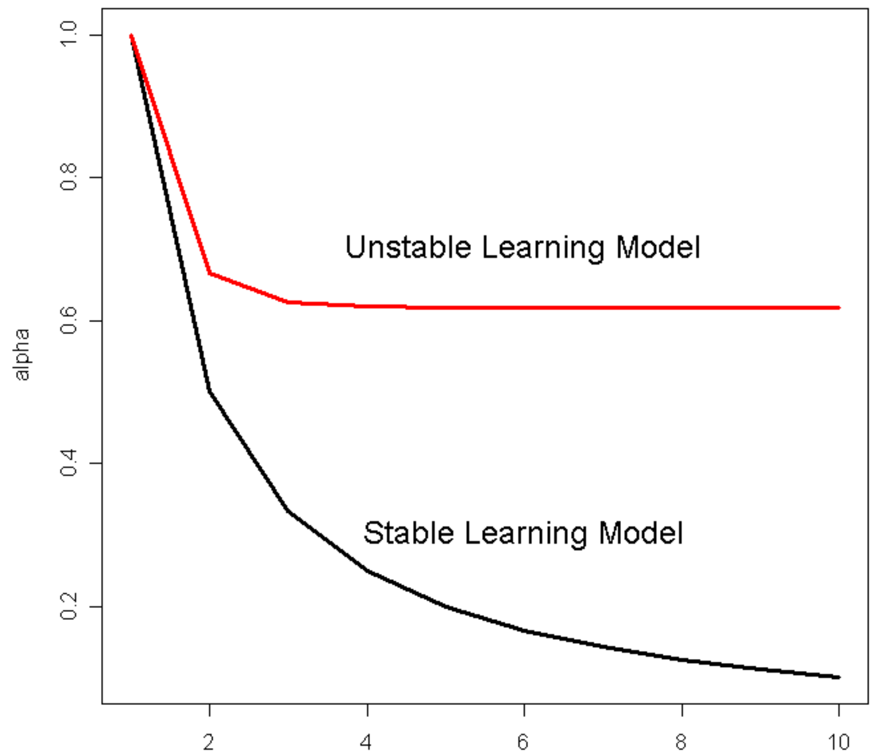

Figure 1: Learning Parameters (α_{it}) Over Time

These observations on how α_{it} changes with each experience also provide insights into how a person updates his or her uncertainty about the firm's quality level. For consumers who act according to the **stable learning model**, the uncertainty about the underlying quality level, U_{it}, asymptotes to zero. This is not true for the unstable case because quality is constantly changing.

2.2.4 Consumer beliefs about the environment

The preceding model development assumes that consumers have perfect knowledge about the ratio of dynamic to stable variance (i.e., D/S) and that they select the correct learning model. These are strong assumptions since a number of empirical studies indicate consumers find it difficult to learn in probabilistic settings and consequently often use an inappropriate schema (Brehmer 1980; Hoch and Deighton 1989). Further, some research suggests that in small decisions with stable outcomes, people try the low-payoff options too often, perhaps suggesting consumers act as if those stable outcomes could have changed (Biele, et. al. forthcoming). Hence, consumers may hold an inaccurate view of whether quality is changing or stable. Such a view in turn influences how they update their beliefs, and ultimately make choices.

This leads us to allow for a mismatch between the consumer's view of the world (unstable or stable) and the actual state of that world. We do this by extending our basic formulation to acknowledge that consumers do not have perfect knowledge about the environment. We suggest that consumers use their prior experiences with other similar situations to infer a value of D/S for this environment. They then blend this prior belief with their current product experience to update their belief about the correct model.

Specifically, we assume that consumers initially have uncertain beliefs about what kind of environment they face, M, and that these beliefs differ across consumers. Thus, each consumer has an uncertain prior belief, p(M), about which model is most appropriate given the environment. Moreover, this uncertainty about the environment implies that consumers have multiple beliefs about quality, conditional on the environment. The consumer's prior belief about the environment, p(M), mixes with the consumer's prior beliefs about quality $b_{it-1}(Q_{t-1} \mid M)$, which are conditioned on the model. This leads to a joint prior belief on M and Q_{t-1}.

2.2.5 Consumer belief updating about quality and the environment

We now integrate our model of how consumers update beliefs about quality with our model of how consumers form beliefs about the environment. Given the joint prior belief about Q_{t-1} and M, we assume consumers learn in a Bayesian manner which model to select. To illustrate, assume that M is a Bernoulli variable selecting between a model with a D/S value of zero (M=0) and model with a positive D/S value (M=1). We use the same assumptions about the delivery process, $p(DE_{it} \mid -)$, and the conditional prior beliefs about the underlying quality, $b_{it-1}(Q_{t-1} \mid M)$, so that these distributions are normally distributed. Finally, we assume that consumers know the true value of S, so that D is recovered exactly, conditional on model. This leads to the following joint and marginal posterior distributions after updating based on a new DE_{it}[5]:

$$p(Q_t, M=m \mid -) \approx [p(DE_t \mid Q_t, S)\, b_{t-1}(Q_t \mid WE_{t-1,m}, U_{t-1,m}, M=m)\, p(M=m)] \quad (7)$$
$$p(Q_t \mid -) = \Sigma_m\, [b_t(Q_t \mid DE_t, WE_{t-1,m}, U_{t-1,m}, S, M=m)\, p(M=m)] \quad (8)$$
$$p(M=m \mid -) = \frac{p(DE_t \mid WE_{t-1,M}, U_{t-1,M}, S, M=m)\, p(M=m)}{\Sigma_k\, [p(DE_t \mid WE_{t-1,m}, U_{t-1,m}, S, M=k)\, p(M=k)]} \quad (9)$$

[5] Note that in what follows we have dropped the (i) subscript for convenience.

These updated distributions, along with equations 3-5, can be used to gain insights concerning the types of learning patterns we should expect from consumers.

First, we examine the effects of model selection when the consumer initially encounters an environment. For simplicity of explication, we assume the consumer has the same prior mean beliefs for their first encounter regardless of which model might hold, i.e., $WE_{t-1,M=0} = WE_{t-1,M=1} = WE_{t-1}$. Based on equation 3 and 5, the conditional updated means ($E[Q_t \mid M] = WE_{t,M}$) are a weighted average of DE_t and WE_{t-1} with the weight on the new experience, DE_t, being greater for the **unstable learning model** than the **stable learning model**. Equation 8 states that the marginal distribution of Q_t is a weighted average of the conditional distributions of Q_t, where the conditioning is on the model being used.

Two implications directly flow from this observation. First, the consumer's best guess for the unobserved quality level after the first encounter, $E[Q_t]$, is between the updated means conditioning on M=0 and M=1. Second, as a result, the transition from WE_{t-1} to $E[Q_t] = WE_t$ will rely more on the observed experience, DE_t ,than predicted by the **stable learning model** and less than predicted by the **unstable learning model**. However, as long as the consumer is uncertain about the environment, such a weighting will be consistent with <u>some</u> **unstable learning model**. Further, even if consumers incorporate the same new experiences, our assumption that consumers have different prior experiences, and hence different prior beliefs, implies that the updated mixture of the stable and **unstable learning models** will also differ between consumers. These different mixtures effectively result in different learning models across consumers (i.e., values of D/S). We state these results as the following propositions:

> P1. Even when the underlying quality level is stable (i.e., D=0), consumers' are more likely to exhibit learning patterns consistent with an **unstable learning model** (i.e., D/S>0) than a **stable learning model** (i.e., D/S=0).
>
> P2. Different consumers will exhibit behavior that is consistent with different learning models (i.e. values of D/S), even given the same experiences with the product.

Taken together, P1 and P2 state that consumers will exhibit heterogeneous learning behavior. However, if consumers learn normatively about the environment they face, consumers' beliefs about a specific environment and the consumer learning models observed in that environment should correspond more closely to the actual environment as they gain more experience (e.g., Rieskamp and Otto 2006). In our setting, learning models are inferred from behavior, while beliefs about the environment are measured directly. However, the result of either approach should be influenced by differences across environments. Specifically,

> P3a. Consumers' experiences will influence how they learn. Those consumers who experience environments with changing quality levels act according to more **unstable learning models** (i.e., higher D/S values) than consumers who experience a stable quality environment.
>
> P3b. Consumers' experiences will influence their beliefs about the environment. Those consumers who experience environments with changing quality levels will believe the environment is more dynamic than those that experience a stable quality environment.

P3 states that consumers are sensitive to the environment they encounter. However, our model implies that consumers are also sensitive to their actual beliefs. Thus, consumers that perceive the environment as changing are more likely to use an **unstable learning model** than those that perceive the environment as stable, regardless of the accuracy of these beliefs. More formally,

> P4. Consumers' beliefs about the environment are related to their learning behaviors. Those consumers who act according to an **unstable learning model** are more likely to perceive the environment as dynamic than consumers who act according to the **stable learning environment**.

2.2.6 Model implications for accuracy and choice

Although P3 is consistent with the idea that consumers will ultimately learn the correct model, our prior discussion implies that consumers' beliefs would rationally be a mixture of the possible models and thus not necessarily close to the model that predicts best given the true market conditions. Hence, we would expect that consumers will differ in their ability to predict what will happen next. The more closely the consumer's learning model at a given point in time matches the actual environment, the more accurate the consumer will be in predicting the experience associated with the next product encounter. More formally,

> P5. Accuracy in predicting what will happen next time is an increasing function of how closely the consumer learning model matches the environment.

Finally, we extend our model to incorporate product preference and choice. There is an extensive literature that connects consumer beliefs about the product to these two output measures (e.g., Rust et al 1999). We write out a consumer utility (V_{ijt}) model for a product j after an experience t as a function of the person's current mean belief and the confidence (C_{ijt})[6]. Substituting in equation 3 for the current mean belief we get

$$V_{ijt} = \gamma_1 * [\alpha_{ijt} * DE_{ijt} + (1 - \alpha_{ijt}) * WE_{ijt-1}] + \gamma_2 * C_{ijt} + \varepsilon_{ijt} \quad (10)$$

Because α_{ijt} is larger the more a consumer acts according to the **unstable learning model,** we should find differences in the relative weights consumers place on their current experience versus their prior mean belief depending on their specific learning model:

> P6. Consumers who act in accordance with an **unstable learning model**, relative to those who act in accordance with a **stable learning model**, will, regardless of the actual environment, place more weight on the current experience and less on the prior belief when making a product choice.

Taken together, P2 and P6 imply that we should observe four types of behavior, as shown in Table 1. From P2 we expect consumers to differ in how accurately their learning models match a given environment. Some consumers will act according to the appropriate learning model and others will act according to an inappropriate learning model. Further, from P6, the differences in models will influence how they make choices about the products. Consequently, we should find that some consumers act in accordance with the environment, make better predictions about future experiences, and as a result make product choices that match their desired attributes. In contrast,

---

[6]Confidence here is conceptually the inverse of the uncertainty about the product offering. Although our model has little to say about the confidence parameter, we include it because we feel it is an important control variable when examining the role of beliefs about Q_t.

consumers who act according to an inaccurate model will make more errors about the product quality and, as a result, more errors in choice. These errors imply that the consumer, if the consumer recognized the true quality of the products, would have made a different choice. We classify these inaccurate consumers as excessively loyal (i.e. they maintain stable beliefs about the product when quality deteriorates), excessively critical (i.e. they maintain stable beliefs about the product when quality improves), or promiscuous (i.e. they change beliefs too much after one bad experience when quality is stable).

Table 1: Outcomes of Learning Model—Environment Matching

Quality Environment	Stable Learning Model	Unstable Learning Model
Stable	Accurate	Promiscuous
Unstable Increasing	Excessively Critical	Accurate
Unstable Decreasing	Excessively Loyal	Accurate

In the following section we present two studies designed to examine these propositions. The overarching design feature of these studies is to place consumers in a setting where they participate in numerous product experiences. After each experience we ask the consumers to indicate their beliefs concerning what they expect will occur in their next experience, i.e. obtain their estimate of WE_{it}. This allows us to determine how consumers blend their prior beliefs with their current experience to form an updated belief structure, i.e. we obtain estimates of α_{it}, and, thus, gain insights into their learning process. We then gain a better understanding of how consumers select a quality updating model and how selecting the wrong model (given the environment) influences behavior.

2.3 Study 1: Updating quality beliefs

One hundred and fifty one undergraduate students participated in a laptop-based research study. The study lasted approximately 30 minutes and participants were paid $7. Participants were asked to imagine that they were taking their car to the CarLube oil change center for an oil change. Participants visited CarLube ten times. The key quality variable was the wait time at CarLube during the oil change. During each visit, the participants "experienced" a wait time where each second waited during the experiment represented one minute at the CarLube shop. In addition, they were told the simulated wait time (i.e., DE_{it}). Prior to each experience, and after the tenth experience, the participants were asked what they expected the wait time would be in the next period (i.e. WE_{it}). Hence, each subject had ten experiences (DE_{it}) and each subject provided eleven WE_{it} estimates. Finally, after the 10^{th} experience we collected measures from subjects about their perceptions of the environment they encountered (i.e., whether they felt it was dynamic or stable).

Participants were randomly assigned to different conditions. In periods 1-4 participants were exposed to one of the four sets of experiences that were generated using Equation 1. In each case the underlying quality level was the same and was held fixed. In addition, the stable variance was set to be low relative to this mean. We use these initial four periods, which correspond to the empirical settings used in existing quality updating research (Rust et al. 1999; Boulding et al. 1999) to see if we could replicate previous results in our setting.

In periods 5-10, we varied the underlying quality environment and thus the experiences (i.e. DE_{it} values) in two ways, across eight conditions in total. First, participants in six conditions continued to see wait times that had a mean value that was indistinguishable from that used for periods 1-4. However, in two of these six conditions the stable variance was substantially increased for periods 5-10. These two higher stable variance conditions had 55 participants. Another 37 participants continued in one of the four stable low conditions. Finally, 59 participants, across the last two conditions, saw wait times in periods 5-10 generated from a process in which the underlying mean was determined via a random walk. Thus, the observed variation in experiences for these 59 participants was due to both a (low) stable variance and the unstable variation associated with the random walk.

Because the experimental conditions were developed by randomly drawing experiences from the two different stochastic quality processes, we statistically analyzed these eight conditions to ensure that they reflected the intended manipulations. For each condition we calculate the BIC for each of the two potential data generating processes and confirmed that the two unstable conditions were most likely to be drawn from a changing quality process and the six stable conditions were most likely to be drawn from the intended stable quality process.[7]

To control for other factors that might impact the updating process, participants were not given any external cues of the wait time change that occurred between periods 4 and 5. Also, they were told that all other aspects of the oil change were completed as expected. As a result, the only source of information about the underlying quality was their pattern of wait time experiences.

We use the measured will expectations to estimate empirical models based on fitting either the **stable** or **unstable learning models**. In addition, for comparison we fit a "theory-less" model that allows consumers to combine their prior belief and current experience in any fashion over time. We also allow consumer heterogeneity in the learning model, enabling us to recover different patterns of updating (α_{it}) across consumers.

Although the data equation for our model is based on equation 3, we initially do not constrain the parameters and assume that errors in the updating process across consumers and over time are uncorrelated:

$$WE_{it} = \beta_{1t} * DE_{it} + \beta_{2t} * WE_{it-1} + \varepsilon_{it}, \qquad (11)$$

[7] Note, however, that although the stable sets of stimuli were most likely to come from a stable stochastic process, and the unstable sets of stimuli were most likely to have come from an unstable stochastic process, there is some probability that the converse could be true. Thus, we might expect some participants to infer the wrong underlying model.

where ε_{it} is distributed normal with mean 0 and variance τ^{-1}. In this unconstrained model we use diffuse normal priors for the β's and, in all relevant models, a diffuse gamma prior for τ^{-1}. For most of our other models, we constrain the coefficients in equation 11 to follow the structural Bayesian updating scheme. To implement these constraints, we assume that the weights β_{1t} and β_{2t} sum to one, that they take values between 0 and 1, and that the transition process for the α_t values follows that defined by either the **unstable** or **stable learning model** (where the only difference is the value of D/S). These constraints are represented as follows:

$$\beta_{1t} = (1-\beta_{2t}) = \alpha_t \quad (12a)$$
$$\alpha_t \in (0, 1) \quad (12b)$$
$$\alpha_t = (\alpha_{t-1} + D/S) \,/\, (\alpha_{t-1} + D/S + 1) \quad (12c)$$

Further, we allow for consumer heterogeneity in the pattern of α values. Specifically, we allow consumers to belong to one of G latent groups with each group having either a D/S value of zero ($g \in sta$) or an uncertain, positive value ($g \in uns$). Each group also has an uncertain initial α value (α_{g1}) that starts the transition process of (12c). More formally,

$$\alpha_{g1} \sim Beta(1,1) \quad (13a)$$
$$D/S_{g \in uns} \sim Gamma(m,c) \quad (13b)$$
$$D/S_{g \in sta} = 0 \quad (13c)$$
$$g_i \sim Categorical_G(P) \quad (13d)$$
$$P \sim Dirichlet_G(1, 1, \ldots) \quad (13e)$$

where m and c are parameters selected to maintain a diffuse gamma prior, G is the number of different groups, and uns and sta are groups with a priori values of D/S that are either uncertain and positive or certain and 0. This specification assumes that each consumer acts according to a single learning model with a single value of D/S over time. However, across consumers the learning models and value of D/S differ. Identification of the values of D/S and α_1 is achieved by pooling over consumers to find similar patterns of response to new experiences over time.

Because, a priori, we theorize that there will be some heterogeneity, we estimate multiple combinations of stable and unstable learning groups and use log marginal likelihoods to determine the best fitting set of groups. We analyze the experiment in two portions: the first four periods and the last six periods (see the appendix for details on the Bayesian estimation).

2.3.1 Results in Periods 1-4

In the first four periods all consumers experience the same low variance around a stable underlying quality. We estimate our empirical models using the methodology described above to determine how well our different models capture the observed consumer behavior. In each case, we report the log marginal likelihoods, the portion of consumers who use a **stable learning model** ($P_{sta} = \sum_{g \in sta} P_g$), and our uncertainty about P_{stat} in terms of the 5% and 95% quantiles for the three types of models (i.e., **stable**, **unstable**, and theory-less).

We start by looking at the unconstrained model (model 1), which we use as our baseline case. Specifically, we assume homogeneity, with all consumers using equation

11 to update their beliefs. However, we do not place any restrictions on β_{1t} and β_{2t}. We first compare this model with homogeneous structural models in which all consumers are constrained to use either the **stable** (model 2) or **unstable** (model 3) learning model. We find the constraints associated with the **unstable learning model** on behavior (model 3) marginally improve the log marginal likelihood (by 2) as compared to the unconstrained model (model 1). In contrast, the constraints of the **stable learning model** (model 2) worsen it substantially (by 21). Hence, if we assume all subjects act according to the same learning model, the **unstable learning model** fits the observed behavior best, meaning zero consumers are classified with the **stable learning model**, even when all consumers experience a stable quality environment. This strongly supports Proposition 1.

Table 2: Model Fit and Portion Stable Learners for Initial Four Periods (Study 1)

Model (Number of Groups)*	Probability of Stable Learning Model**	Log Marginal Likelihood
1. Theory-less	N/A	-1551
2. Unstable (0) – Stable (1)	100%	-1572
3. Unstable (1) – Stable (0)	0%	-1549
4. Unstable (1) – Stable (1)	67% (.43, .80)	-1504
5. Unstable (2) – Stable (1)	25% (.10, .42)	-1463
6. Unstable (2) – Stable (2)	28% (.13, .45)	-1454
7. Unstable (2) – Stable (3)	31% (.16, .49)	-1450
8. Unstable (2) – Stable (4)	33% (.18, .54)	-1453
9. Unstable (3) – Stable (2)	19% (.04, .38)	-1448
10. Unstable (3) – Stable (3)	**23% (.07, .43)**	**-1445**
11. Unstable (3) – Stable (4)	27% (.10, .49)	-1449
12. Unstable (4) – Stable (3)	20% (.04, .42)	-1449

* - Unstable(x) – Stable (y) is a model with x Unstable groups and y Stable groups
** - Values are posterior mean with the 5% and 95% quantiles given in parentheses.

Next, we allow for heterogeneity in behavior. In search of the best fitting model, we estimate multiple models that allow various degrees and types of group heterogeneity. We find the best fitting model is model 10. This model has three **stable learning** groups (differing only in their initial priors for α_{g1}) and three **unstable learning** groups (differing in initial priors and values of D/S). In this case, 23% of consumers are classified as acting according to a **stable learning model**, i.e., the posterior mean of P_{stat} = 0.23. Hence, even with heterogeneity we find support for Proposition 1. This indicates that the **unstable learning model** better reflects learning behaviors than the **stable learning model** in the early periods. Further, we find that introducing heterogeneity in learning models increases fit substantially (by 104), strongly supporting Proposition 2: heterogeneity exists in the way consumers learn from the same experiences.

2.3.2 Results in Periods 5-10

We next analyze the data from periods five to ten when consumers encountered service experiences that were one of the following: stable with a small variance, stable with a large variance, or unstable. We report the fit and proportion of stable learners for each of these three conditions across multiple numbers of latent classes in Table 3. Based on log marginal likelihood, we find the best fitting model to have three **unstable groups** and one **stable group** for the unstable condition, three **unstable** and two **stable** groups for the stable high condition, and two **unstable** and two **stable groups** for the stable low condition. These models, respectively, have 12% (unstable condition), 38% (stable high condition), and 37% (stable low condition) of consumers using the **stable learning model** at the mean of the posterior distribution of P_{stat}.

To determine if these proportions differ, and thus provide support for P3a, we compare our estimates for the proportion of consumers who are classified as **stable** (P_{sta}) for the three different conditions. We do this by comparing the posterior distributions for each condition. Specifically, we subtract the P_{sta} posterior values for the unstable condition from those of the two stable conditions. We denote these differences as $\delta_{stable\ low}$ and $\delta_{stable\ high}$ respectively, and find that $\delta_{stable\ high}$ has less than 2.5% of the posterior below 0 and $\delta_{stable\ low}$ has less than 6% of the posterior below 0. These results support P3, suggesting that consumers learn over time, through experiences alone, about their learning environment.

To provide additional evidence for P3a, we examine the population D/S values for our three conditions. We do this by constructing a posterior distribution for the median value of D/S across consumers for each condition. The means of these posterior distributions are 0.34, 0.17, and 0.27, respectively, for the dynamic, stable high, and stable low conditions. The differences in means between the dynamic and stable conditions are highly significant (p-value < .01) in both cases. This suggests that the average consumer differs across conditions (underlying quality environments) in how unstable her learning behaviors are and serves as further evidence for P3a.

Table 3: Model Fit and Portion Stable Learners Last Six Periods (Study 1)

Condition	Model (Number of Groups)*	Probability of Stable learning model**	Log Marginal Likelihood
Dynamic	Unstable (1) – Stable (1)	18% (.06, .35)	-856
	Unstable (2) – Stable (1)	13% (.05, .25)	-823
	Unstable (2) – Stable (2)	16% (.06, .28)	-823
	Unstable (3) – Stable (1)	**12% (.01, .24)**	**-820**
	Unstable (4) – Stable (1)	10% (.01, .23)	-820
Stable High	Unstable (1) – Stable (1)	48% (.28, .72)	-628
	Unstable (2) – Stable (1)	26% (.06, .55)	-604
	Unstable (2) – Stable (2)	47% (.24, .67)	-592
	Unstable (2) – Stable (3)	50% (.28, .70)	-595
	Unstable (3) – Stable (2)	**38% (.15, .59)**	**-586**
	Unstable (3) – Stable (3)	42% (.21, .62)	-586
	Unstable (4) – Stable (2)	38% (.16, .58)	-586
Stable Low	Unstable (1) – Stable (1)	78% (.60, .90)	-337
	Unstable (1) – Stable (2)	78% (.62, .91)	-334
	Unstable (2) – Stable (1)	22% (.02, .55)	-333
	Unstable (2) – Stable (2)	**37% (.10, .69)**	**-327**
	Unstable (3) – Stable (2)	32% (.07, .62)	-327
	Unstable (2) – Stable (3)	45% (.17, .74)	-327

* - Unstable(x) – Stable (y) is a model with x Unstable groups and y Stable groups

** - Values are posterior mean with the 5% and 95% quantiles given in parentheses.

Thus far we have inferred the consumer's choice of learning models from their observed behavior. However, we also explicitly asked our participants to tell us whether and how they thought the underlying quality was changing. Specifically, after participants experienced all ten encounters, we presented them with their predicted and experienced wait times for the latter periods. We asked them to assign 100 points to five possible options to reflect the (quality) environment to which they were responding. The options included (a) trends, (b) mean shifts, (c) random changes in mean, (d) no change, and (e) changes in variability[8], with the specific terminology based on verbal protocols gathered during a pretest presented in the appendix. We classified the five options as corresponding to either the unstable (a, b, and c) or stable (d and e) environments.

We aggregate over periods five to ten to form an average proportion of perceptions about the environment that are unstable for each participant. While this perceptual measure is fixed for an individual, our measure of the consumer's learning model is uncertain and varies over the posterior draws. For each posterior draw and corresponding assignment of individuals to learning models, we estimate the effect of learning model and condition on consumer perceptions of the environment. We estimated this relationship by drawing 100 simulations from the posterior distribution for a linear model with perceptions of environment as the dependent variable and the conditions and whether we classified them as being in an unstable learning model (denoted **unstable)** as dummy independent variables[9]. We repeat this process for each posterior sample of the group identities and then pool across the regression posteriors for each of the learning model posterior draws to form the marginal distribution of the coefficients. In this way, we are able to capture model uncertainty from both stages.

---

[8] Note this changes the variability around a stable mean.

[9] To draw these simulations, we use a standard Gibbs sampler starting at the maximum likelihood estimates.

This analysis lets us examine two propositions. First, we find significant main effects for the conditions, which support P3b. The estimates for condition are 0.40, 0.50, and 0.76, respectively, for the stable low, stable high, and dynamic conditions. Forming pair wise differences, the differences between the unstable and each of the two stable conditions have less than a 1% probability of not being greater than zero, whereas the difference between the stable high and stable low conditions has less than a 6% probability of not being greater than zero.[10] This result provides strong evidence that participants can appropriately perceive whether the environment is changing based on the experiences they face.

Second, we can examine P4, that consumer learning patterns reflect actual differences in how consumers perceive the environment. We do this by noting the mean effect of the dummy variable **unstable** is 0.10 with less than 9% of the probability mass below zero. Further, the effect of the variable **unstable** without controlling for condition is more significant. In this case, the mean effect is 0.19 and less than 2% of the probability mass is below zero. Hence, what we estimate as the best fitting learning model correlates with how subjects perceive the environment. This provides evidence that consumers match their learning model with their beliefs about the environment and supports P4.

---

[10] These results also hold when the unstable learning model dummy is not included, but the probability of differences between the two stable conditions drops below significant levels.

Finally, we turn to P5. We define the correct model as the model (either unstable or stable learning) the consumer should choose if she knew which environment she faced with perfect knowledge. We are interested in knowing whether consumers using the incorrect model have greater predictive error of what will happen on their next experience occasion (PE) than consumers using the correct model. To test this proposition, we examine whether, controlling for condition effects, the predictive error for the incorrect model is greater than those of the correct model.

We construct the predictive error measure by summing the squared errors between the WE_{it-1} and DE_{it} for periods five to ten. We use a log transform to better accommodate a linear model. While PE is fixed for an individual, which model an individual uses is uncertain (i.e., varies over the posterior distribution). Like the previous analysis, we estimate a linear model for each posterior sample of the learning model variable and then combine the samples. In this case, the linear model has PE as the dependent variable with dummy independent variables for each condition and for whether the individual uses the wrong model. The results are shown in Table 4.

As shown in the first three rows of estimates in Table 4, we find strong evidence that predictive error increases as the variation in the underlying quality process increases. To examine the validity of P5, we examine the distribution of the wrong model effect. Specifically, we examine whether the effect of having the incorrect model is positive, i.e., that having the incorrect model causes one to have more predictive error.

We find that the portion of the distribution that is below zero is less than 5%, supporting P5. This result also supports the idea that the learning models we estimate reflect differences in how accurately consumers learn in their environment.

Table 4: Effect of Wrong Model on Predictive Error (Study 1)

Parameter	Coefficient Values**
Stable Low Intercept	1.28 (1.08, 1.48)*
Stable High Intercept	2.95 (2.78, 3.13)*
Dynamic Intercept	3.80 (3.69, 3.90)*
Wrong Model Effect	0.27 (0.03, 0.49)*

* Denotes less than 5% of the sample is below zero. Differences between conditions are significant in all cases with less than 5% probability below zero.

** - Values are posterior medians with the 5% and 95% quantiles given in parentheses.

2.3.3 Summary of Study 1 findings

Consumers appear to act as if they use Bayesian updating schemes to learn about the quality level. We find these schemes fit the data better than an unconstrained, a-theoretical model. We also find that, on average, consumers act in a manner more consistent with the **unstable learning model** than the typically assumed **stable learning model**. However, not all consumers act in the same way; we find heterogeneity in how consumers respond to the same stimuli. Initially, when faced with a stable quality environment, more consumers respond with the **unstable learning model,** but a significant (approximately 25%) portion respond with a **stable learning model**. After some experiences, we find that consumers are sensitive to their environment.

Specifically, more consumers respond to unstable quality environments by using an **unstable learning model** than if they face stable quality environments, and they are more likely to report that they believe the quality level is unstable. Moreover, consumers' beliefs about the environment and their learning models are consistent. Finally, we find the predictive error is greater for consumers that apply the wrong learning model than for those that act according to the correct model. In summary, we find empirical support for propositions 1-5. This overall pattern of results strongly supports that consumers behave as if they are using our learning model. In the next study, we seek to extend these results and add choice and preference measures, so that we can explore how learning patterns influence choice, i.e., test P6.

2.4. Study 2: Making product choices

This study also involves consumer experiences where the underlying quality of the product is either stable or changing. However, the experimental conditions in Study 2 differ from Study 1. The main difference is that participants experience two different products. Of the two products, one (the benchmark) product is always portrayed in a stable quality condition across the two experimental conditions, while the other (the focal) product is manipulated as having either a stable or unstable quality process. This greatly adds to the complexity of the learning context for participants, since now they have to keep track of the experiences generated by two products instead of by one. It also means that the perceived quality level could change even when it is objectively stable, if consumers define quality as a relative concept (Boulding et al. 1993).

We use such a setting for two reasons. First, it allows us to replicate our results in a different, more complex setting. Second, it enables us to study choice behavior between two competing products. In what follows, we explain the experimental method in detail, present the analysis and results, and conclude with an overall discussion of Study 2.

2.4.1 Experimental method

The participants were 155 undergraduate students who took part in a laptop-based experiment and were compensated between $5 and $8 depending on the accuracy of their answers, making the study incentive compatible. The cover story stated that the University was experimenting with providing broadband services over power lines (BBPL). BBPL was explained as an experimental technology, which currently exists, but is still not commercially viable. Part of the cover story indicates that two different software companies are vying for a contract with the University to host BBPL and that BBPL-related software has a major influence over the performance of the service. Participants are told the two companies provided the University free software versions for use during the testing stage and that the participants are part of a test group set up by the University to gauge the quality of the service.

Participants sequentially used both services to download the same large data file. During the course of the study they were asked to do this 10 times, each time alternating between the software from the two rivals. As in study 1, participants not only experience the simulated download times but are also told how long it takes to download the file (DE_{it}). After each participant experiences the jth service, she is asked what they expect the download time will be next time they use this service (WE_{ijt}) and how uncertain they are about this prediction. In addition, similar to the approach in Rust et al. (1999), after the 4th, 5th, 7th, 9th and 10th experiences of the two rival services, participants are asked to indicate relative preference (allocation of 100 points) between the two products. The stimuli are fully counterbalanced to control for order effects relating to the benchmark product versus the focal product. No significant order effects were found, and hence in the analyses that follow, order is ignored.

Three features of these ten pairs of experiences are noteworthy. First, for 115 of the 155 participants the focal service experiences came from one of three different patterns of data that were generated to reflect an unstable process, where the mean quality changed each period but the overall mean of the experiences over the ten periods was equal to the benchmark focal service mean. For the remaining 40 participants the focal service experiences came from a stable process, with a mean equal to the benchmark service but the variance of the service was two standard deviations greater than that of the benchmark service. Second, the total variation of the observed experiences associated with the focal service was approximately equal, regardless of whether the underlying process was stable or unstable. Third, in periods 5 and 10 participants saw identical decreases in download speeds from the previous period for both the focal and the benchmark services. To make this common decrease in download speeds noticeable, we had to select a decrease that was greater than one standard deviation from the current mean for the high variation, but stable, focal service. Thus, the change in download speed for the benchmark service was over two standard deviations from its current (stable) mean.

2.4.2 Replication of relevant propositions

We explore the robustness of our Study 1 findings by applying the same essential empirical analysis to Study 2 data with several necessary adjustments to reflect the differences in the experiments. First, we center our attention on periods 5 to 9. We do this for two reasons: (1) to allow some time for participants to learn which environment they are encountering, and (2) to avoid the large unusual decreases between periods 4 and 5 and periods 9 and 10. Second, we estimate two independent learning models for each consumer corresponding to the two services. Third, we estimate our participants' responses to the benchmark product taking into account their focal experiences. Specifically, in addition to analyzing the unstable and stable focal conditions, we analyze the benchmark paired with stable and benchmark paired with unstable conditions. For each of these four groups, we find the best fitting model. Table 5 presents the model fit and portion stable learners in each condition.

Table 5: Model Fit and Portion Stable Learners by Condition for Periods 5-9 (Study 2)

Condition	Model (Number of Groups)*	Probability Stable learning model**	Log Marginal Likelihood
Focal – Unstable	Unstable (1)—Stable (1)	9% (.04, .17)	-2735
	Unstable (2)—Stable (2)	31% (.12, .45)	-2691
	Unstable (3)—Stable (2)	27% (.08, .43)	-2684
	Unstable (2)—Stable (3)	33% (.17, .46)	-2686
	Unstable (3)—Stable (3)	**32% (.16, .46)**	**-2680**
	Unstable (3)—Stable (4)	34% (.20, .47)	-2685
	Unstable (4)—Stable (3)	32% (.16, .45)	-2684
Focal – Stable	Unstable (1)—Stable (1)	35% (.10, .62)	-1021
	Unstable (2)—Stable (1)	16% (.15, .41)	-1017
	Unstable (1)—Stable (2)	42% (.19, .66)	-1017
	Unstable (2)—Stable (2)	**27% (.07, .54)**	**-1012**
	Unstable (3)—Stable (2)	23% (.05, .47)	-1019
	Unstable (2)—Stable (3)	32% (.12, .56)	-1013
	Unstable (3)—Stable (3)	30% (.11, .54)	-1013
Benchmark – Paired Unstable	Unstable (1)—Stable (1)	9% (.03, .18)	-2556
	Unstable (2)—Stable (1)	7% (.01, .15)	-2551
	Unstable (1)—Stable (2)	31% (.11, .54)	-2551
	Unstable (2)—Stable (2)	**21% (.05, .45)**	**-2549**
	Unstable (3)—Stable (2)	18% (.04, .39)	-2552
	Unstable (2)—Stable (3)	27% (.09, .49)	-2550
	Unstable (3)—Stable (3)	26% (.08, .48)	-2550
Benchmark – Paired Stable	Unstable (1)—Stable (1)	43% (.05, .90)	-849
	Unstable (2)—Stable (2)	46% (.11, .83)	-845
	Unstable (3)—Stable (2)	39% (.09, .75)	-845
	Unstable (2)—Stable (3)	53% (.18, .85)	-845
	Unstable (3)—Stable (3)	**52% (.18, .82)**	**-843**
	Unstable (3)—Stable (4)	57% (.25, .84)	-844
	Unstable (4)—Stable (3)	50% (.17, .80)	-845

* - Unstable(x) – Stable (y) is a model with x Unstable groups and y Stable groups
** - Values are posterior mean with the 5% and 95% quantiles given in parentheses.

The results suggest clear support for heterogeneity within condition (P2). Looking across conditions, most consumers respond with an **unstable learning model**. Only one condition has more than 50% of consumers using a **stable learning model**: the benchmark condition paired with the stable focal condition. In the complex environment of Study 2, we take this as evidence in support of most consumers using the unstable model (P1). We also note that this result occurs after four experiences within a consistent environment and hence the consumers had more time to learn about the environment than what was allowed in our analyses in Study 1.

Support for the environment influencing model choice (P3a) is less clear, however. To begin, we analyze the probability of a participant being classified as stable (P_{sta}), which is the sum of the aggregate probabilities of belonging to one of the stable learning groups. All of these values overlap greatly across the four conditions, despite apparent mean differences. The pair wise differences between the focal patterns have more than 10% of the posterior on the opposite side of zero from the mean difference. Hence, it appears that there is little difference between the responses to quality environments in terms of the portion selecting a zero value for D/S. This does not support P3a.

However, as in Study 1, we also examine the average person's unstable model in terms of D/S in each condition (i.e., the median value of D/S from the distribution of individual D/S values). The mean of these D/S values are 3.98 and 1.86 respectively for the focal unstable and focal stable conditions. These means differ significantly (p-value < .01). Combining these two results related to P3a suggests an interesting nuance. Consumers may not differ across environments in terms of the probability of using **an** unstable model, but the environment seems to affect **the** specific unstable model people use. This pattern of results provides support for P3a.

Finally, from Table 5, we note that the focal pattern (unstable or stable) appears to influence the way consumers learn about the benchmark pattern. Recall, the benchmark pattern experiences are the same regardless of the focal pattern. However, the mean probability of using a **stable learning model** differs between the benchmark paired with unstable ($P_{sta} = 0.21$) and the benchmark paired with stable ($P_{sta} = 0.52$). The difference between these two probabilities has less than 7% of the distribution below zero. Further, the same analysis on the benchmark conditions of the average subject results in values of D/S equal to 6.30 and 0.39, respectively, for the benchmark patterns paired with unstable and stable focal patterns. These means differ significantly (p-value < .01). This pattern of results suggests a potentially important role for the influence of competitor product offerings on consumers' perceptions of the stability of the underlying quality level of a product, and thus consumer learning.

We test our hypothesis on the prediction errors (i.e., P5) using the same basic analysis plan as in Study 1. The results are presented in Table 6. The effect of the wrong model coefficient is significant and positive with the effect having over 99.9% of its posterior above zero.

Table 6: Effect of Wrong Model on Predictive Error (Study 2)

Parameter	Coefficient Values**
Benchmark-Stable Intercept	8.69 (8.51, 8.86)*
Benchmark-Dynamic Intercept	8.57 (8.39, 8.72)*
Focal-Stable Intercept	11.49 (11.30, 11.68)*
Focal-Dynamic Intercept	10.26 (10.17, 10.36)*
Wrong Model Effect	0.25 (0.08, 0.42)*

* - Denotes less than 5% of the sample is below zero.
** - Values are posterior medians with the 5% and 95% quantiles given in parentheses.

2.4.3 Preference results

We use the preference data collected after experiences 5, 7, 9, and 10 to test P6. We treat the preference data as a finer measure of choice and investigate effects on preferences. Our preference model assumes an underlying utility specification with additive linear components as in equation 10. For the focal product, this model suggests that as α_{it} increases, the marginal effect of DE_{it} increases, and the marginal effect of WE_{it-1} decreases. To link this prediction to the learning models, we note that **unstable learning models** have higher α values than **stable learning model**s, assuming a similar initial value of α.

We estimate equation 10 with several alterations that lead to the following estimating equation:

$$Pref_{it} = \boldsymbol{\theta_0}\, \mathbf{Prod_{it}} + \theta_1(a_{iFt}*DE_{iFt}) + \theta_2((1-a_{iFt})*WE_{iFt-1}) + \theta_3 C_{iFt} + \theta_4 WE_{iBt} + \theta_5 C_{iBt} + \varepsilon_{it} \quad (10e)$$

We note the changes. First, our dependent variable, Pref, measures consumer i's preference for the focal product relative to the benchmark product. Thus it captures the difference between the utilities of the two products (V_{iFt}-V_{iBt}), where F is for the focal product and B for the benchmark product. Second, we use WE_{it} rather than WE_{it-1} and DE_{it} for the benchmark product.[11] Third, we include control variables. Specifically, we include a vector of product-time dummies, $\mathbf{Prod_{it}}$, that control for the average difference in preferences between each pair of products for each period and consumer confidence measures, C_{iFt} and C_{iBt}, that control for any uncertainty related effects. Fourth, we capture the effects of the consumer's learning model by interacting DE_{ijt} with the estimated value of α_{ijt} (a_{ijt}) and WE_{it-1} with the estimated value, (1- a_{ijt}). These interacted values become the data. Thus, given a positive value for the coefficients on these interaction terms, consumers with higher values of α_{ijt} put more emphasis on the current experience and less weight on their prior beliefs. Because we hold fixed the experiences (DE_{ijt}) across individuals within a condition, the DE_{ijt} is a linear combination of the product-time control variables. As a result, information introduced through the interaction term with DE_{ijt} is due primarily to differences in learning parameters.

We estimate the model for preferences using equation (10est). We assume the stochastic term ε_{it} is distributed normally with mean zero and variance τ_v^{-1}. To complete

[11] The results for the primary variables of interest are the same even when the two terms are used instead of WE_{it}, but the effect of the benchmark product disappears. We conjecture this arises because of little variation, relative to the focal product, in delivered experiences for the benchmark product over encounters.

the statistical model, we assume reference priors, $p(\gamma,\tau_v) \sim 1/\tau_v$. Like in the estimation of predictive errors and uncertainty, we estimate the model for each posterior sample of the individual learning group identities and pool across these samples to examine the validity of P6.

Table 7: Preference Model Results (Periods 5, 7, 9, and 10)**

Variable***	Coefficient for Focal Product	Coefficient for Benchmark Product
WE_{it}	N/A	-0.014 (-0.027, -0.002)*
C_{it}	0.257 (0.196, 0.318)*	-0.298 (-0.363, -0.233)*
$(1- a_{it})\ WE_{it-1}$	0.021 (0.006, 0.035)*	N/A
$a_{it}\ DE_{it}$	0.019 (0.005, 0.034)*	N/A

* - Denotes less than 5% of posterior was on unexpected side of zero

** - Values are posterior medians with the posterior 5% and 95% quantiles in ( )s.

*** - Note that a_{it} is the estimated values of α_{it} for the focal product.

We present the key marginal distributions for the preference model in Table 7. The control variables (the Benchmark WE_{it} and C_{it} for both brands) all differ from zero and are negative reflecting the fact that increases in the consumers expected mean quality belief and certainty for the benchmark product lowers the relative preference for the focal product. We note that the effect of the confidence a consumer has is strong and approximately the same across the Focal and Benchmark products. The focal product coefficients for the estimated learning parameters (a_{it}) interacted with DE_{it} and WE_{it-1} are positive with very little overlap with zero (less than 3%). This result indicates that preferences are influenced by consumer learning models in the expected direction. The effect of current experiences is increasing in α_{it} and the effect of will expectations decreasing in α_{it}, indicating support for P6.

2.4.4 Summary of Study 2 findings

The results of this experiment build upon the results of Study 1, in that we replicate all the relevant findings from the first experiment in a new, much more complex setting. This increases our confidence that the **unstable learning model** is a good as-if model for learning behaviors. Further, in support of P6 we find that consumers acting according to the **unstable learning model** differ from those acting according to the **stable learning model** when making choices in that they place relatively more emphasis on the most recent experience and less emphasis on the prior quality belief.

2.5 Discussion

We introduce and analyze a normative model of learning when quality is changing. Analysis of this model reveals that consumers should update their quality beliefs more readily when quality is changing than when it is stable. We show the **stable learning model**, commonly used in the literature, is a special case of our more general model. We also develop propositions about consumer behavior based on our model and then conduct two experiments to test these propositions.

Our results reveal that, empirically, the **unstable learning model** fits actual consumer belief updating better than the **stable learning model**, even when true quality levels are stable. This result holds even after we allow for heterogeneity in consumer learning. Despite this apparent bias towards **unstable learning models**, we find that consumers are influenced by their environment and adjust which model they apply due to the types of experiences they encounter. This sensitivity is not uniform across a population of consumers and, as a result, some consumers use a learning model that does not correspond with the quality environment they experience, i.e., they have the wrong model. We demonstrate that these consumers with the wrong model have higher predictive error and that the learning models we estimate for subjects tend to match their perceptions of the environment. Finally, we find that consumers' learning models influence whether they weigh their prior beliefs or new experiences more when making a product choice. We also demonstrate that this effect on choice holds even when controlling for the confidence consumers have about the product quality. This overall pattern of results provides strong support that consumers are aware of their environment and act according to our learning model. Ultimately, we find that consumers account for potentially changing quality by altering how they learn.

Returning to Table 1, we find that while an important proportion of consumers act according to an accurate model, many more act inaccurately. Most of these inaccurate consumers, when facing stable quality, appear to be less loyal (i.e., they give less weight to their prior beliefs about the brand) than suggested by the **stable learning model**. However, some consumers, when facing changing quality, place a heavier emphasis on prior beliefs about quality than is normative, assuming perfect knowledge of the environment. These consumers are either 1) excessively loyal, if actual quality has decreased (i.e., they do not sufficiently update their beliefs in a negative direction) or 2) excessively critical, if actual quality has increased (i.e., they do not sufficiently update their beliefs in a positive direction).

These findings point to at least three direct managerial implications. First, because consumers differ in how they learn, firms may benefit from segmenting consumers along this dimension. For example, when launching a new product, managers may want to target typical **unstable learners**, whereas brand loyalty efforts might be more effective if targeted at typical **stable learners**. Here, additional research is necessary to explore how to identify a consumer's preferred learning model from more easily observed data.

Second, since we find that consumers can differentiate between different quality environments from the experiences they encounter, changing quality levels lead to more unstable learning and lower weight on the past beliefs. This change in weight reduces the utility of the brand from consistent beliefs, a type of brand equity (e.g., Erdem, et al. 1999; Rust, et al. 1999). Hence, firms need to recognize that while increasing quality perceptions, they might be altering their ability to maintain and develop brand loyalty. While our research does not address communications, per se, we do identify what might be a new goal from such communications: to influence the learning models consumers are using.

Third, our findings from Study 2 suggest that the focal product influences learning about the benchmark product. In particular, the results suggest that competition might influence not only the beliefs about quality (e.g., Boulding, et al. 1993), but also the way consumers learn about quality levels over time. This competitive influence demands managerial attention to competitor changes in quality, with several direct implications. First, if a competitor's quality increases this will increase the likelihood that consumers will switch to an **unstable learning model** and act more promiscuously regarding the firm's stable quality product. This unstable learning also creates an opportunity for the firm to respond with improvements in its own mean quality level. However, if quality improvements are not possible, it is especially important to deliver consistent quality to head off switching because of consumer promiscuity due to one bad experience. Second, even when competitor quality decreases, the firm could be hurt because consumers again are more likely to treat the focal firm's quality as changing. Hence, paradoxically, when competitor quality deteriorates, it is possible the firm might observe less brand loyal customers. Future research is needed to examine the extent and nature of these influences and how firms might manage these effects.

Our results also have potential implications for research at the intersection of choice models and consumer learning. A growing body of literature applies the **stable learning model** to myopic (e.g., Roberts and Urban 1988) and forward-looking consumers (e.g., Erdem and Keane 1996; Crawford and Shum 2005). Our results suggest consumers often appear to use an **unstable learning model** even when the quality level is stable. This result, in combination with the fact that in today's environment underlying quality levels of products and services are rapidly changing, seems to indicate that we should see increased attention to **unstable learning models** in the choice model setting.

Finally, our model should be useful to more deeply investigate influences on the selection of learning models. Future research might address a variety of questions in this domain. Are single, large disconfirming experiences enough to push consumer model selection towards an unstable model? Do firm communications play a role in influencing the learning models consumers select? Given more rapidly changing products, do consumers tend to use **unstable learning models** more now than they did in the past? Do older consumers that grew up in more stable environments tend to use **stable learning model**s more than younger consumers? If so, is brand loyalty more easily established in older than younger consumers? The world is changing, and with those changes come many opportunities to explore interesting, and managerially important, questions.

3. Essay Transition

In the first essay, we use experiments to elicit expectations and manipulate experiences. We find that consumers even when product quality is stable, act as if the product's quality levels are changing. This finding potentially has implications for Bayesian learning models that are applied to choice data. In the second essay, we explore these implications using simulated data.

In essay one, we adopt the notation found in the literature on quality expectations (e.g., Boulding et al. 1999). This notation differs from that found in the literature on choice data (e.g., Erdem 1998) and we adopt notation found in the choice literature for this second essay. As a result, the notation differs between the two essays. However, the fundamental unstable learning model remains the same.

4. Essay 2: Choice Models With Unstable Consumer Learning -- An MCMC Solution and Monte Carlo Experiment

4.1 Introduction

Consumers often think of products and services in term of underlying attributes such as quality, durability, and ease of use since these are the attributes that affect their interactions (experiences) with the product offering. In many instances these underlying attributes are not directly observable. Instead, consumers use their and others' experiences with the product to gradually learn the true levels of these attributes. This essay is concerned with a) modeling how consumers learn about an attribute level of a product offering that is not directly observable and b) developing a methodology for estimating the parameters of this learning process. Following the extant literature on belief updating (Erdem and Keane 1996), we characterize a person's belief in terms of a probability distribution and refer to the product offering attribute of interest as "quality." However, our model and discussion is applicable to any tangible or intangible attribute that is imperfectly observable.

Our exploration takes place within an environment in which consumers make choices among competing product offerings. These product offerings do not always deliver the same experiences; instead, from one purchase occasion to another the experienced quality differs. Since consumers are not able to directly observe the next experience prior to purchasing, they must make purchase decisions based on expectations of that next quality experience.

These quality expectations are formed based on past experiences and are influenced by two fundamental sources of uncertainty for consumers. First, consumers can be uncertain about the underlying quality of a product because their experiences (or other signals) are noisy, and thus do not reveal the true level of the attribute for the product. This variability around the true (mean) quality of the product can be due to perceptual error, service variation, or manufacturing variability.

Second, and our point of departure from the existing literature, consumers' uncertainty about a product can be associated with the belief that the product's true (mean) quality level changes over time. Such changes in objective quality have been documented by Mitra and Golder (2007) and can be attributed to changes in the underlying delivery process (manufacturing or service quality process) or variability in the context in which the delivery process occurs. For example, the service quality of airlines has decreased and increased over the years due to a range of contextual causes (e.g., security requirements) and firm actions (e.g., online ticketing, cost reductions, etc).

More importantly, regardless of whether the product is actually changing, if individuals attribute the variation in experienced quality levels as due to changes in true product quality, then they will react to this second source of uncertainty. For example we found in Essay 1 that the learning process for many consumers was better characterized by an unstable learning model even though the underlying quality was stable.

Such beliefs were also consistent with the results found in a number of experiments reported by Erev and Haruvy (2008). In that paper they describe a series of "clicking paradigm" experiments where decision makers learn over time based only on the information made available from the outcome of their last choice. Specifically these decision makers sit at a computer screen and choose (click on) one of two buttons. Each button has a payoff associated with it and the decision maker observes (and receives) the payoff associated with the chosen button. This process of choosing is done multiple times, in these cases either 200 or 400 times. The participants' goal is to learn about their environment so that they can choose the correct button. More specifically their goal is to maximize their total payoff over the trials.

In these experiments the payoff associated with a button could vary or be constant over these multiple experiences. However, one of the two buttons always has a higher expected payoff. Consequently, if the person learned fully about the underlying latent value of each button, the person's "optimal" action would be to always choose the button with the higher expected payoff. For example in one of the experiments reported by the two researchers, the high payoff button (H) gave a payoff of 21 with a probability of 50% and 1 otherwise, i.e. the expected payoff of H was 11. The other button (L) gave a payoff of 10 with certainty. In this particular case after about 20 trials the decision makers stabilized on the proportion of times they selected button H at .57. Table 8 presents the results of a number of different variations of this general setting.

Table 8: Experimental Results

Variance around True Mean		Difference in Expected Payoffs	Percent Selecting H
H	L		
1.6	0	0.2	61
9.6	0	0.2	27
1.6	1.6	0.05	51
357	17.7	0.35	50
10	0	1.0	57
0	9	1.0	65
0	0	1.0	87

Source: Erev and Haruvvy (2008)

As seen from this table the proportion of times button H was chosen decreased with the variability of the H payoff relative to the L button payoff and increased with the difference in the two expected payoffs. Such a finding is not too surprising and indicates that the participants were responsive to the different manipulations. However, what is surprising (at least if one believes consumers ultimately learn whether the world is stable or not) is that over these multiple trials the decision makers did not consistently choose the H button. This is particularly surprising for the condition where there was no variation in either button. Such behavior is consistent with the premise that participants felt that there was a chance that the underlying environment might change and thus they continued to sample the L button to see if its expected payoff had increased.

The above empirical results lead us to hypothesize that individuals tend to believe that there are two types of variability when they have repeated exposure to product offerings. The first is the variability around the underlying true quality level of the offering. In the click experiments this variability is due to the stochastic nature of the experiment, e.g. a 50% of one outcome and a 50% of a second outcome. In the service settings of Essay One, such as repeatedly visiting a fast food outlet, the variability might be due to variations in the abilities of the specific service provider, the time of day, or the presence or absence of other customers.

The second variability is associated with consumers' beliefs that the underlying true quality level of a product offering (including a gamble) can change over time. This belief causes an individual to interpret some of the variation in quality as shifts in the mean level of quality. For example, in Essay One, these changes could reflect a belief that wait times have increased or download speeds have decreased. Previous Bayesian models of consumer learning ignore this second type of uncertainty and we incorporate this possibility into a choice model.

Throughout, we assume that consumers act as Bayesians, conditional on their belief about the environment. Evidence suggests that most consumers are well-approximated by Bayesian rules, even though some consumers react too much or too little to new information (El-Gamal and Grether 1995). Our approach allows consumers to respond as Bayesian learners that have inaccurate beliefs about whether quality is changing or not. We show that a rational consumer with inaccurate beliefs should use a different learning model from the standard one found in the literature. With these inaccurate beliefs, consumers overreact to each new experience relative to the "rational expectations" Bayesian learning model (i.e., the stable learning model).

In the next section, we develop integrated choice and learning models based on a) this unstable assumption about consumers and b) the common assumption of a stable environment. In the third section, we describe the structural estimation of these models. In the fourth section, we discuss a Monte Carlo simulation study that demonstrates how well these models recover the true parameters when the data generating process reflects (does not reflect) the model assumptions. Specifically, we estimate the parameters of both consumer learning processes in simulated environments in which we vary whether the consumers actually learn assuming the underlying mean is stable or unstable when in fact the underlying mean is stable. Finally, we conclude the paper by discussing our results, their implications, and potential future research avenues.

4.2 Model

We develop a model of choice with consumer uncertainty and learning. Consider a market with consumers i=1, . . ., I, product offerings j=1, . . , J, and purchase opportunities t=1, . . . T. Consumers may be imperfectly informed about the quality of the product offering and thus use these interactions with the product offerings to obtain fallible signals about the products' true attribute levels.[1] We define this imperfect use experience process as

$$X_{ijt} = A_{jt} + x_{ijt} \quad (1)$$

where X_{ijt} is the quality level consumer i perceives as receiving from product offering j at time t assuming the consumer interacts with the product at time t, A_{jt} is the true quality level for product offering j at time t, and x_{ijt} is an i.i.d. random error term associated with the variability in the delivery process and/or in the consumer's perceptual errors. More specifically, following the literature on Bayesian learning models, we assume the random errors x_{ijt} are distributed normally with constant variance that can differ across product offerings. Hence,

$$x_{ijt} \sim N(0, \sigma^2_{xj}) \quad (2)$$

where σ^2_{xj} is the variance of use experiences for product j. We assume that these stochastic elements are i.i.d. across consumers, product offerings, and time. The experience variability captures how noisy the use experiences are as signals of the latent attribute with higher variability indicating noisier signals.

4.2.1 Unstable and Stable Environments

The existing literature assumes $A_{jt} = A_j$, that is, the true attribute level is not changing. We refer to the environment that arises from this assumption as a **stable**

[1] We could equivalently describe the experiences X_{ijt} as outcomes of a stochastic quality process. In this conceptualization, the A_{jt} are the time-product specific means of the stochastic quality process, i.e., they represent the mean quality level of the process. With this noted, we henceforth refer to A_{jt} as the "true" quality or attribute levels for simplicity.

environment. If the true attribute level changes over time, we use a t subscript on A_{jt}[2] and refer to this changing environment as an **unstable environment**.

To capture the time varying true (mean) quality in a relatively unrestrictive way, we assume the future changes are equally likely to be positive or negative. Specifically, we assume

$$A_{jt} = A_{jt-1} + a_{jt} \tag{3}$$

where a_{jt} are mean zero i.i.d. stochastic shocks to the true quality levels. This shock (like the A_{jt}'s) is unobserved to the consumer and can be caused by a range of firm or environmental causes. For example, the firm may cut costs or staff affecting the service quality process, or the firm may change the ingredients or manufacturing process. Alternatively, something in the environment may change, altering the true quality levels such as the introduction of a new highway that increases how crowded a store gets or a new security protocol that influences airline check-in processes.

We build upon a well-developed Bayesian forecasting literature, by assuming the stochastic terms a_{jt} are normally distributed.

$$a_{jt} \sim N(0, \sigma^2_{Uj}) \tag{4}$$

2 Erdem et al (2005) assume that relative quality is fixed, but absolute quality could be adjusting.

where σ^2_{Uj} is the variance of the stochastic changes in the quality levels A_{jt}. It is the presence or absence of σ^2_{Uj} that determines whether the environment is stable or unstable (i.e., if $\sigma^2_{Uj} = 0$, then the environment is stable; when $\sigma^2_{Uj} > 0$, it is unstable).

Though we introduce the unstable environment, we will assume that the environment is stable for the remainder of the paper (i.e., $A_{jt}=A_j$) and drop the t subscript on A. Our intent in introducing the time-varying quality levels is to motivate that consumers **might act as if** the environment is unstable.

4.2.2 Unstable and Stable Consumer Learning Models

We next present learning models that representing optimal learning based on the consumer assuming the quality level is stable or unstable. We refer to consumer learning models in which consumers assume that the true quality level is fixed (i.e., that they are in a stable environment) as **stable learning models**. We refer to learning models in which consumers believe true quality is changing (i.e., that they are in an unstable environment) as **unstable learning models**. Thus, the defining quantity for whether a learning model is stable or unstable is the consumer's assumption about the unstable variance, σ^2_{Uj}. More specifically, if the consumer assumes $\sigma^2_{Uj} = 0$, then she uses a stable learning model, and if the consumer assumes $\sigma^2_{Uj} > 0$ then she uses an unstable learning model. To be clear, the difference between the stable and unstable learning models is in the consumer's assumption about the unstable variance, which is not necessarily reflected in the actual environment.

Note that by taking this approach we do not assume consumers have rational expectations about whether quality is changing or not. Instead, even though we will always assume that the true quality is stable (i.e., $A_{jt} = A_j$), consumers could act as if the environment is unstable and use an unstable learning model. In Essay One, we present empirical evidence and a rational model that suggests consumers may behave as though the environment is unstable when in fact it is stable. We also earlier discussed other experiments that are compatible with the premise that consumers may not have rational expectations, but instead may be characterized by beliefs about the environment that do not match actual conditions. In particular, consumers may believe the outcomes of a product or option are changing when, in fact, they are constant. This mismatch suggests that learning models may benefit from incorporating assumptions that are counter to the actual state of the environment. In our Monte Carlo experiment, we manipulate the environment and the consumer learning model in order to determine how a mismatch would impact the recovery of the structural estimates. This investigation is important, because if, as we show in essay one, consumers use inaccurate learning models, then such mismatches are likely to occur.

Before proceeding to the beliefs and updating process, we note two important quantities that we assume the consumer knows: σ^2_{Uj} and σ^2_{Xj}. The assumption that the variance of experiences, σ^2_{Xj}, is known is found in virtually every application of Bayesian learning models to choice data. Thus, our further assumption of a known σ^2_{Uj} is a direct extension of this existing assumption. In the technical appendix, we show why this assumption is relatively innocuous for our application.

4.2.2.1 Consumer Beliefs and the Updating Process

After each experience with a product offering, we assume consumers update their prior beliefs about the true quality levels through Bayesian learning. Previous literature suggests that the Bayesian updating mechanism provides a reasonable approximation to observed choice (Erdem 1998; Erdem and Keane 1996; Roberts and Urban 1988) and belief behaviors (Rust et al. 1999; Boulding et al 1999). Following this literature, we assume that each consumer's prior belief about the quality levels of product offering j is initially normally distributed. Coupling this assumption with the assumption that the use experiences are normally distributed implies that all subsequent consumer beliefs are also normally distributed. Throughout the following development, we will assume the consumer believes the environment is unstable and thus uses an unstable learning model. This is the most general of our models and the application to a stable environment or learning model is direct.

We start our model development by noting consumers hold four different beliefs during the updating process. The first is the consumer's posterior belief, about the quality level at t-1 held at time t-1, i.e., before the individual has a product experience at time t, but after updating based on the experience at time t-1. This belief is denoted

Posterior Belief: $\pi_{it-1}(A_{jt-1}) = N(E_{it-1}(A_{jt-1}), \sigma^2_{Aijt-1})$ (5a) where E_{it-1} (--) is consumer i's best (mean) guess about the quality level for product offering j at time t-1 and σ^2_{Aijt-1} captures the person's uncertainty about this underlying true quality.

The second represents the consumer's belief of the quality level next period, prior to an experience in that period. The consumer assumes that in the next period, the quality level may have changed ($A_{jt} = A_{jt-1} + a_{jt}$). In the consumer's model, these changes (a_{jt}) are mean zero, so that they do not affect the expectation of the belief, i.e., $E_{it-1}(A_{jt}) = E_{it-1}(A_{jt-1})$. However, the variance of the forward-looking belief will increase by σ^2_{Uj}, reflecting the consumer's assumption of unstable variance:

Prior Belief: $\pi_{it-1}(A_{it}) = N(E_{it-1}(A_{jt}), \sigma^2_{Aijt-1} + \sigma^2_{Uj})$ (5b)

The third belief represents the consumer's forecasted experience, X_{ijt} and is the belief the consumer uses to make his or her next purchase. The expectation of this forecast belief does not differ from the mean of the prior belief $\pi_{it-1}(A_{it})$. To see this, note that $X_{ijt} = A_{jt-1} + a_{jt} + x_{ijt}$ and that like the a_{jt}, the x_{ijt} are mean zero, so that they do not affect the expectation of the belief. Hence, the expectation of the forecast is equal to the consumer's posterior expectation for the product offering's quality level. However, the variance of the forecast belief increases from the variance of the prior belief by the variance of the experiences, σ^2_{Xj}. Thus, $Var_{it-1}[X_{ijt}] = Var_{it-1}[A_{jt-1} + a_{jt} + x_{ijt}] = \sigma^2_{Aijt-1} + \sigma^2_{Uj} + \sigma^2_{Xj}$. As a result, the uncertainty of the forecast distribution of the next experience is greater than the variance of the t-1 posterior or prior beliefs of A_{jt}. Hence, Forecast of Experience: $\pi_{it-1}(X_{ijt}) = N(E_{it-1}(A_{jt}), \sigma^2_{Aijt-1} + \sigma^2_{Xj} + \sigma^2_{Uj})$ (5c)

The final belief of interest is the consumer's updated (posterior) belief after having a product offering experience at time t. We denote this belief as follows:

Posterior Belief: $\pi_{it}(A_{jt}) = N(E_{it}(A_{jt}), \sigma^2_{Aijt})$ (6)

This posterior belief will initiate a new set of prior and forecast beliefs for period t.

The difference between the posterior belief at time t-1 (5a) and the posterior at time t (6) captures the belief updating process. This difference is due to two sources of change. First, as represented in the prior belief (5b), looking ahead to period t increases the variance of the belief due to potential changes in the true quality. Second, as represented in (6), a new experience provides information about this new quality level that informs the uncertainty and expectation of the new belief. The link between these two means and variances can be directly specified using either multivariate normal theory or the Kalman filter (West and Harrison 1997).

Before discussing these linkages, we note the variances contained in (5a-c) differ for the stable learning model. Because there no $\sigma^2_{Uj} = 0$ in the stable model, equation (5a) and (5b) are identical, so that the initial posterior is the prior. Further, (5b) and (5c) do not contain the σ^2_{Uj}.

4.2.2.2 The Updating Equations

If the consumer does not have an experience with product j in time period t then the new and previous posterior beliefs have the same expectation (i.e., $E_{it}(A_{jt})= E_{it-1}(A_{jt-1})$. This might occur, for example, when the person has an experience with some other product k at time t. However, if the consumer experiences product j at time t, the new and old expectation will differ (i.e., $E_{it}[A_{jt}] \neq E_{it-1}[A_{jt-1}]$). The update in expectations is represented as follows:

$$E_{it}[A_{jt}] = E_{it-1}[A_{jt}] + D_{ijt}\, \beta_{ijt}\, (X_{ijt} - E_{it-1}[A_{jt}]) \tag{7}$$

where β_{ijt} is a Kalman gain coefficient obtained from the Kalman filtering algorithm, and is bounded between 0 and 1 (i.e., the new expectation is a convex combination of the previous expectation and the new experience), and D_{ijt} is an indicator variable set to 1 when the consumer encounters a product offering experience and zero otherwise. We discuss these gain coefficients in the following subsection, but note that the Kalman gain coefficient differs depending on whether the learning model is stable or unstable.

Similarly, the change in the variance of the posterior belief depends on whether the consumer has an experience with the product. As shown in the appendix, the consumer should use equation (8b) to update her uncertainty if she did not have an experience with the product and equation (8a) if she did.

$$\sigma^2_{Aijt} = (1- \beta_{ijt}) (\sigma^2_{Aijt-1} + \sigma^2_{Uj}) = \beta_{ijt} \sigma^2_{Xj} \quad (8a)$$

[3]

$$\sigma^2_{Aijt} = \sigma^2_{Aijt-1} + \sigma^2_{Uj} \quad (8b)$$

where β_{ijt} is the same Kalman gain coefficient as found in (7). Note that the equations are the same for both the stable and unstable model, but that in the stable learning model the $\sigma^2_{Uj}=0$ so that the variance only changes if the consumer has an experience. Also, note that if the consumer has a product experience (i.e., gains new information) the variance normally will decrease[4]. In contrast, without new information the variance always increases. The new belief doesn't have new information to adjust the prior (5b), so the variance of the new posterior is equal to that of the prior. This increase arises because the consumer recognizes that the quality level might have changed since the last period. In other words the unobserved change in the true quality levels makes the old belief about quality less relevant to the next experience.

4.2.2.3 Kalman Gain Coefficients in Stable and Unstable Learning

In the stable learning model (i.e., where consumers believe $\sigma^2_{Uj}=0$), the Kalman gain coefficient is

$$\beta_{ijt} = \sigma^2_{Aijt-1} / (\sigma^2_{Aijt-1} + \sigma^2_{Xj}). \quad (9)$$

In contrast with unstable learning, (i.e., where consumers believe $\sigma^2_{Uj} > 0$) the Kalman gain coefficient is

---

[3] Note that $\beta_{ijt} \sigma^2_{Xj} = (1- \beta_{ijt}) \sigma^2_{Aijt}$ (the form that has appeared in most Bayesian learning models) when the mean quality levels are constant (i.e., $A_{jt}=A_j$).
[4] In the unstable learning model this is the case for sufficiently large initial variance σ^2_{Aij0}. In the stable learning model, this is always true.

$$\beta_{ijt} = (\sigma^2_{Aijt-1} + \sigma^2_{Uj}) / (\sigma^2_{Aijt-1} + \sigma^2_{Xj} + \sigma^2_{Uj}). \tag{10}$$

The only difference between these formulations is the addition of σ^2_{Uj} in both the numerator and denominator. This addition will cause the consumer to have a fundamentally different response to new and old information over time. To clarify the roles of the different parameters, we simplify notation (dropping i and j and inserting S or U for the stable or unstable learning models, respectively) and re-parameterize as follows:

$$\beta_{0(S)} = \sigma^2_0 / (\sigma^2_0 + \sigma^2_X)$$

$$\beta_{0(U)} = (\sigma^2_0 + \sigma^2_U) / (\sigma^2_0 + \sigma^2_U + \sigma^2_X),$$

$$\varrho = \sigma^2_U / \sigma^2_X.$$

Then, when the consumer has an experience

$$\beta_t = \beta_{t-1} / (\beta_{t-1} + 1) \tag{9a}$$

$$\beta_t = (\beta_{t-1} + \varrho) / (\beta_{t-1} + \varrho + 1). \tag{10a}$$

Equations (9a) and (10a) make clear that just two quantities can characterize the transitions with an experience. In the unstable case, when the consumer does not have an experience, the "unobserved" Kalman gain coefficient, β_{t-1}, increases by ϱ. In the stable model, without an experience the coefficient doesn't change. Thus, the entire path of Kalman gain coefficients can be characterized by these two quantities. It is instructive to discuss the values of these Kalman gain coefficients with repeated experiences to make clear the difference between an unstable learning process and a stable learning process.

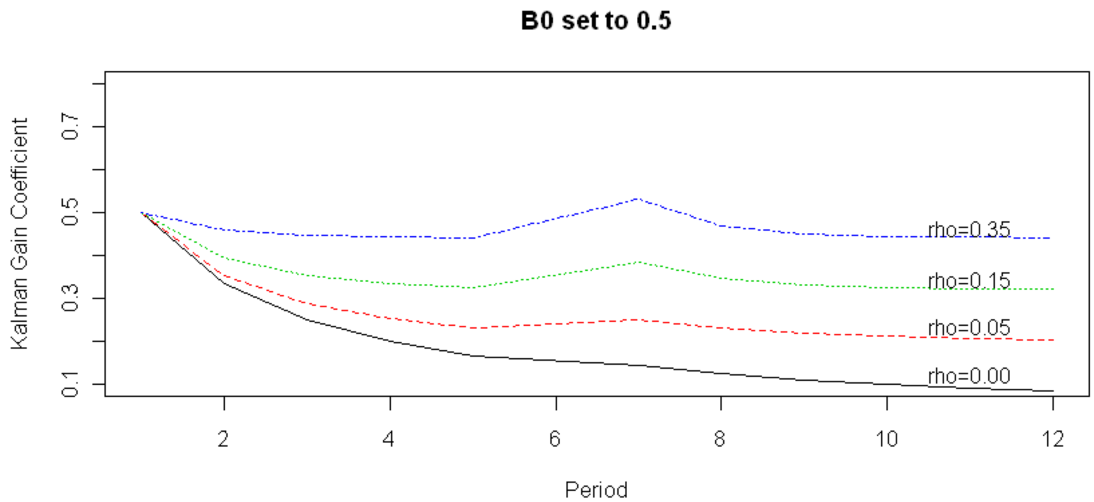

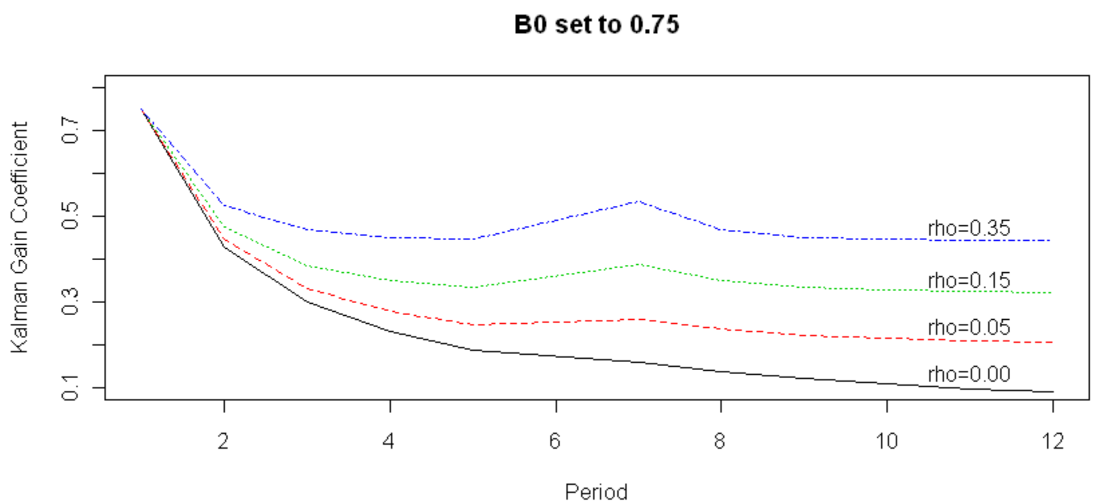

Figure 2: Kalman Gain Coefficients

Figure 2 depicts β_{ijt}, assuming an experience in each period, for t=1, . . ., 5 followed by a period with no experience, and another six periods of experiences. In the first panel, the β_0 is 1/2 and in the second panel β_0 is 3/4. In both panels the ϱ takes on values 0 (i.e., stable learning), .05, .15, and .35 as indicated on the curves. We note three fundamental properties of these curves. First, the β_0 only has a notable influence on the first few coefficients. Second, after several experiences the unstable and stable learning models diverge as the β_{ijt} under stable learning continues to decrease, while those of the unstable learning cases rapidly approach a positive asymptotic value. The level of that asymptote is determined by the value of ϱ. Third, with values of ϱ greater than zero, the coefficient actually increases between periods 5 and 7 due to not having an experience in period 6, as expressed in equation (8b).

The β_{ijt}'s influence both how much the expectation shifts from one time to the next and how much the variance decreases over time. With greater β_{ijt}'s the consumer learns more from each new experience, but retains more uncertainty about that belief. Hence, these β_{ijt}'s can be viewed as representing the learning process with greater values implying more learning from the experience (and still more uncertainty about the true state of the world).

4.2.3 Learning Models and Brand Equity

Following existing literature, (e,g, Erdem et al. 2004, Keller 1993), a consumer's belief structure can be viewed as a measure of "brand equity." Brands with higher equity will have either higher expected attribute levels and/or lower perceived risk, ceteris paribus. We do not propose this as a full measure of brand equity, but instead, following existing literature (e.g., Erdem 1998), offer it as a way to examine the influence of the learning model on some types of brand equity. We explore how these measures differ in the two different learning models.

First, consider the reaction of consumers to a prior bad experience. As can be seen in (7) the form of the consumer's mean belief is the same across the learning models. However, the reaction to the same experience differs. Specifically, as shown above, for a fixed initial uncertainty and experience variability, the β_{ijt} will be higher under the unstable learning model than the stable learning model. As a result, consumers using an unstable learning model will react more to the most recent experience and weigh the past less than will consumers using the stable model. This difference persists and even grows with more experiences.

Second, when a consumer does not have an experience, in the unstable learning model, that consumer will react to future experiences more, meaning that periods of no purchase erode brand equity. In contrast, stable consumers are unaffected by this lack of experience, implying that no purchase has no effect on the expected qualities and brand equity. Thus, the learning model that consumers use affects both how much brand equity can be accumulated and how long that equity persists.

4.2.4 Consumer expected utilities

We assume U_{ijt}, the utility consumer i perceives at time t that she will get from purchasing brand j on the next occasion, depends on the perceived attribute level for the experience X_{ijt+1} and price P_{ijt}[5]. We specify a flexible linear utility formulation.

---

[5] It is possible to include other observable attributes. We do not do so in order to concentrate our attention on the learning model.

$$U_{ijt} = \tau\, E_{it}[X_{ijt+1}] + \alpha\, P_{ijt} + \varepsilon_{ijt}, \qquad (11)$$

where α and τ are the price sensitivity and utility weight parameters assumed to be common across consumers.[6] The variable ε_{ijt} is a stochastic component specific to a purchase decision and captures random, unobservable (to the researcher) preference shocks that are known by the household at the time of purchase.

Since the consumer does not know the quality level of the experience she will have at time t+1 given the choice of the product offering at time t, the consumer must base the purchase upon the expected experience. From equation (5c), we have $E_{it}[X_{ijt+1}] = E_{it}[A_{jt}]$. We note that this formulation does not differ between the stable and unstable learning models. The two models only differ in the changes in expectations due to product experiences or lack thereof.

4.3 Data and Structural Estimation

Thus far we have put forth a model of how consumers should react to a new experience when their goal is to learn about the underlying quality level of a firm's product offering. We next turn our attention on how one might estimate this learning model from observed data. In this section, we discuss the data needed for this estimation and the estimation model we use to recover structural parameters of the consumer's learning model.

---

[6] Note that it is a fairly direct extension in theory to incorporate heterogeneity in these parameters. We attempt to maintain a parsimonious model in order to focus attention on the central learning models.

4.3.1 Simulated Data

Hypothetically, at least three different information sets could be used to estimate our model of consumer learning. The first uses data similar to that found in essay 1 where the consumer provides expectation data, $E_{it-1}[X_{ijt}]$, prior to each purchase occasion. Such data are best collected in an experimental setting and allow for a much more nuanced modeling of the process consumers use to learn. However, with expectations data alone the connection to choice and marketing actions is less direct. More importantly, panels of expectation data for individuals are rarely available (for a somewhat different use of panels of belief data, see Chintagunta et al. 2007). The second approach, and the one most often found in the marketing literature (e.g. Erdem and Keane 1996; Erdem 1998), is to collect purchase behavior data along with pricing data and infer the consumers' quality expectations from their purchases. Although such data are readily available, without any measures of expectations, identification of all the model parameters is not possible without simplifying assumptions (Shin, Misra, and Horsky 2007).

The third approach, and the one taken here, assumes the researcher has the same panel data on choices and prices as mentioned in the second approach but supplements these data with a single period of self-explicated quality expectations. To date the only researchers to use such data are Shin, Misra, and Horsky (2007). With this rarity noted we believe it should be relatively easy for managers to gain a single time slice of expectation data for a panel of individuals, thereby allowing one to marry the benefits of each approach. Further, we note that many applications use an initialization period to set priors for consumers (e.g., Mehta et al. 2004). These initialization periods generate similar information for the model to that involved in our use of a single period of expectation data (i.e., an informative prior about each consumer's initial expectations). Thus, our use of the initial expectation data is not conceptually divergent from existing practice.

Before proceeding to the estimation models, we note the implications of this dataset for what is directly observable. First, beyond the first (initialization) period, we will not observe any consumer expectations or experiences. Thus, the $E_{it}(A_j)$ and the X_{ijt} values are unobserved. Second, we observe choices, which provide information about the expected utilities. It is through these indirectly observed expected utilities that we must recover the expectations and experiences over time. In section 4.3.3 we discuss the identification of these parameters from the data.

4.3.2 Structural Estimation Models

Past research on Bayesian learning models has tended to use simulated maximum likelihood (SML) to recover the structural parameters. We develop an estimation approach using Bayesian methods. This approach both allows straightforward estimation of consumers' expectations and provides a valuable, yet tangential outcome—a direct, functional link between Bayesian learning models and a class of time-series models, the dynamic linear model (DLM) (see West and Harrison 1997), which has been applied in a number of applications in marketing. Finally, we discuss identification for our model.

4.3.2.1 Estimation Models Under Different Assumptions

In the model development section, we discussed two types of consumer learning models (unstable and stable) and two types of potential environments (unstable and stable). As a researcher, we could assume any of these environments and consumer learning models. This produces a 2x2 of possible models to estimate as depicted in Table 9. Each of the cells is a label we will apply to succinctly describe the assumptions of the estimation model (mod). For instance, SSmod assumes that the environment is stable and the consumers' learning models are stable as well.

Table 9: Environment and Consumer Learning Models

	Consumer Learning Models	
Environment	**Stable**	**Unstable**
Stable	SS	SU
Unstable	US	UU

We will focus our attention on models SSmod and SUmod. We focus here for two reasons. First, the cases with an unstable environment were found to not be well-identified from the types of data typically available in choice applications. Second, in most applications the claim has been that quality is, in fact, stable in the timeframes examined. Thus, for most applications, quality is likely not changing, even though consumers may act as though it is. Thus, we explore the situations in which quality is in fact stable, but for which consumers may act as if it is unstable.

We now describe the consumer initial beliefs and the reparameterizations and normalizations used in the SSmod and SUmod models.

4.3.2.2 Consumer Initial Beliefs

We need to specify for each consumer her initial belief, i.e. we must specify an expectation and variance. We use the single set of expectation data, WE_{ij}, to directly identify the initial expectation of each consumer's belief. We assume a survey instrument is used to collect the self-explicated data. We assume that this survey produces an unbiased estimate of the consumer's true expectation. We also assume the variance associate with this survey estimate is known.[7] Further, we assume that the measurement errors are normally distributed, so that the researcher has a prior for each consumer's true expectation, $E_{i0}[A_j]$, that is normally distributed as follows:

$$p_{ij}(E_{i0}[A_{j0}]) \sim N(WE_{ij}, \sigma^2_{ME}) \qquad (12)$$

where σ^2_{ME} is the known measurement error variance. We note that, in most choice models that incorporate learning, some initialization is done to set the initial information about consumer's expectations. Thus, our initialization is quite similar in function to current practice.

[7] The extension to make this variance uncertain is straightforward, but diverts attention from the central components of the learning model and requires additional expectation data per individual.

We assume all consumers share a common variance associated with a given product. That is, the initial variance of the belief about quality is assumed constant across individuals:

$$\sigma^2_{Aij0} = \sigma^2_{0j} \qquad (13)$$

These shared initial variances become heterogeneous as soon as the individuals make different product decisions, which then affect the variance of the belief.

4.3.2.3 Reparameterization and Normalizations

We make several simplifying assumptions and normalizations to facilitate identification. First, as noted above, we assume that consumers' uncertainty is homogenous initially, that is $\sigma^2_{Aij0} = \sigma^2_{0j}$. Second, we assume that the non-purchase option represents an outside good, which we normalize to a constant zero utility over time. Third, we fix the value of τ in equation (12) to 1. This ensures identification of the scale of the unobserved quality. Fourth, we assume that the variance of the experiences, $\sigma^2_{Xj} = \sigma^2_X$, is the same across products in the category. Fifth we assume that the unstable variance, $\sigma^2_{Uj} = \sigma^2_U$, is the same across products in the category. Together assumptions four and five imply that consumers use a single learning rule for all products within a category. Finally, to aid in the mixing of the sampler, we conduct two reparameterizations that were introduced in section 4.2.2.3. Specifically, we set $\rho = \sigma^2_U/\sigma^2_X$ and $\beta_{0j} = (\sigma^2_{0j} + \sigma^2_U)/(\sigma^2_{0j} + \sigma^2_U + \sigma^2_X)$.

4.3.2.4 The Likelihood and Prior

As mentioned earlier, our observed data to be used as the dependent variable are choices. Thus we need to use a choice model to infer the consumer's underlying utility model parameters. From equation (12) we see that the choice likelihood depends on the distributional assumption for **ε**, the vector of idiosyncratic, time-varying preference shocks. Two standard assumptions are that these errors are i.i.d. extreme value and i.i.d. normal. The former implies the probability of consumer i choosing brand j at time t takes the form of a multinomial logit choice probability (MacFadden 1974) and the latter leads to a multinomial probit formulation which can be estimated using data augmentation (McCulloch and Rossi 1994). We use the latter assumption, which facilitates conditional normal analysis.

To the researcher, the likelihood of consumer i purchasing product offering j at time t, q_{ijt}, is conditional on the parameters α, $\sigma^2{}_X$, ϱ, $\{A_j\}_j$,$\{\beta_{0j}\}_j$, and the expectations, $\{E_{it}[A_j]\}_{ijt}$. We have already shown in equations (9a) and (10a) that β_{ijt} is a deterministic function of β_{0j}, ϱ, and $\mathbf{D}_i$, where $\mathbf{D}_i$ is the past purchase history for consumer i. Thus, with these β_{ijt} and equation (7), one can construct the conditional distributions of $E_{it}[A_j] \mid E_{it-1}[A_j]$. These distributions are conditional on the parameters $\sigma^2{}_X$ and A_j and the prior distribution of the initial expectation, p_{ij}.

We now specify the priors for the structural parameters. We use diffuse conjugate normal priors for the price sensitivity, α, and the true qualities, A_j. For the variance of the experiences we use an improper "reference prior", σ^2_X. For the initial learning rate, β_{0j}, we use a uniform(0,1) prior. Finally, for ϱ we use a uniform(0,1000) prior. While the upper limit is artificial, in practice this upper limit is never approached, and the prior is not binding on that side.

For posterior inference, we use a block Gibbs sampler with conditionally conjugate draws for all but the parameters that form the updating coefficients, β_{ijt}, for which Metropolis-Hastings steps are necessary because of the non-linear form. We describe the sampler in the technical appendix.

4.3.3 Mapping between BLM and DLM with linear utility

We begin by defining a general dynamic linear model (DLM) system. Based on a slight modification of the univariate DLM of West and Harrison (1997, p. 102), the DLM equations are defined as follows:

$$Y_t = F_t' \theta_t + g X_t + v_t,\ v_t \sim N(0, V) \quad (14)$$
$$\theta_t = G_t \theta_{t-1} + h Z_t + w_t,\ w_t \sim N(0, W_t) \quad (15)$$
$$(\theta_0 \mid D_0) \sim N(m_0, C_0) \quad (16)$$

This system is defined by the 'tuple at each time period $\{F_t, g, h, V, G_t, E_t, W\}_t$. We will ignore the third equation in what follows, but its role is easy to establish once the mapping is clear.

To achieve a mapping between the Bayesian learning model (BLM) and the dynamic linear model (DLM) requires that a linear utility formulation is used (i.e., no quadratic term). We further do not include heterogeneity in the price sensitivities or uncertainty weights to simplify notation, and note that incorporating heterogeneity does not alter the fundamental mapping. We restate as reference the key simplified equations here for the unstable learning model:

$$E_{it-1}[U_{ijt}] = \tau\, E_{it}[X_{ijt}] + \alpha\, P_{ijt} + \varepsilon_{ijt} \tag{17}$$
$$E_{it}[X_{ijt}] = (1 - D_{ijt}\beta_{ijt})\, E_{it}[X_{ijt-1}] + \beta_{ijt}\, D_{ijt}\, (A_j + x_{ijt}) \tag{18}$$

First, we note that the stable learning model is a simplification of this system of equations and, thus, is automatically mapped into a DLM if this system is. Second, note the following terms are stochastic and normally distributed with mean zero: ε_{ijt} and x_{ijt}. Third, note the mapping is as follows: 1) Equation (14) maps to Equation (17) and 2) Equation (15) maps to Equation (18). Now, we present the matrix representation of this mapping, using as shorthand $z_{ijt} = E_{it}[X_{ijt}] = E_{it}[A_{ijt}]$.

$$Y_t = E_{it-1}[U_{ijt}]$$
$$V = I$$
$$\theta_t = (z_{11t}, z_{12t}, z_{21t}, z_{22t}, \ldots, z_{N1t}, z_{N2t})'$$
$$\mathrm{diag}(F_t) = (\tau, \tau, \ldots, \tau)$$
$$\mathrm{diag}(g) = (\alpha, \alpha, \ldots, \alpha)$$
$$h = (A_{1t}, A_{2t}, A_{1t}, A_{2t}, \ldots);$$
$$\mathrm{diag}(Z_t) = (D_{11t}\beta_{11t}, D_{12t}\beta_{12t}, \ldots, D_{N1t}\beta_{N1t}, D_{N2t}\beta_{N2t})$$
$$X_t = (P_{11t}, P_{12t}, P_{21t}, P_{22t}, \ldots, P_{N1t}, P_{N2t})'$$
$$\mathrm{diag}(G_t) = ((1-D_{11t}\beta_{11t}), (1-D_{12t}\beta_{12t}), (1-D_{21t}\beta_{21t}), \ldots, (1-D_{N1t}\beta_{N1t}), (1-D_{N1t}\beta_{N2t}))$$
$$\mathrm{diag}(W_t) = (D_{11t}\beta_{11t}, D_{12t}\beta_{12t}, \ldots, D_{N1t}\beta_{N1t}, D_{N2t}\beta_{N2t})^* \sigma^2_X$$

where diag(*) indicates the diagonal of a diagonal matrix. Thus, the BLM does map into the DLM, but to a very specific one. It has a time-varying system transition matrix that changes according to the estimated values of $\beta_{ijt}(\beta_{0j}, \varrho)$. Further, it has a very specific time-varying system transition variance matrix, W_t, that is a function of $\beta_{ijt}(\beta_{0j}, \varrho)$ and

σ^2_X. Thus, the theory about how consumers learn refines the model assumptions about the random intercepts. These constraints are so specific that the BLM is not a submodel of perhaps the most widely known DLM in marketing (see van Heerde, Mela, and Manchanda 2004). That DLM is a submodel of the broader DLM framework because it employs constant system variance matrix, W, and constant system transitions, G.

We also note that the fundamental mapping is not affected by including consumer uncertainty or quadratic utility of quality in the expected utility. However, in the quadratic case, the mapping is no longer linear and the DLM is not the relevant comparison. Instead, the broader framework of a particle filter replaces the Kalman filter as the core linking mechanism in order to accommodate the non-linear relationship in the system equation and a broader dynamic non-linear framework is necessary (see, for example, Godsill, Doucet, and West 2004).

4.3.4 Identification

In this section, we describe what information in the data helps to identify the parameters of the SSmod and SUmod models. Before we begin, however, we first note several design decisions that we took in order to avoid potential identification problems. First, we have included measures of initial expectations. Consequently, unobserved initial cross-sectional variation in expectations cannot explain any overestimation of learning. Second, we have not included heterogeneity in pricing and quality weights in order to avoid potential unobserved heterogeneity/state dependence misidentification.

In our Monte Carlo study, we will restrict our attention to data generating processes that do not contain this particular form of unobserved heterogeneity. With this noted, we now turn to the identification of our parameters.

First, the identification of price sensitivities is straightforward from choices conditional on relative prices. Thus, so long as sufficient variation in choices and relative prices across the different options is present in the data, the price sensitivities should be identified. Note this would also be true for additional observable time-varying covariates, such as display.

Turning to the parameters related to learning, identification occurs less directly and ultimately all of the parameters are linked together. However, several parameters have distinct sources of identification. First, recognize that the A_j's are the means of the unobserved experiences, and, as a result, the means of the expected utility intercept. In the stable learning model, the A_j's determine the average market shares after a large number of consumer experiences with the product. Thus, the A_j's can be identified from long run average market shares across consumers. However, in the unstable learning model, variation around the A_j's persist due to the positive convergent values of β_{ijt}. Still the intercept of quality expectations, since the quality is in fact stable, will be centered at the true quality. Consequently, the A_j's in the unstable model can be identified after a relatively small number of consumer experiences with the product from the mean of average market shares, where the mean is taken over multiple periods.

Next, consider first the stable learning model, in which ϱ is not present. In this case, σ^2_X and β_{0j} jointly define the path from the initial expected value of consumer expectations to the final distribution of expectations across consumers that is centered at the A_j's. These transitions are only parametrically defined. Thus, the recovery of σ^2_X and the β_{0j}'s from the observable data is essentially a curve-fitting process. The β_{0j}'s heavily influence the first few expectation changes (the early impact on choice conditional on past purchases), while the σ^2_X largely determines for how long past purchases have an observable impact on choice.

Now consider the unstable model. Note that ϱ, which is only present in the unstable learning model, has a distinct effect from the β_{0j}'s and σ^2_X when the consumer does not have an experience with the product. Without an experience the consumer's uncertainty changes without changing the expectation so that each non-purchase occasion for product j increases the change in expectations when a purchase finally occurs. In other words, the consumer's weight on the new experience differs depending on the consumer's current uncertainty, and this uncertainty depends on ϱ. This increase is uniform across consumers, products, and missed purchase occasions, so that with a sufficient pool of these occasions the effect of ϱ on expectations can be distinctly identified.

4.3.5 Summary

In this section we have defined and elaborated on the versions of an estimation model that allow stable and unstable consumer learning models as applied to choice data. We have discussed the SSmod and SUmod models. In the following section, we discuss the convergence behavior of the SSmod and SUmod samplers as well as their ability to recover the true parameters of the environment for which they were intended.

4.4 Monte Carlo Study

Having specified the model and estimation process, we turn to applying the model to simulated data. The goal of these simulations is two-fold. First, we provide evidence about the identification of parameters for both the SUmod and SSmod. Second, we explore the potential bias that arises from making incorrect assumptions about the learning model that consumers use. Specifically, we show Monte Carlo results for a 2 x 2 experiment in a stable environment in which consumers actually use either a stable or unstable learning model, and in which the researcher assumes consumers either use a stable or unstable learning model.

This section proceeds as follows. First, we describe the design and measures we use to analyze the Monte Carlo study. Second, we demonstrate the convergence of the samplers used to simulate the posterior distribution of the parameters. Third, we present the results of the Monte Carlo study.

4.4.1 Monte Carlo Study Design and Measures

The Monte Carlo study examines the performance of the two models, SSmod and SUmod in two different conditions--when consumers actually act as if the world is stable or unstable. We label these conditions as Scond and Ucond respectively. In both cases, the actual quality of the products is stable. However, in the Ucond case, consumers defy this truth and act as if the products are unstable. Because in all cases the actual quality is stable, we drop the leading S and use only U or S to indicate the assumption about how consumers behave.

Thus, we consider a 2x2 cross of the model and condition, where the condition determines how the data are generated (i.e. the quality levels, consumer experiences, consumer learning from these experiences and consumer choices based on this learning) and the model determines which assumption about consumers is contained in the estimation learning model. In two cells of the study, the model and condition match, whereas in two of the cells the model and condition do not match (see Table 10).

Table 10: Monte Carlo Experimental Design

Actual Consumer	**Model Assumptions**	
	Stable	**Unstable**
Stable	Scond, Smod	Scond, Umod
Unstable	Ucond, Smod	Ucond, Umod

We draw a number of simulated data sets that differ due to both stochastic elements and due to the specific parameters chosen. Each dataset assumes there are two products within the product category and 200 consumers making choices every period—where choice can be no choice. These simulated datasets have three stochastic components. More specifically, each experience has an unobserved (to the researcher) random shock, each choice has unobserved (to the researcher) utility shocks, and each consumer has initial random expectations about the product qualities. Consumers' initial expectations are drawn from a normal distribution with variance of .75 in all cases. The center of this distribution depends on the dataset as explained below. The random quantities make each dataset unique in the observed choices and initial expectations. In addition, this randomness leads to heterogeneity across consumers in the unobserved Kalman gain coefficients and quality expectations.

Table 11: Parameter sets for Monte Carlo study[1]

Parameter	Set 1	Set 2	Set 3[2]
A_1	(1.50), 1.35, 1.65	(1.75), 1.60, 1.90	(1.50), 1.35, 1.65
A_2	(1.00), 0.85, 1.15	(1.00), 0.85, 1.15	(1.00), 0.85, 1.15
σ_X	(0.75), 0.60, 0.90	(1.50), 1.35, 1.65	(0.75), 0.60, 0.90
β_{01}	(0.75) , 0.60, 0.90	(0.80), 0.65, 0.95	(0.75) , 0.60, 0.90
β_{02}	(0.75) , 0.60, 0.90	(0.80), 0.65, 0.95	(0.75) , 0.60, 0.90
ϱ	(0.50), 0.35, 0.65	(0.20), 0.05, 0.35	(0.50), 0.35, 0.65
α	(-1.50), -1.65, -1.35	(-1.50), -1.65, -1.35	(-1.50), -1.65, -1.35

[1] Values in parentheses are held constant while the other parameters change.
[2] Set 3 is a replication of the parameter values in set 1, but with the average across consumers of the $\{E_{it}(A_j)\}$'s equal to the average of the true product qualities, rather than the true product qualities (the assumption in Set 1.

Finally, we systematically vary the parameters across a range of plausible values. Table 11 presents the values of the parameters used in each dataset. Each column represents one set of datasets. The values in parentheses for each parameter are the "base" values that are held constant when the other parameters are changed. Each set manipulates all 7 parameters and uses three levels of each parameter, the base level and two different levels. The shaded areas in Table 12 indicate for two parameters how this works. The center cell represents the "no change" set in which none of the values changes, whereas the other four cells indicate change cases. Thus, each set contains a total of 15 simulated datasets (2 change cases per parameter plus the no change case) out of the possible 2187 data sets, i.e. 3 to the seventh power. Also, because ϱ is fixed to zero in the stable cases, the Smod cells have three no change sets.

Table 12: Parameter values used for β_{01} by ϱ (example pair)

$\varrho \setminus \beta_{01}$	0.60	(0.75)	0.90
0.35			
(0.50)			
0.65			

Note: Shaded areas have simulated datasets, white areas do not.

The first two sets make the same assumption about the distribution of the initial consumer expectations, $\{E_{it}(A_j)\}$. Specifically, the average of these expectations across consumers is assumed to equal the true quality. Thus, the only initial variation from true quality originates from initial errors arising from heterogeneity across consumers. The third set of datasets has an alternative assumption about the initial expectations--that the average across consumers of the initial expected quality equals the average of the two true product qualities. In this case, initial beliefs differ from true quality in a systematic fashion with a downward bias for the high quality product (product 1) and an upward bias for the low quality product (product 2). We test this alternative assumption because a similar assumption has been used in the past literature.

To examine each cell in the study, we select four measures of the performance of the model along each of the parameters. In all cases, we use a thinned sample from the posterior distribution after the initial burn-in. First, we calculate for each parameter the portion of runs for which the true parameter falls within the estimated 90% credible interval. This measure provides a sense for whether the estimation model is recovering the true parameters within the uncertainty the model assigns to those parameters. Second, we consider how much uncertainty about the parameters the estimation model retains after incorporating the data. We measure this by examining the average credible interval size for each parameter. Third, we calculate the root mean squared error. Fourth, we calculate the average bias. For the latter two measures, we use both the posterior mode and median as the point estimates. While the mean, median and mode provide similar results for most parameters, the parameter ϱ has a long right tail in most applications and the mean and median produce larger bias than the mode.

4.4.2 Convergence of Samplers

In this section, we detail the performance of the samplers described in the technical appendix. We apply each to a simulated dataset that is representative of the datasets used in the broader Monte Carlo study. The datasets are generated so that what the consumers actually do and what the estimation model assumes is the same (i.e., for the Smod the condition is Scond and for the Umod the condition is Ucond).

For both cases, we examine the convergence and autocorrelation in the Markov chain. To examine convergence we do both visual inspection of sequence plots and more formal tests. To examine autocorrelation, we examine the sequence plots and plot the autocorrelation function. We find that convergence in both models is rapid, but that high autocorrelations make long chains necessary.

4.4.2.1 The Smod Sampler

Figure 3 presents sequence plots of the parameters thinned at 250. Visual inspection suggests that convergence is achieved very rapidly. In addition, we do several formal tests of convergence. First, we use diagnostics proposed by Heidelberger and Welch (1981) and Geweke (1992), both of which confirm stationarity. Further, we examine values of the Kolmogorov-Smirnov statistic for a variety of thinning levels and initial burn-in sizes to establish that convergence is rapid, but that autocorrelation requires long chains to ensure the sample space is traversed sufficiently.

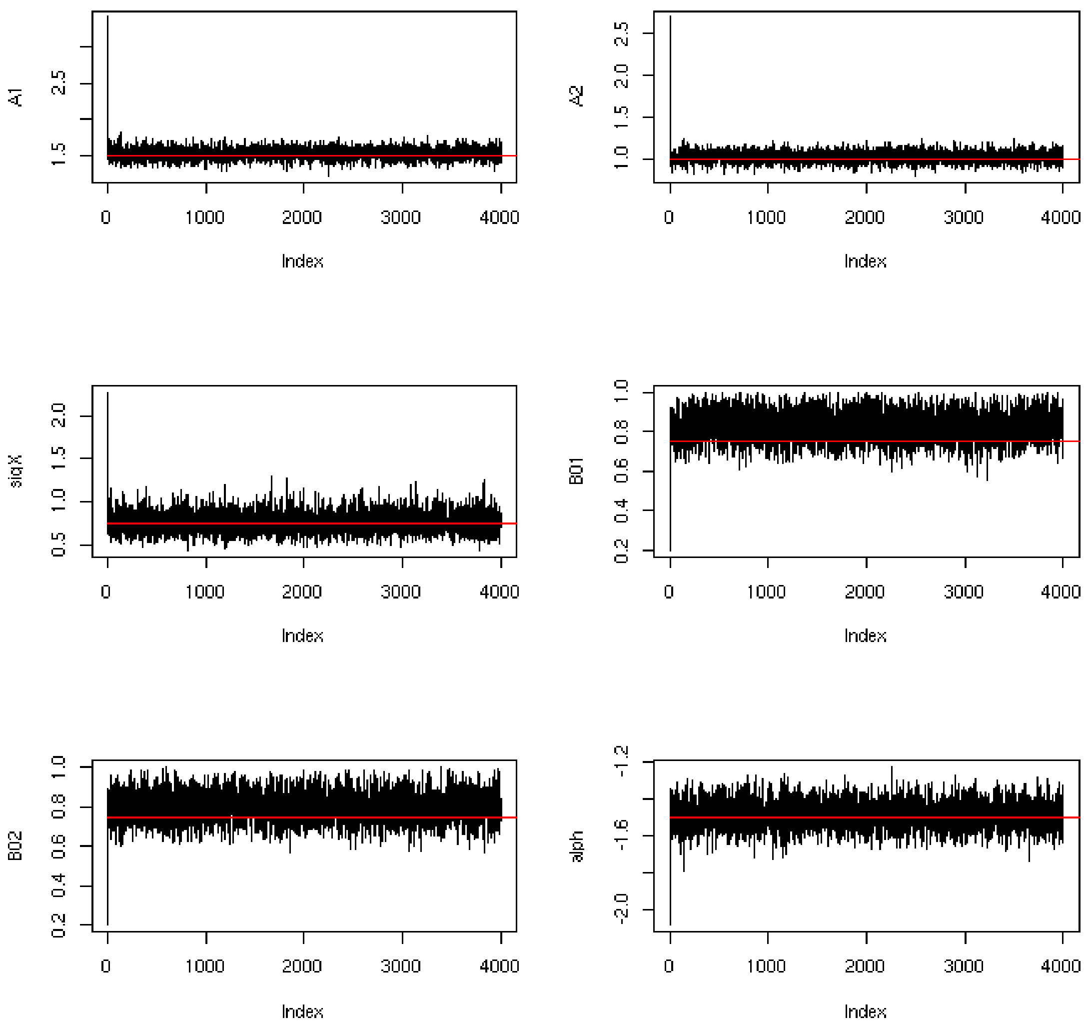

Figure 3: Sequence plots of structural parameters thinned at 250

Examination of the autocorrelation function indicate significant correlations up to approximately 200 lags (see Figure 4). This suggests as our analysis of the K-S statistics, that a relatively long chain is necessary for inference. We note, however, that once thinned at 250, autocorrelation no longer appears to cause significant problems as depicted in Figure 5.

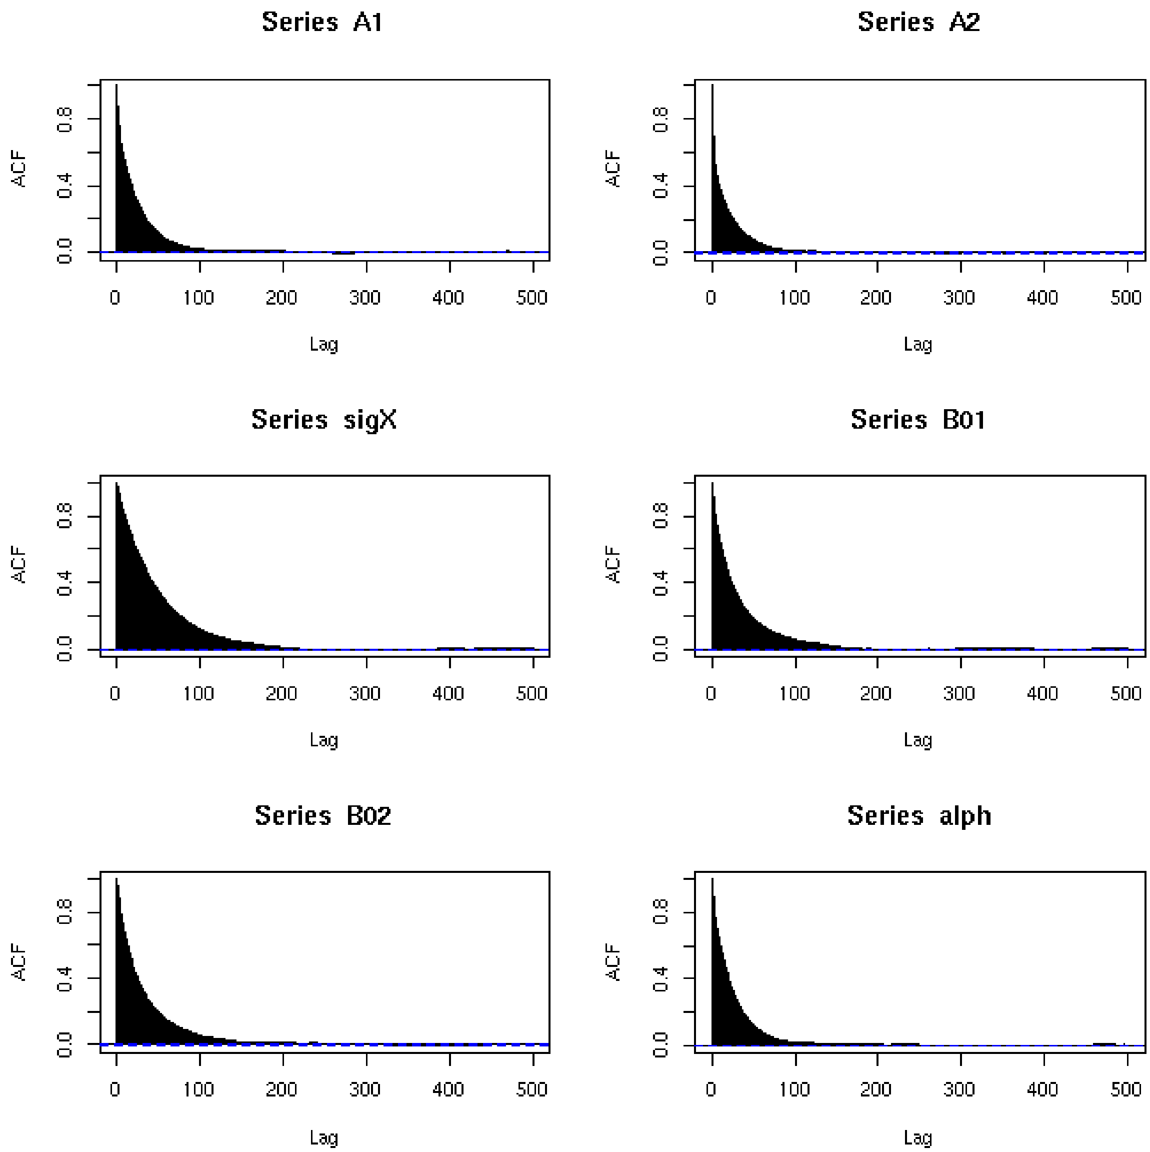

Figure 4: Autocorrelations prior to thinning

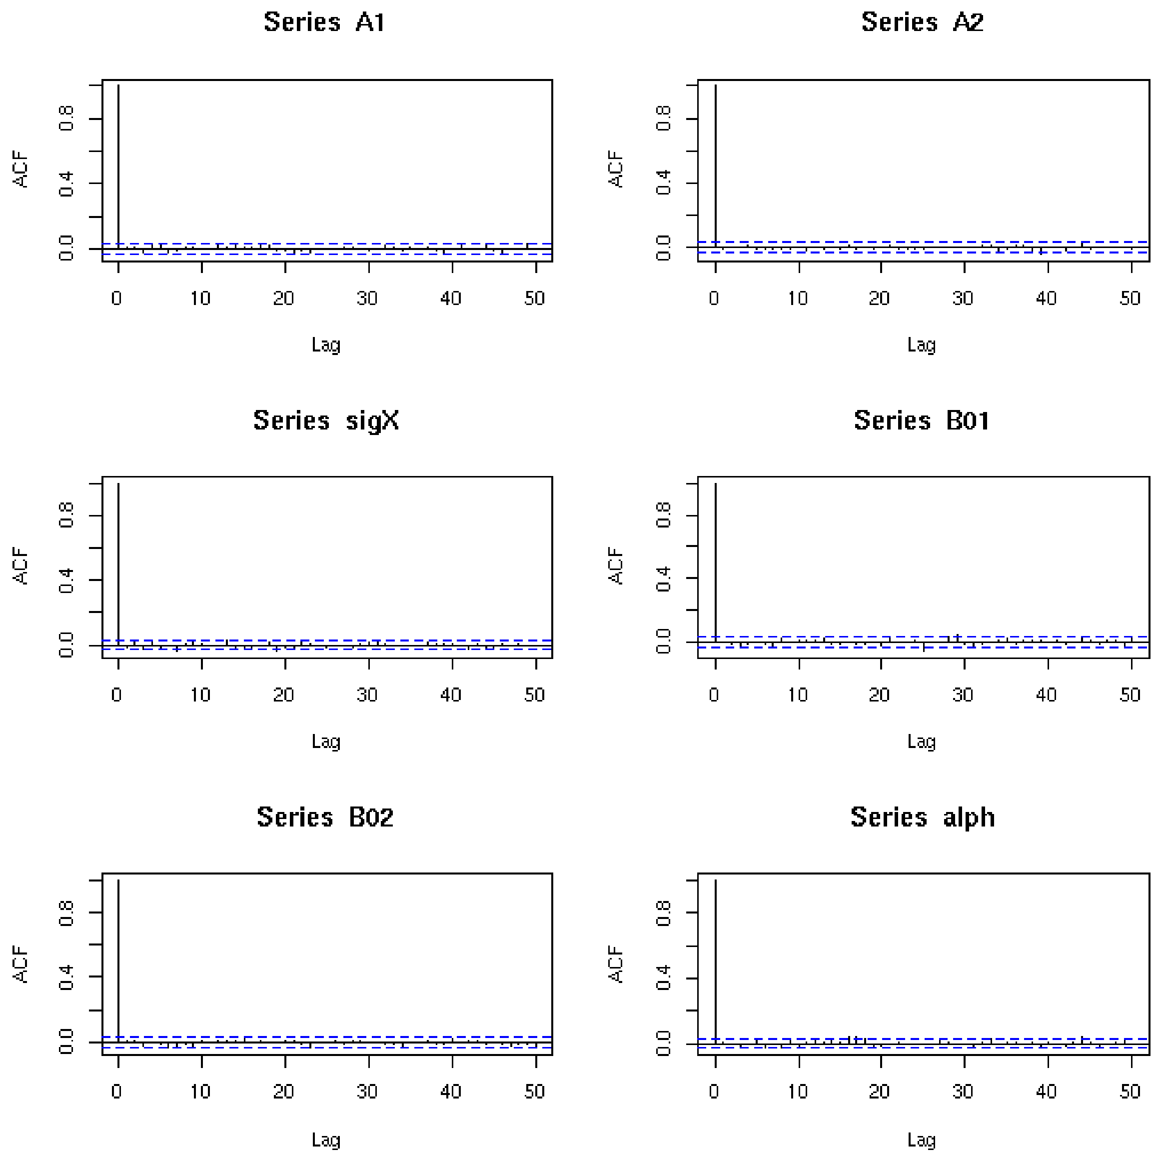

Figure 5: Autocorrelations after thinning at 250.

4.4.2.2 The Umod Sampler

Figure 6 presents sequence plots of the parameters thinned at 500. Like the Smod, visual inspection and formal tests indicate convergence, but with high autocorrelations.

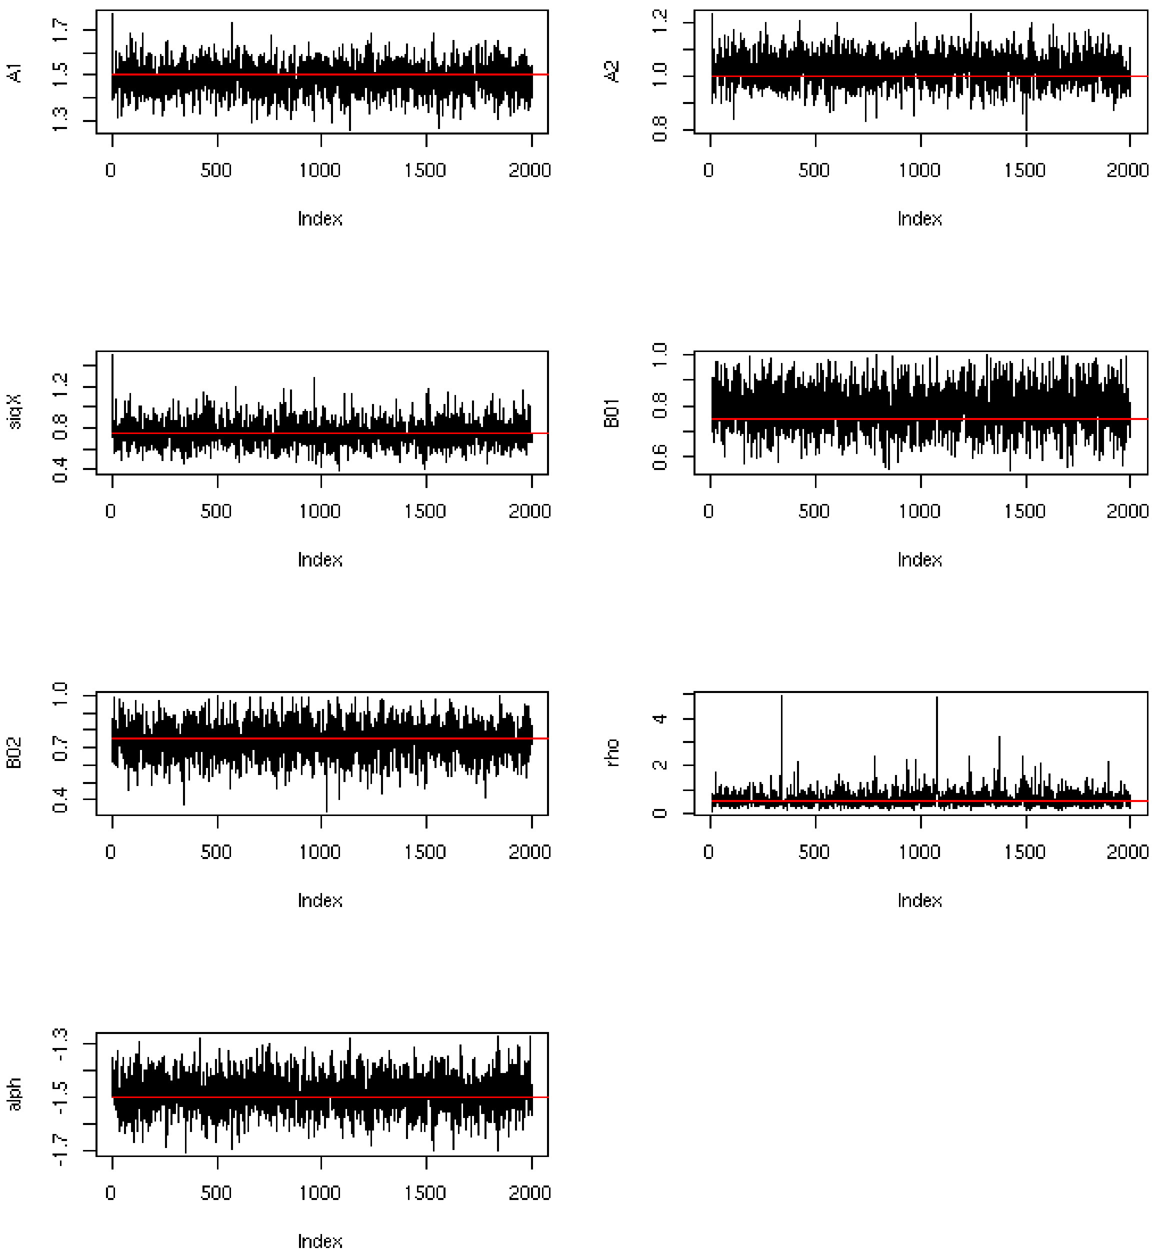

Figure 6: Sequence plots of structural parameters thinned at 500

Examination of the autocorrelation function indicate significant correlations up to nearly 2000 lags (see Figure 7). Once thinned at 500, autocorrelation no longer appears to cause significant problems as depicted in Figure 8.

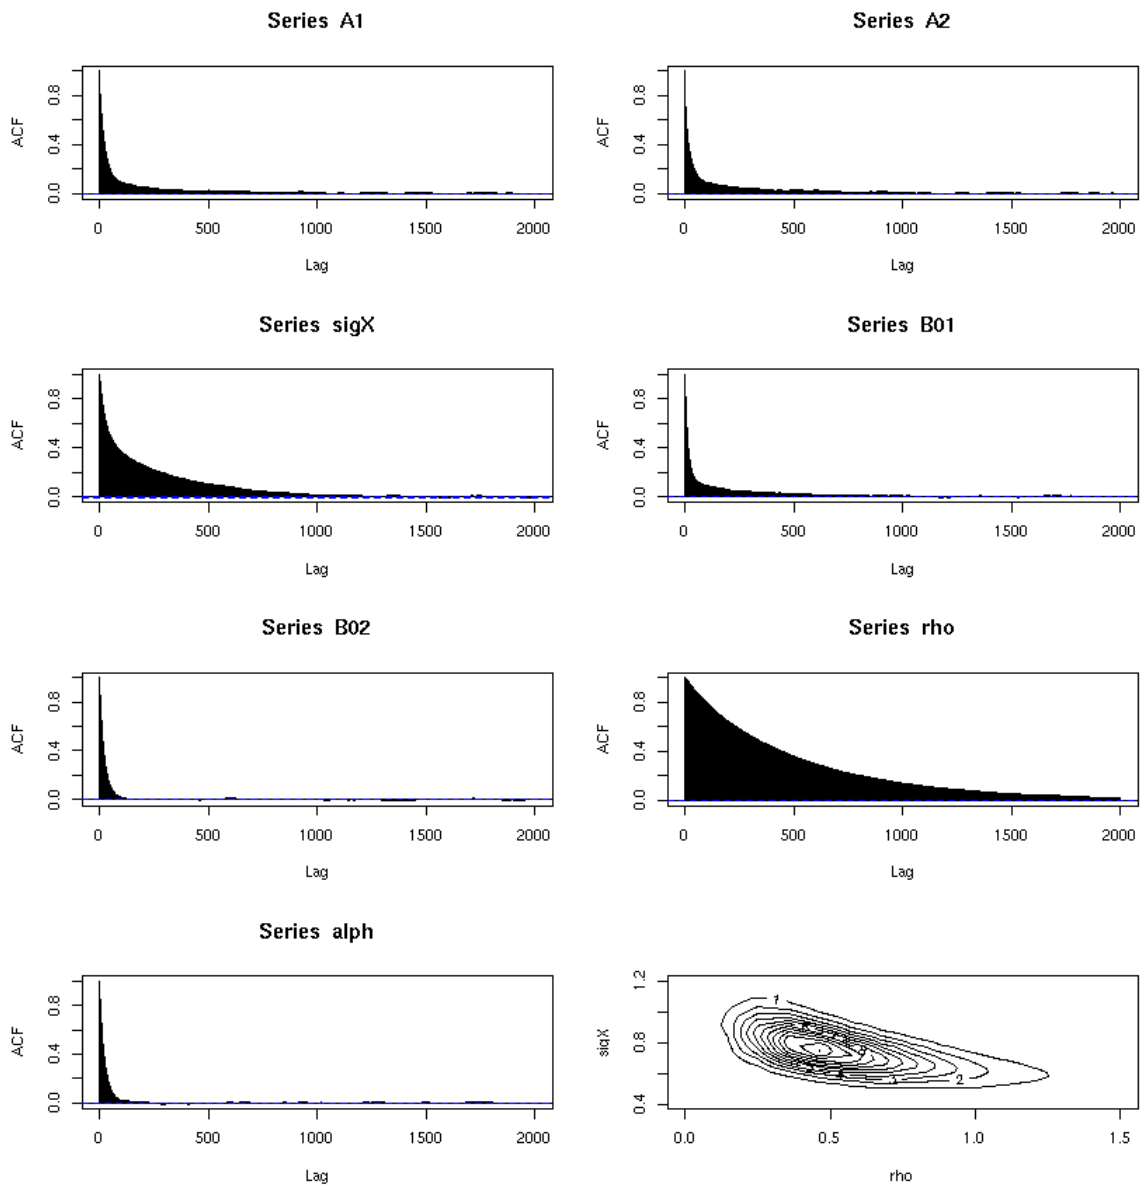

Figure 7: Autocorrelations prior to thinning

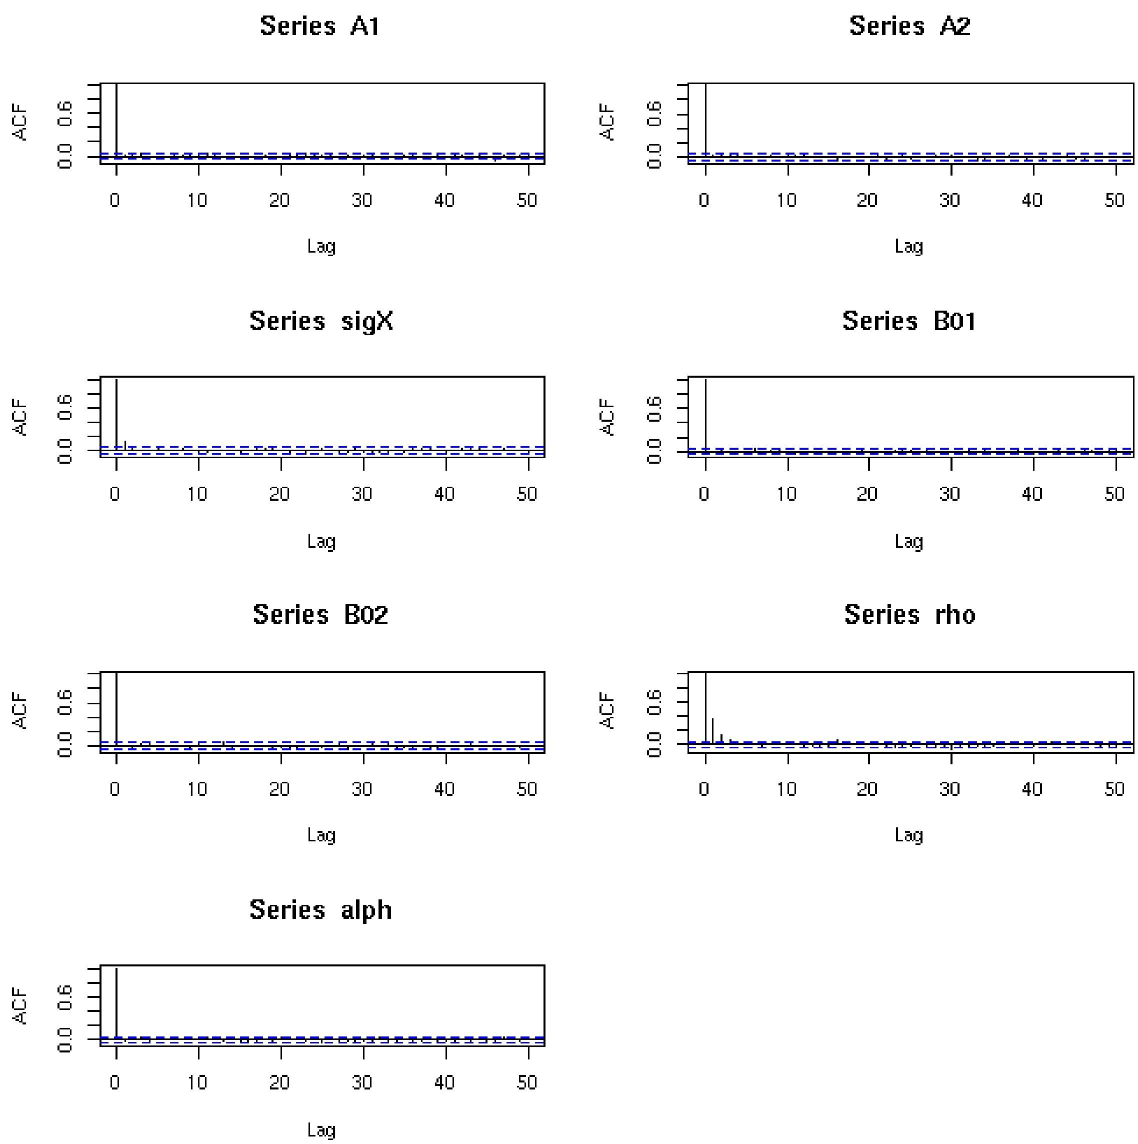

Figure 8: Autocorrelations thinned at 500

4.4.2.3 Summary

We find that convergence is rapid, but high autocorrelations make long chains necessary in order to ensure independent samples for inference. The sequence plots of parameters (Figure 3 and 6) also indicate that the convergent region covers truth for all parameters in both case. In the next section, we analyze the results of our Monte Carlo study to examine the performance of both the Smod and Umod in recovering parameters with different values of the true parameters.

4.4.3 Monte Carlo Study Results

In what follows we present and discuss the results of the Monte Carlo study. For each dataset, we draw 200,000 post-convergence posterior simulations from the sampler. For inference, we use a sample of 400 after thinning by 500 to approximate an independent sample. We present the cell level results in Table 13.

We begin by describing the recovery of the first six parameters (A_1, A_2, σ_X, β_{01},β_{02},α) in the models that match the actual condition (i.e., (Scond, Smod) and (Ucond, Umod)). In these two cells, all six parameters are recovered quite well from the model. The 90% credible intervals include the true parameter above or approximately at the 90% level that is expected. The root mean squared errors for the two models are similar and not problematically large. The bias is small or negligible in all of these cases and again similar across cells. Further, the credible intervals suggest levels of uncertainty that would allow identification of significant differences between the true qualities, but not between the initial learning parameters. Hence, when applied to data that matches the model assumptions, both models perform similarly well for these six parameters.

We now turn to the parameter ϱ in the (Ucond, Umod) cell. The true parameter for ϱ is inside the credible interval approximately the expected portion of the time. The average credible interval is large however, given the true values of ϱ. This large uncertainty about the distribution of ϱ in turn increases the root mean square error. The bias at the median is large and positive, but the bias at the mode is small and negligible. This bias at the median reflects a long right tail apparent in the posterior distribution of ϱ. In summary, ϱ is recovered reasonably well, but the data appears to provide relatively less information about the parameter compared to the other parameters.

Table 13: Monte Carlo Study Results

		Model Matches Actual Condition		Model Does Not Match Condition	
Parameter/Measure		Scond,Smod	Ucond,Umod	Scond,Umod	Ucond,Smod
%True in 90%CI	A_1	100	100	96	22
	A_2	87	93	91	9
	σ_X	93	93	93	56
	β_{01}	93	89	96	80
	β_{02}	98	100	96	87
	α	96	96	96	89
	ϱ		91	0*	
90%CI Length	A_1	0.22	0.22	0.23	0.22
	A_2	0.19	0.19	0.19	0.20
	σ_X	0.53	0.49	0.52	0.77
	β_{01}	0.24	0.29	0.24	0.24
	β_{02}	0.23	0.26	0.23	0.24
	α	0.22	0.22	0.22	0.21
	ϱ		0.79	0.04	
$\sqrt{}$MSE(Median)	A_1	0.049	0.045	0.052	0.17
	A_2	0.059	0.052	0.054	0.20
	σ_X	0.14	0.140	0.140	0.47
	β_{01}	0.065	0.094	0.070	0.085
	β_{02}	0.065	0.069	0.066	0.076
	α	0.052	0.052	0.051	0.070
	ϱ		0.220	0.016	
$\sqrt{}$MSE(Mode)	A_1	0.050	0.048	0.053	0.17
	A_2	0.059	0.052	0.054	0.19
	σ_X	0.130	0.140	0.130	0.43
	β_{01}	0.069	0.100	0.070	0.087
	β_{02}	0.071	0.078	0.074	0.079

	α	0.052	0.050	0.054	0.070
	ϱ		0.150	0.011	
Bias(Median)	A_1	0.010	0.008	0.023	-0.160
	A_2	-0.014	-0.006	0.001	-0.180
	σ_X	0.041	0.013	0.053	0.340
	β_{01}	-0.010	-0.007	-0.015	-0.019
	β_{02}	0.013	0.013	0.010	-0.024
	α	0.012	0.012	0.009	0.048
	ϱ		0.092	0.013	
Bias(Mode)	A_1	0.008	0.009	0.023	-0.160
	A_2	-0.013	-0.007	0.002	-0.180
	σ_X	0.009	-0.004	0.036	0.310
	β_{01}	-0.015	-0.004	-0.015	-0.022
	β_{02}	0.012	0.016	0.008	-0.026
	α	0.013	0.010	0.009	0.046
	ϱ		0.014	0.008	

*Note that it is not possible for the ϱ estimate to cover truth in the Scond, Umod cases, since it is bounded below at 0.

Next, we turn to the cases when the model does not match the actual conditions. To begin, we note that the (Scond, Umod) cell is similar to the cases where the model matches the actual conditions for the first six parameters. In addition, the model responds effectively to zero values of ϱ, despite this value being a boundary condition. Here, the estimated value of ϱ is quite small--small enough to allow accurate recovery of the remaining parameters. Thus, the results from this cell suggest very little is lost by applying the Umod to the case when consumers behave as if quality is stable.

Finally, we turn to the (Ucond, Smod) cell. It is this cell, where we expect the greatest potential problems. The question we pose is what, if any, biases are introduced into estimates if consumers act as if products are changing, but the model assumes they act as if products are stable.

We find that the initial learning rates and the pricing coefficient is recovered in a similar quality to the other cells. However, the true qualities and the variance of experiences are not. First, the true qualities are biased downwards by approximately -0.15. Hence, the model implies the products are less intrinsically attractive than the products actually are by between 10 to 15%. Second, the variance of the experiences is biased upwards by nearly 50%. Hence, the model implies that the product experience is less reliable as a signal (or perhaps the product is simply less reliable) than it actually is. Further, it implies that consumers have more to learn about the product than they actually do.

4.5 Discussion and Conclusions

In this essay, we introduce and explore the implications of an unstable learning model as applied to choice data. We develop this model to allow consumers to act as if products are changing even when they are in fact stable. We present a Bayesian approach to estimating the model that directly links Bayesian learning models to the time-series technique, the dynamic linear model (DLM). We further show through this link that Bayesian learning models are a specific constraint on this broader framework. Finally, we conduct a Monte Carlo investigation to determine how well the model performs as compared to a variant of the more commonly applied stable learning model.

We find that both the stable and unstable learning models perform well when the way consumers learn matches the assumptions of the model. Further, we find that the unstable model performs as well as the stable model when consumers act as if the quality is stable. More specifically, when the unstable model does not match how consumers actually learn, the model estimates a sufficiently small value for the uniquely unstable parameter, ϱ, that the remaining parameters are recovered just as well as in the stable model. However, when the stable model does not match how consumers actually learn, the parameter estimates are biased. In particular, the model estimates the true qualities as too low and the variance of experiences as too high. Thus, we demonstrate that the unstable model is a pure improvement over the stable model for the addition of a single parameter.

Further, we suggest that to the extent that learning models inform estimates of brand equity and decisions to develop, overcome, or exploit brand equity, the biases in true quality and experience variability that can arise from the use of the stable learning model can prove problematic. First, both biases are likely to cause managers to underestimate their brand equity. This could lead them to invest in building further brand equity. However, the unstable learning model would reveal that such efforts may be misaligned, since the problem is not low brand equity, but a poor estimate of brand equity. This is because in the unstable learning model the time between purchases increase consumer uncertainty. On potential implication is that the firm should focus on reducing the time between experiences. This can be done by getting the consumer to purchase the brand more frequently or by sending other messages that help consumers to recognize that quality has not changed.

Second, because the variance of experiences is estimated as too low, managers may incorrectly assume that quality is too unreliable and may make investments to rectify this lack of reliability. Such investments may be less effective than expected. Third, these estimates of high variance of the experiences also suggests that consumers have more to learn about the products than they actually have to learn. Instead, consumers are relearning the same things over and over, because they assume the products may have changed. This problem for exploiting brand equity can be an opportunity for launching new products. In particular, consumers will learn about new products more rapidly than existing learning models imply.

Finally, we suggest several avenues of future research. First, though we intentionally restricted the model to be parsimonious, various straightforward extensions would be necessary in order to use the model in an application. The specific extensions would depend on the application, but two generic ones are obvious. These extensions include incorporating unobserved heterogeneity in price sensitivities and increasing the number of products from two to many. Both of these modifications are straightforward given the existing model structure. Further, one would need to develop the measurement model for the initial expectations. This measurement model is not heavily restricted; the formulation used in the estimation model allows any type of covariance structure between the initial expectations of consumers without fundamentally altering the sampler.

Second, to apply the model to experimental data to explore the structural role of learning without requiring expectations data, some extensions may be valuable. In particular, incorporating heterogeneity in the learning parameters, ϱ and β_0 would allow a the model to identify the important heterogeneity across consumers in the way that they learn from product experiences. Based on our investigation, such individual level learning parameters would require significantly more time periods of data to recover parameters with any meaningful informativeness. The work of Biele et al. (forthcoming) presents a good approach for conducting such experiments.

5. Essay One Appendices

5.1 Estimation

In all of the models in essay one, we use Bayesian estimation as implemented in the WinBUGS software. We thin samples based on the autocorrelation function and report results using these thinned samples. We are generally very conservative in throwing out early samples to ensure that only samples from a converged chain are used in analysis. Multiple MCMC chains are used to help diagnose convergence and ensure we have achieved sufficient mixing in the latent class models. Finally, in our latent class models, we allow simulation without constraining the model to an identified space. This can improve simulation speed, but introduces a label switching problem. Where necessary, we employed label correcting methods suggested by Stephens (2000) to better interpret the results.

Log marginal likelihoods are calculated using the harmonic mean estimator (Raftery 1996). We use large samples and bootstrapped subsamples to examine robustness to low probability draws. Deviance plots were also used to verify differences between model fits identified using log marginal likelihoods.

We fit the model to multiple values of G and combinations of stable and unstable groups. For the first four periods of Study 1, this process produced Table 2, which we use to compare marginal likelihoods and select the best model, three unstable and three stable groups. In fact, once two unstable and one stable learning groups are allowed, the substantive conclusions do not differ dramatically. After this point, the additional groups appear to be enabling more dispersion of the α_1 and D/S parameters. For the later periods, we do this same process. Here, the number of subjects in each cell is smaller and so we identify fewer groups as the best fit.

In cases with larger numbers of groups in periods 5-10, sampling became problematic with diffuse priors. Some groups were assigned no individuals on some draws (i.e., a model with one less group was being selected by the routine). To accommodate these situations, we use a slightly less diffuse prior on D/S. In these cases, we use a gamma(1, m) prior, where m is less than or equal to 1. We use multiple values of m to ensure that the prior is not overly restrictive.

5.2 Wording of Environmental Beliefs Question

The following question was asked after all experiences occurred as a retrospective debrief:

> Please rate the statements below by allocating exactly 100 points among the statements. You may allocate 100 points among as many statements as you feel appropriate.
>
> A. The average wait time seemed to exhibit an increasing or decreasing trend that I expected to continue.
> B. The average wait time seemed to be changing in a random fashion. Although I think the average wait time was changing, I thought that it (average wait time) might go either direction.
> C. The average wait time seemed to have shifted and was a new average level that I expected might continue.
> D. The average wait time did not seem to be changing and the degree of variation in the wait times seemed consistent.

E. The average wait time did not seem to be changing, but the degree of variation seemed to have increased.

6. Essay Two Appendices

6.1. Proofs of Unstable vs. Stable Learning Proposition

Let $\beta_{ijt(U)}$ be the value of β_{ijt} from an unstable learning model and $\beta_{ijt(S)}$ be the value of β_{ijt} from a stable learning model. Further assume that σ^2_{Aijt-1} and σ^2_{Xj} are the same across both learning models.

6.1.1 Proof A1.1 $\beta_{ijt(U)} > \beta_{ijt(S)}$

The proof arises from the observation that a variance must be positive and by direct inspection of equations 9 and 10 underlying $\beta_{ijt(U)}$ and $\beta_{ijt(S)}$. Thus, for the same starting uncertainties and experience variabilities, the unstable learning model has larger Kalman gain coefficients than the stable learning model.

6.1.2 Proof A1.2 $\lim_{k\to\infty} \beta_{ijk(U)} - \beta_{ijk(S)} > 0$

This is easy to see based on A1.1 and noting that with no experience the unstable model Kalman gain coefficient increases, whereas the stable learning model's stays constant.

For the last proof we assume that $\beta_{ijt-1(S)} = \beta_{ijt-1(U)}$.

6.1.3 Proof A1.3 $\beta_{ijt(U)} - \beta_{ijt-1(U)} < \beta_{ijt(S)} - \beta_{ijt-1(S)}$

We use equations 9a and 10a and note that $\sigma^2_{Uj}/\sigma^2_{Xj}$ is always positive since both are positive variances and hence $\beta_{ijt(U)} > \beta_{ijt(S)}$. With this fact in hand, A1.3 is direct.

6.2 Assumptions that σ^2_{Uj} and σ^2_{Xj} are known

We first note that under the reparameterizations and normalizations, the assumption that σ^2_{Uj} and σ^2_{Xj} are known amounts to only assuming that σ^2_{X} is known, since the unstable variance is counterfactual (i.e., it doesn't affect actual quality levels). Further, with our reparameterization, we have σ^2_{X} as distinct from the β_{0j}'s and ϱ. Since only the β_{0j}'s and ϱ determine the Kalman gain coefficients as shown below, the parameter σ^2_{X} that is recovered from the structural estimation actually has no effect on the learning model, except through the variance of the experiences. In other words, that parameter is not known to consumers!

To see that the σ^2_{X} parameter doesn't affect the Kalman gain coefficients of the learning model note that it does not appear in equations 9a and 10a. Thus, only the initial value, β_{0j}, and ϱ are necessary to determine every subsequent coefficient.

However, there is one sense in which the assumption of known σ^2_{Uj} and σ^2_{Xj} is binding. This is the sense in which consumers are not adjusting the value of ϱ and or β_{ijt-1} in response to observed experiences. Such an adjustment process would insert complex non-linear perturbations in the path of Kalman gain coefficients that would be a function of the unobserved experiences. Because these perturbations would be relatively small in comparison to the primary effect of a constant ϱ, we view this as a relatively innocuous assumption.

`

6.3 Gibbs sampler and conditional distributions

To conduct inference about the structural models, we use block Gibbs sampling methods. In this section we discuss the blocks of the samplers and where insightful the exact conditional distributions. The differences between the models are relatively minor, so we illustrate with the SUmod and discuss modifications to sample from the SSmod. The sampler has five essential blocks and we describe each below.

1. The first block is the data augmentation step (Tanner and Wong 1987). In this step, the choices are augmented with unobserved "utilities" so that an underlying normally distributed variable Y is available for the remaining sampling steps. This block is conditional on all data and parameters in the model. The distribution, $p(\mathbf{Y} \mid -)$, is a truncated normal distribution, where chosen options have truncated utilities with values greater than the other options and non-chosen options have truncated utilities with values less than the chosen option. The means for the Y_t normal variates are $F' \theta_t + g X_t$ and the variances are the identity matrix.

2. The second block is the forward filter backward sampling algorithm to update the θ_t. We employ the well-established forward-filtering backward-sampling routine used in the dynamic linear model (DLM) approach to sampling (West and Harrison 1997). The sampling follows from the mapping of the Bayesian learning model into the DLM as given in section 4.3.3. This results in a series of conditionally normal distributed random draws of the θ_t.

3. The third block is the draw of α. These draws arise from a conditional normal linear model with known variance based on the underlying utilities. The conditional normal linear model is represented in Equation (17), where the θ_t are treated as data. The variance is known since the augmented data, Y_t, have an identity variance matrix.

4. The fourth block contains the terms associated with the experience distribution, (i.e., A_j's and σ^2_X). The A_j's are conditionally normal based on Equation (15) and the σ^2_X is conditionally gamma from the same equation. We draw these separately, but jointly they represent a normal linear model set of draws.

5. The fifth block contains Metropolis-Hastings random walk samplers for the β_{0j}'s and ϱ. All affect the likelihood only through their influence on the β_{ijt}'s. Thus, the only equation that is included is Equation (14), where the θ_t are treated as data. For these samplers, we use an adaptive first stage followed by a longer series of draws from a stable sampler. The candidate samplers are truncated so that only valid values of the parameters are potentially sampled.

Bibliography

Ackerberg, Daniel A. (2003), “Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination,” *International Economic Review*, 44(3), 1007-1040.

Akçura, M. Tolga, Füsun F. Gönül, and Elina Petrova (2004), “Consumer Learning and Brand Valuation: An Application on Over-the-Counter Drugs,” *Marketing Science*, 23(1), 156-169.

Barry, Donald M. and Gordon F. Pitz (1979), "Detection of Change in Nonstationary, Random Sequences," *Organizational Behavior and Human Performance*, 24, 111-125.

Biele, Guido, Ido Erev, and Eyal Ert (forthcoming), “Learning, Risk Attitude, and Hot Stoves in Restless Bandit Problems,” *Journal of Mathematical Psychology*.

Boulding, William, Ajay Kalra, Richard Staelin, and Valarie A. Zeithaml (1993), “A Dynamic Process Model of Service Quality: From Expectations to Behavioral Intentions,” *Journal of Marketing Research*, 30(1), 7-27.

---, Ajay Kalra, and Richard Staelin (1999), “The Quality Double Whammy,” *Marketing Science*, 18(4), 463-484.

Brehmer, Brendt (1980), “In One Word: Not from Experience,” *Acta Psychologica*, 45, 223-241.

Brown, Eric R. and Alice L. Bane (1975), "Probability Estimation in a Chance Task with Changing Probabilities," *Journal of Experimental Psychology: Human Perception and Performance*, 104(2), 183-187.

Camacho, Nuno, Bas Donkers, and Stefan Stremersch (2008), "Salience in Physician Decision-Making," *Informs Marketing Science Conference Presentation*, Vancouver.

Chan, Tat, Chakravarthi Narasimhan, and Ying Xie (2007), "Impact of Treatment Effectiveness and Side Effects on Prescription Choices," Working paper.

Ching, Andrew (2003), "A Dynamic Oligopoly Structrual Model for the Prescription Drug Market After Patent Expiration," Working Paper.

Chintagunta, Pradeep K., Renna Jiang, and Ginger Z. Jin (2007), "Information, Learning, and Drug Diffusion: the Case of Cox-2 Inhibitors," Working paper.

Coscelli, Andrea and Matthew Shum (2004), "An Empirical Model of Learning and Patient Spillovers in New Drug Entry," *Journal of Econometrics*, 122, 213-246.

Crawford, Gregory S. and Matthew Shum (2005), “Uncertainty and Learning in Pharmaceutical Demand,” *Econometrica*, 73(4) 1137-1173.

Dudycha, Arthur L., Myron G. Dumoff, and Linda W. Dudycha (1973), "Choice Behavior in Dynamic Environments," *Organizational Behavior and Human Performance*, 9, 328-338.

Edgell, Stephen E. (1983), "Delayed Exposure to Configural Information in Nonmetric Multiple-Cue Probability Learning," *Organizational Behavior and Human Performance*, 32, 55-65.

Erdem, Tulin (1998), “An Empirical Analysis of Umbrella Branding,” *Journal of Marketing Research*, 35(Aug), 339-351.

--- and Michael P. Keane (1996), “Decision Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets,” *Marketing Science*, 15(1), 1-20.

---, ---, T. Sabri Oncu, and Judi Strebel (2005), “Learning about Computers: An Analysis of Information Search and Technology Choice,” *Quantitative Marketing and Economics*, 3, 207-246.

---, Michael P. Keane, Baohong Sun (2007), "A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality," Working paper.

---, Joffre Swait, Susan Broniarczyk, Dipankar Chakravarti, Jean-Noel Kapterer, Michael Keane, John Roberts, Jan-Benedict E. M. Steenkamp, and Florian Zettelmeyer (1999), “Brand Equity, Consumer Learning, and Choice,” *Marketing Letters*, 10(3), 301-318.

---, Ying Zhao, and Ana Valenzuela (2004), "Performance of Store Brands: A Cross-Country Analysis of Consumer Store-Brand Preferences, Perceptions, and Risk," *Journal of Marketing Research*, 41(Feb), 86-100.

Erev, Ido and Ernan Haruvy (2008), "Learning and the Economics of Small Decisions," Draft chapter submitted to the second volume of *The Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth.

Gelman, A. (1996) “Inference and Monitoring Convergence,” in *Markov Chain Monte Carlo in Practice*, Boca Raton: Chapman & Hall/CRC, 131-144.

Godsill, Simon J., Arnaud Doucet, and Mike West (2004), "Monte Carlo Smoothing for Nonlinear Time Series," *Journal of the American Statistical Association*, 99(465), 156-168.

Hoch, Stephen J. and John Deighton (1989), “Managing What Consumers Learn from Experience,” *Journal of Marketing*, 53(2), 1-20.

Israel, Mark (2005), “Services as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance,” *American Economic Review*, 95(5), 1444-1463.

Iyengar, Raghuram, Asim Ansari, and Sunil Gupta (2007), "A Model of Consumer Learnign for Service Quality and Usage," *Journal of Marketing Research*, 44(Nov), 529-544.

Keller, Kevin Lane (1993), "Conceptualizing, Measuring, and Managing Customer-Based Brand Equity," *Journal of Marketing*, 57(Jan), 1-22.

Kopalle, Praveen K. and Donald R. Lehmann (1995), "The Effects of Advertised and Observed Quality on Expectations about New Product Quality," *Journal of Marketing Research*, 32(3), 280-290.

---, and Donald R. Lehmann (2006), "Setting Quality Expectations When Entering a Market: What Should the Promise Be?" *Marketing Science*, 25(1), 8-24.

Kumar, Piyush, Manohar U. Kilwani, Maqbool Dada (1997), "The Impact of Waiting Time Guarantees on Customers' Waiting Experiences," *Marketing Science* 16(4), 295-314.

Lindberg, Lars-Ake and Berndt Brehmer (1976), "Transfer in Single Cue Probability Learning," *Organizational Behavior and Human Performance* 16, 177-192.

Massey, Cade and George Wu (2005), "Detecting Regime Shifts: The Causes of Under and Overreaction," *Management Science*, 51(6), 932-947.

Mehta, Nitin, Surendra Rajiv, and Kannan Srinivasan (2003), "Price Uncertainty and Consumer Search: A Structural Model of Consideration Set Formation," *Marketing Science* 22(1), 58-84.

---, Surendra Rajiv, and Kannan Srinivasan (2004), “Role of Forgetting in Memory-Based Choice Decisions: A Structural Model,” *Quantitative Marketing and Economics*, 2, 107-140.

Meyer, Robert J. (1981) "A Model of Multiattribute Judgments under Attribute Uncertainty and Informational Constraint," *Journal of Marketing Research*, 18(4), 428-441.

Meyer, Robert J. and Arvind Sathi (1985a), "A Multiattribute Model of Consumer Choice during Product Learning," *Marketing Science*, 4(1), 41-61.

Mitra, Debanjan and Peter N. Golder (2006), "How Does Objective Quality Affect Perceived Quality? Short-term Effects, Long-term Effects and Asymmetries," *Marketing Science*, 25(3), 230-247.

Muchinsky, Paul M. and Arthur L. Dudycha (1975), "Human Inference Behavior in Abstract and Meaningful Environments," *Organizational Behavior and Human Performance*, 13, 377-391.

Narayanan, Sridhar and Puneet Manchanda (2007), "Heterogeneous Learning and the Targeting of Marketing Communication for New Products," Working paper.

Osborne, Matthew (2006), "Consumer Learning, Habit Formation, and Heterogeneity: A Structural Examination," Working Paper.

Peterson, Cameron R., Kenneth R. Hammond, and David A. Summers (1965), "Multiple Probability-Learning with Shifting Weights of Cues," *The American Journal of Psychology*, 78(4), 660-663.

Raftery, Adrian E. (1996), "Hypothesis Testing and Model Selection," in *Markov Chain Monte Carlo in Practice*, W. R. Gilks, S. Richardson, and D. J. Spiegelhalter eds., Boca Raton: Chapman & Hall, 163-187.

Rapoport, Amnon (1979), *Response Models for Detection of Change*, D. Reidel Publishing Company, Heningham, Mass.

Rieskamp, Jörg and Philipp E. Otto (2006), "SSL: A Theory of How People Learn to Select Strategies," *Journal of Experimental Psychology: General*, 135(2), 207-236.

Roberts, John H. and Glen L. Urban (1988), "Modeling Multiattribute Utility, Risk, and Belief Dynamics for New Consumer Durable Brand Choice," *Management Science*, 34(2), 167-185.

Ruffner, John W. and Paul M. Muchinsky (1978), "The Influence of Shifting Cue Validity Distributions and Group Discussion Feedback on Multiple Cue Probability Learning," *Organizational Behavior and Human Performance* 21, 189-208.

Rust, Roland T., Jeffrey Inman, Jianmin Jia, and Anthony Zahorik (1999), "What You Don't Know about Customer-Perceived Quality: The Role of Customer Expectation Distributions," *Marketing Science*, 18(1), 77-92.

Shin, Sangwoo, Sanjog Misra, and Dan Horsky (2007). "Disentangling Preferences, Inertia, and Leaning in Brand Choice Models," Working paper.

Sriram, S., Pradeep Chintagunta, and Ramya Neelamegham (2006), "Effects of Brand Preference, Product Attributes, and Marketing Mix Variables in Technology Product Markets," *Marketing Science*, 25(5), 440-456.

Stephens, Matthew (2000), "Dealing with Label Switching in Mixture Models," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 62(4), 795-809.

Summers, David A. (1969), "Adaptation to Change in Multiple Probability Tasks," *The American Journal of Psychology*, 82(2), 235-240.

Surowiecki, James (2004), "The Decline of Brands," *Wired*, 12(11-November), 2 pages.

Szymanowski, Maciej and Els Gijsbrechts (2008), "Conditional Cross-Brand Learning: When Are Private Labels Really Private?" Working paper.

Tanner, Martin A. and Wing Hung Wong (1987), "The Calculation of Posterior Distributions by Data Augmentation," *Journal of the American Statistical Association*, 82(398), 528-540.

Wang, Lei K., Eric Anderson, and Karsten Hansen (2007), "Retail Brand Equity and Consumer Learning," Working paper.

West, Mike and Jeff Harrison (1997), *Bayesian Forecasting and Dynamic Models, Second Edition*, New York: Springer.

Zhao, Yi, Ying Zhao, and Kristiaan Helsen (2007), "Consumer Learning in a Turbulent Market Environment: Modeling Consumer Choice Dynamics in the Wake of a Product Harm Crisis," Working paper.

Biography

Mitchell Lovett was born in Ann Arbor, Michigan in 1974 and lived in Michigan and Ohio for the first 22 years of his life. He received his Bachelors of Arts with majors in Mathematics, Economics, and German from Ohio Wesleyan University, where he graduated Summa Cum Laude and earned membership to Phi Beta Kappa. After college, he worked for three years at the Center for Creative Leadership in Greensboro, North Carolina. He met Cali Thompson, who he later married. He followed her to Boise, Idaho and received an assistantship to attend Boise State University, where he earned his Masters of Business Administration. It was there, through his interaction with his Professors Jason MacDonald and Shikhar Sarin, that he became excited about conducting academic research in marketing. These interactions led to two conference presentations and ultimately an article in *The Journal of the Academy of Marketing Science*.

After completing his MBA, Mitchell attended Duke University to pursue the PhD in Business Administration. During the program, he was nominated to be Duke's attendee at the American Marketing Association's Sheth Foundation Doctorial Consortium in 2007 and, with Professor Christine Moorman, earned a grant-bearing Outstanding Silver Award in the Institute for the Study of Business Marketing Research Challenge. While finishing his dissertation, Mitchell, started his new job as an Assistant Professor at the Simon Graduate School of Business in the University of Rochester.