2018

Essays On Dynamic Updating Of Consumer Preferences

Tong Lu
University of Pennsylvania, tonglu@wharton.upenn.edu

Essays On Dynamic Updating Of Consumer Preferences

Abstract

Consumers dynamically update their preferences over time based on information learned through product search and consumption experiences, particularly in online media. Using three unique datasets from different domains, we address specific ways in which firms can use rich information about their customers' behaviors to improve (1) the visual display of products on a webpage in online shopping, (2) predictions of new product adoption in online gaming, and (3) the timing of product release in online learning. First, we explore how consumers visually search through product options using eye-tracking data from two experiments conducted on the websites of two online clothing stores, which can inform retailers on how to position products on a virtual webpage. Second, we examine how consumers' variety-seeking preferences change depending on past consumption outcomes within the context of an online multi-player video game, which can be used to improve predictions of new product adoption. Third, we use clickstream data from an online education platform to test theories of goal progress, knowledge accumulation, and boundedly rational forward-looking behavior, which can be used to explain binge consumption patterns and inform content providers on the best way to structure and release content. In each of these three projects, we build a mathematical model of individual decisions, with the parameterization grounded in theories of consumer behavior, and we demonstrate through in-sample prediction that our model is able to capture specific heterogeneous patterns within the data. We then test that our model is able to make out-of-sample predictions related to managerial interventions, and empirically verify our predictions using either lab experiments or new field data following a natural experiment policy change.

Degree Type
Dissertation

Degree Name
Doctor of Philosophy (PhD)

Graduate Group
Marketing

First Advisor
Eric T. Bradlow

Second Advisor
J. Wesley Hutchinson

Keywords
binge consumption, bounded rationality, natural experiment, online education, online gaming, online retail

Subject Categories
Advertising and Promotion Management | Education | Marketing | Statistics and Probability

ESSAYS ON DYNAMIC UPDATING OF CONSUMER PREFERENCES

Tong Lu

A DISSERTATION

in

Marketing

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy

2018

Supervisor of Dissertation   Co-Supervisor of Dissertation

_____   _____

Eric T. Bradlow   J. Wesley Hutchinson
Professor of Marketing   Professor of Marketing

Graduate Group Chairperson

_____

Catherine Schrand, Professor of Accounting

Dissertation Committee:

Robert J. Meyer, Professor of Marketing
David R. Bell, Professor of Marketing
Colin F. Camerer, Professor of Behavioral Economics, California Institute of Technology

ESSAYS ON DYNAMIC UPDATING OF CONSUMER PREFERENCES

Dedicated to my mom and dad, my boyfriend, and my two cats, Allison and Simon

ACKNOWLEDGEMENT

ABSTRACT

ESSAYS ON DYNAMIC UPDATING OF CONSUMER PREFERENCES

Tong Lu

Eric T. Bradlow

J. Wesley Hutchinson

Consumers dynamically update their preferences over time based on information learned through product search and consumption experiences, particularly in online media. Using three unique datasets from different domains, we address specific ways in which firms can use rich information about their customers' behaviors to improve (1) the visual display of products on a webpage in online shopping, (2) predictions of new product adoption in online gaming, and (3) the timing of product release in online learning. First, we explore how consumers visually search through product options using eye-tracking data from two experiments conducted on the websites of two online clothing stores, which can inform retailers on how to position products on a virtual webpage. Second, we examine how consumers' variety-seeking preferences change depending on past consumption outcomes within the context of an online multi-player video game, which can be used to improve predictions of new product adoption. Third, we use clickstream data from an online education platform to test theories of goal progress, knowledge accumulation, and boundedly rational forward-looking behavior, which can be used to explain binge consumption patterns and inform content providers on the best way to structure and release content. In each of these three projects, we build a mathematical model of individual decisions, with the parameterization grounded in theories of consumer behavior, and we demonstrate through in-sample prediction that our model is able to capture specific heterogeneous patterns within the data. We then test that our model is able to make out-of-sample predictions related to managerial interventions, and empirically verify our predictions using either lab experiments or new field data following a natural experiment policy change.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF ILLUSTRATIONS

ix

CHAPTER 1 : Introduction

How many decisions does a consumer make in a day? The answer depends on the granularity at which we define decisions. Economists tend to focus on decisions at relatively low temporal resolutions, including purchases or investments that unfold over days or even years. Meanwhile, psychologists may consider the neural and psychological processes involved in decision-making as decisions themselves. These decisions may be conscious or unconscious and can occur within seconds or fractions of a second. The work presented in this dissertation examines decisions that span the range of temporal resolutions, from split-second eye movements occurring within a 5-minute shopping trip, to clickstream decisions that determine the consumption of digital media content tracked across several weeks or even years. Figure 1 provides a summary of the types of data and temporal resolutions addressed in the following chapters of this dissertation.

Figure 1: Examples of consumer decisions across a range of temporal resolutions

Years/Days
(jobs, houses, cars)

Low

Minutes/Hours
(retail shopping trips)

Seconds
(mouse clicks, food choices)

**Temporal
Resolution of
Decision
Processes**

100 - 500 ms
(eye movements during
information search)

5 - 50 ms
(motor reflexes, image
processing, neuron firing)

High
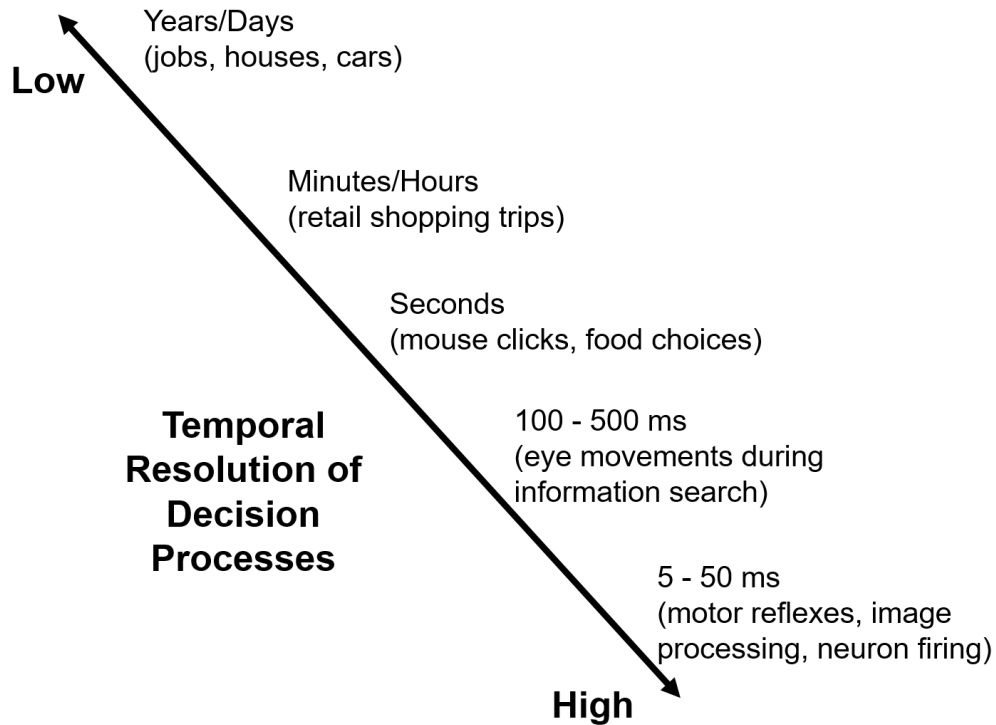
In marketing, it is common to treat each purchase as a single decision, but retail firms are increasingly interested in collecting more detailed information about their customers' consumption patterns that precede the final outcome. Advancements in process tracing methods enable both practitioners and researchers to track behaviors that lead up to a final purchase for material goods, as well as customer engagement in experiential goods. That is, firms are interested in the precursors to sales.

For example, clickstream data is commonly used as a measure of online activity such as shoppers' decisions to enter a store's website and examine particular product pages before checking out (Bucklin and Sismeiro 2003; Moe and Fader 2004; Montgomery et al. 2004) or when consuming digital media, such as online video streaming (Schweidel and Moe 2016). RFID (radio-frequency identification) sensors placed on store shelves or shopping carts can analogously measure a shopper's foot traffic through different areas of a brick-and-mortar store (Hui, Bradlow, and Fader 2009; Hui et al. 2013; Larson, Bradlow, and Fader 2005). Within each webpage or physical area of a store, we can measure the behavior of shoppers at even higher temporal resolution using eye-tracking equipment to observe their eye movements as they collect product information. Eye-tracking allows researchers to determine exactly what shoppers were looking at during each moment and make inferences about their knowledge of products and consideration sets (Chandon et al. 2009; Shi, Wedel, and Pieters 2013; Stüttgen, Boatwright, and Monroe 2012; Yang, Toubia, and de Jong 2015).

Using three unique datasets from different product domains that track consumer decisions measured at different temporal resolutions, we address specific ways in which firms can use this rich information to respond to the changing preferences of consumers. Using this data, we model how consumers are *dynamically* updating their preferences over time based on information gained when they learn about products through information search and experienced consumption. By empirically modeling these dynamic preferences, we can suggest ways for retailers to respond in terms of the product layout in an online store, predictions of whether or not consumers will adopt new versions of the product, and decisions about

the timing of product release.

In Chapter 2, we focus on information search on the landing page of an online store in order to test boundedly rational models of eye movements and choice. Like the physical entrance of a brick-and-mortar store, the landing page of an online retailer draws shoppers deeper into the website, and landing page decisions are likely to have important downstream consequences. We explore the internal decision processes and external characteristics of the shopping environment that determine how extensively shoppers search on the landing page. Our findings are based on eye-tracking data collected during two incentive-compatible online shopping experiments. We build a sequential sampling model of information search that captures eye fixations as decisions, with the shopper ending search by clicking on a fixated link when its attractiveness crosses a decision threshold. Our baseline model is able to make accurate in-sample predictions regarding search length, as well as which links shoppers fixate on and ultimately click on. In addition, we find that extending this baseline model with forward-looking and threshold-learning components improves model fit and in-sample predictions. We also conduct parametric sensitivity analyses to estimate a "rationalizing exchange rate" for the costs and benefits of eye fixations. Finally, we use counterfactual simulations to predict the effects of product layout on search in Experiment 1 and empirically verify our predictions in Experiment 2, which provides implications for the design of online store displays.

In Chapter 3, we focus on modeling variety-seeking preferences to predict new product adoption within the context of online video games. As consumers engage in experiential products, they may exhibit variety-seeking behavior by switching frequently between different options within the same category, or exhibit inertial behavior by consistently choosing the same option. We propose that variety-seeking may be characterized by both the incidence of option switching as well as the distance between options, as measured by the number of shared and unshared attributes. We model consumers with variety-seeking preferences that vary in response to consumption outcomes, which indicate the quality of the

consumer's experience for a specific consumption occasion. We test our model within the context of an online video game, in which players choose between different map environments for each round of play and consumption outcomes can be measured objectively by a player's performance during the round. A unique feature of our context is that the firm released five sets of expansion packs containing new maps at various dates following the initial release of the base game. We demonstrate that including our dynamic variety-seeking parameterization in the map choice model can help improve out-of-sample predictions of whether or not players "adopt" these new products by playing on the new maps following their release.

In Chapter 4, we focus on content consumption within the context of online education. Online education possesses unique features such as scheduled content tied to learning assessments that make it an appropriate testing ground for behavioral theories of how consumers progress towards their goals and accumulate knowledge. Using clickstream data from Coursera, an online educational platform, we build a model that captures individual decisions about which course to consume at each moment in time, whether the content is a lecture or a quiz, and when to take breaks of different lengths. There are three key features of our model: First, individuals may be heterogeneously motivated to progress towards the goal of completing the course by watching lecture videos and taking quizzes. Second, individuals can accumulate knowledge through lectures in order to pass quizzes. Third, individuals can be forward-looking in a boundedly rational way. We demonstrate that these model features are able to capture patterns in the data that we characterize as "binge learning," a topic of recent significant academic and practical importance. We also use our model to conduct counterfactual simulations to determine how the timing of content release (a key decision variable for online platforms) affects consumption and knowledge accumulation, and empirically test these predictions using data following a policy change on the Coursera platform.

Across these three chapters, we are able to observe at relatively high temporal resolu-

tion customers' behaviors during product search (Chapter 2) or how they engage in online experiential content (Chapters 3 and 4). By modeling how people incorporate product information from either search or through consumption experiences into their present and future preferences and choices, we can provide recommendations for improving the sequence of product offerings. In Chapter 2, by "sequence" we refer to the positioning of products on the landing page of an online store. Specifically, we test for whether the ordering of products on the page (i.e., best products shown first vs. worst products shown first) affects the length of search on a webpage and what link the shopper chooses to click on. In Chapters 3 and 4, a "sequence" refers to the temporal sequence and timing of product release (e.g., maps or online course content). We plan to explore how these sequences affect consumption choices, the rate of consumption, as well as more downstream behaviors like performance, enjoyment, and future product purchases.

CHAPTER 2 : Information Search on the Landing Page: Boundedly Rational

Models of Eye Movements and Choice

2.1. Introduction

Marketers have long held that increasing foot traffic within physical stores and increasing the click through rate within virtual environments usually leads to more purchases. Brick-and-mortar retailers create attention-grabbing displays to draw people farther into the store and consider more products (Pegler 2006). Similarly, the landing page of an online retailer can draw people deeper into the website. Online retailers will often fill the landing page with attractive product images that link to more detailed product information, or with category links that help guide shoppers to more organized groups of products. Thus, the landing page is especially important because the first click reflects both the shopper's preferences and the retailer's desire to influence downstream shopper decisions.

Despite these intuitions and the rapid growth of online retail, to our knowledge, no work has been done on how people make decisions on the landing page of an online store. We address this gap in the literature by modeling the internal decision processes and external characteristics of the shopping environment that determine how shoppers search on this first page before clicking on either a product or a category link that takes them deeper into the virtual store. How long shoppers spend in certain areas of a store may have downstream consequences regarding their consideration sets and final purchases. For example, Hui et al. (2013) find that increasing travel distance within a grocery store leads to more unplanned purchases, and Bronnenberg, Kim, and Mela (2016) find strong state dependence in the search path through the attribute space of digital cameras in online search such that early search is highly predictive of the characteristics of the camera eventually purchased.

To study how shoppers search on the landing page of an online store, we conducted two incentive-compatible eye-tracking experiments with female undergraduates shopping for 5 minutes at the website of an online clothing retailer. In both experiments, shoppers be-

gan at the landing page of the website that displayed a grid of featured products, plus links to available product categories (see Figure 2). Using eye-tracking methods, we were able to observe exactly where on the page participants were looking at each moment in time.

Figure 2: Layout of American Apparel landing page for Experiment 1
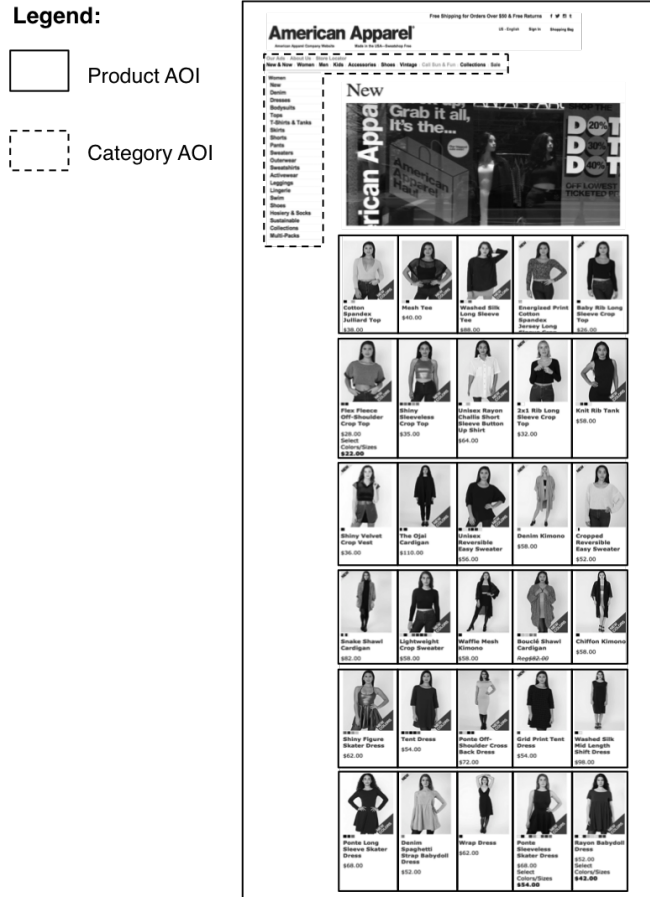
In order to describe our units of analysis, we introduce here some terms commonly used in the eye-tracking literature. The recorded stream of eye movements consists of *fixations*, defined as when the eyes are still for at least 50 milliseconds, and *saccades*, defined as when the eyes are rapidly moving. Only fixated visual information can be processed at a detailed level, while no information is processed during saccades. Thus, we use fixations as a measure of visual search on the landing page. Each fixation provides high resolution information from a spatial area of about 2 degrees of visual angle; outside of the fixated area, the quality

of information acquisition is highly degraded. For example, a single fixation can encode 5-7 letters when reading text and one product image when identifying image features (Rayner 1998).

To determine what information shoppers gain from each fixation, we divided up the landing page into *areas of interest (AOIs)*, each containing a link to a specific product ("Product AOI") or more general category of products ("Category AOI"), as shown in Figure 2. Thus, each fixation corresponds to an AOI "choice" that provides a small sample of visual information. Sampled information accumulates over time and provides a low-level record of the search process. Note that the AOIs that we used were relatively small, so two or three fixations would be required to sample all the information in the AOI, and we did not further subdivide Product AOIs by product image and price or the Category AOI by specific links.

2.1.1. Research Contributions

In our work, we make four main research contributions. First, we use eye-tracking data to investigate the decisions that shoppers make on a real retailer's webpage at a more granular level compared to clickstream data. We primarily focus on modeling shopper behaviors on the landing page, and we find that these behaviors are predictive of downstream decisions that occur later on in the shopping trip. Specifically, in both of our experiments, we find that additional fixations on the landing page reduce both the number of subsequent webpages visited and subsequent spending. These findings suggest that the initial decision about how to leave the landing page is an important determinant of subsequent search and that, consistent with Hui et al. (2013), perhaps it is better to send shoppers deeper into the store than to encourage them to linger on the landing page.

In line with the emerging stream of research that treats visual search patterns as part of the decision-making process (Chandon et al. 2009; Shi, Wedel, and Pieters 2013; Stüttgen, Boatwright, and Monroe 2012; Towal, Mormann, and Koch 2013; Yang, Toubia, and de Jong 2015), we model eye fixations on the landing page as a series of "split-second" decisions

8

(50-500 milliseconds) that depend on the shopper's valuations of the AOIs and the effort of making eye movements within a sequential sampling framework (Busemeyer and Townsend 1993; Gluth, Rieskamp, and Büchel 2013; Otter et al. 2008; Ratcliff and Smith 2004). We assume that shoppers acquire information incrementally by sampling small amounts of information with each eye fixation (Krajbich, Armel, and Rangel 2010; Towal, Mormann, and Koch 2013). When the attractiveness of an AOI exceeds a decision threshold, the shopper clicks on a link. This baseline model is able to capture the heterogeneity of landing page search characteristics across shoppers, specifically in terms of search length (i.e., total number of fixations) and the fixation pathway (i.e., the series of fixations and the clicked link).

Second, we demonstrate that goodness-of-fit and in-sample predictions can be improved with model extensions that capture cognitive processes described in prior work on "boundedly rational" decision making (Gigerenzer and Goldstein 1996; Kahneman 2003; Simon 1955). We allow individuals to be one-step forward-looking when making fixation decisions by taking into account the probability of subsequently clicking on an AOI, either when deciding which AOI to fixate on or when making click/continue decisions (De los Santos and Koulayev 2017; Gabaix et al. 2006; Yang, Toubia, and de Jong 2015). We also allow individuals to change their decision thresholds as a function of search length or based on the value of AOIs seen so far, which is consistent with explore-exploit strategies of information search (Dellaert and Haübl 2012; Huang and Hutchinson 2013; March 1991).

Third, we assess the optimality of the estimated search process for each person. Note that we do not assume in our model that shoppers search optimally, and in fact the decisions made by shoppers in our data are more complicated that those assumed in standard analyses of optimal sequential search (e.g., Weitzman 1979; Lippman and McCardle 1991; see Section 2.4. for extended discussion). Instead, bounded rationality is tested by using parametric sensitivity analyses to estimate a "rationalizing exchange rate" for the costs and benefits of eye fixations. This allows us to assess whether shoppers were close to optimal search; for

example, could a shopper have clicked on a better AOI by searching longer with a higher decision threshold, or gotten a similar outcome by searching less with a lower decision threshold.

Fourth, one question of managerial interest is whether retailers can change the layout of the landing page to directly influence how long shoppers spend on the page, thereby affecting downstream behaviors. Thus, using the model estimated with Experiment 1 data, we conducted a set of counterfactual simulations to predict shopper behaviors on the landing page if products were positioned in a different order based on individual-level preferences. In Experiment 2, using a different online retailer, we empirically tested the counterfactual predictions by including two experimental conditions where the ordering of products on the landing page was customized for each shopper based on her preferences. The results of this experiment were consistent with the counterfactual predictions.

The plan for the rest of this chapter is as follows. Section 2.2. gives an overview of our decision threshold model of eye fixation and click decisions on the landing page. Section 2.3. lays out the specification of the baseline model and extensions. Section 2.4. describes the setup and results from Experiment 1. Section 2.5. describes the setup and results from Experiment 2. Section 2.6. concludes with discussion.

2.2. Model Overview

We are interested in modeling the pathway of eye fixations on different AOIs on the landing page of an online store. We treat fixations as a sequence of split-second decisions leading up to a final decision to click on a link within an AOI to move on to a different part of the website. Path data (e.g., grocery shopping, eye-tracking, web browsing, etc.) can be used by marketers to understand the dynamic decision processes of consumers as they interact with their virtual or physical shopping environment (see Hui, Fader, and Bradlow 2009 for an overview).

2.2.1. Sequential Sampling Models of Visual Attention and Information Search

We use eye fixations on specific AOIs as a measure of visual search on the landing page. After each fixation on an AOI, we assume that shoppers accumulate information and update their beliefs about the attractiveness of the fixated AOI through a sequential sampling process (Otter et al. 2008; Ratcliff and Smith 2004; Townsend and Ashby 1983). In our models, each fixation is followed by an implicit click/continue decision and the shopper clicks on an AOI once the attractiveness of one of the AOIs exceeds a decision threshold. Hui, Bradlow, and Fader (2009) use a similar approach to model how shoppers move through different areas of a physical store, with each visit to a specific area or "node" followed by a checkout/continue decision, and the attractiveness of different nodes updated after each visit.

Sequential sampling models have been shown to be psychologically and neurologically plausible accounts of how consumers make fast decisions (split-second or within a few seconds; see Otter et al. 2008 for a review). Random walk models like Decision Field Theory (Busemeyer and Townsend 1993), diffusion models (Krajbich, Armel, and Rangel 2010; Ratcliff and Tuerlinckx 2002; Satomura, Wedel, and Pieters 2014), and Poisson race models (Huang and Hutchinson 2008; Van Zandt, Colonius, and Proctor 2000) all fall under the umbrella of sequential sampling models. Treating the decision process as a stochastic accumulation of information can explain the differences in reaction times for simple versus complex choices, speed-accuracy tradeoffs, memory retrieval, and serial position effects. Branco, Sun, and Villas-Boas (2012) describe a similar model for product search (albeit across a much longer time horizon) in which product knowledge is accumulated over time until it reaches a threshold for purchase or no purchase.

2.2.2. Eye Movements as Split-Second Decisions

We extend upon traditional sequential sampling models by assuming that consumers' latent evolving preferences can be measured explicitly via eye movements and that shoppers make active split-second fixation decisions based on the attractiveness of the AOIs and the

effort of making eye movements towards the AOIs. In contrast, a passive accumulation process is assumed by most sequential sampling models. Prior work has treated fixation pathways as dependent on the value, position, and visual salience of options (Chandon et al. 2007; Towal, Mormann, and Koch 2013), as following specific patterns within latent states (Liechty, Pieters, and Wedel 2003; Shi, Wedel, and Pieters 2013; Stüttgen, Boatwright, and Monroe 2012; Wedel, Pieters, and Liechty 2008), and as a one-step forward-looking utility maximization process (Yang, Toubia, and de Jong 2015).

In our model, each fixation event occurs in two stages. In Stage 1, the shopper chooses an AOI to fixate on. In Stage 2, she decides whether or not the attractiveness of the currently fixated AOI exceeds a decision threshold. If it does, she clicks on the link within the AOI; if not, she moves on to the next fixation, which may be on the same or a different AOI. In the baseline model, the threshold is static and the probability of fixating on an AOI depends on its attractiveness and the effort of making the eye movement towards it.

2.2.3. Boundedly Rational Model Extensions

In addition to the baseline model, where the decision threshold is estimated as a single static heterogeneous non-AOI specific parameter, we also test extensions of the model that are based on empirical and theoretical findings about boundedly rational search processes.

First, we allow individuals to be one-step forward-looking, either when making their Stage 1 fixation decisions or their Stage 2 click/continue decisions. In the Stage 1 extension, when shoppers decide which AOI to fixate on, they may take into consideration the probability of actually clicking on the AOIs. In the Stage 2 extension, when shoppers consider whether to click on the currently fixated AOI or to continue search, they may take into consideration the future expected attractiveness of the clicked AOI if they were to make one *more* fixation, under the assumption that this fixation would be their last opportunity to search.

Gabaix et al. (2006) demonstrated that this type of limited time horizon forward-looking model predicted a sequence of choices that more closely resembled how lab participants

behaved during a product search task, compared to the sequence predicted by infinite time horizon index strategies (Weitzman 1979; Gittins 1979). Yang, Toubia, and de Jong (2015) applied this approach to visual search within a conjoint experiment. Additionally, since shoppers in our experiments can fixate on AOIs as many times as they want in any order, and accumulate information with each fixation, modeling the full state space using a structural dynamic programming approach is infeasible. De los Santos and Koulayev (2017) faced a similar challenge in modeling consumer choices of search refinement tools during online hotel search and also implemented a one-step lookahead approach.

Second, we allow the decision threshold to change after each fixation. Lab participants engaged in stylized search tasks often exhibit declining thresholds (Adam 2001; Brown, Flinn, and Schotter 2011; Koulayev 2014), which may be explained by focus on elapsed time or rising search costs, but also represents a reasonable heuristic that more time spent searching implies lower values of available options. Also, consumers may decrease their decision threshold when shifting from explore mode to exploit mode (Huang and Hutchinson 2013; March 1991), or from search mode to choice mode (Dellaert and Häubl 2012). Thus, we test the general phenomenon of declining thresholds by estimating the slope on the decision threshold across fixations; a negative slope corresponds to a decreasing threshold and therefore an increasing probability of clicking on a link within an AOI to leave the landing page.

We also allow the decision threshold to change based on the attractiveness of AOIs seen so far to test a specific set of theories where individuals are influenced by recently seen options. This includes the "bounce rules" described by Hey (1982, 1987) where individuals establish a reference point by initially inspecting a small set of options, as well as the local behavioral influences modeled by Häubl, Dellaert, and Donkers (2010), where individuals respond excessively to the attractiveness of current and recent options, as well as the contrast between them.

2.2.4. Relationship to Optimal Search

The decision threshold model that we propose shares similarities and differences with classic problems of sequential search. Morgan and Manning (1985) model search as multiple draws from a single distribution, with each draw corresponding to a realized outcome that the individual can accept or reject. Similarly, our shoppers can make multiple samples of the same option to gather information (i.e., by re-fixating on an AOI). However, our model also needs to account for shoppers searching across multiple *options*, which is more similar to the Weitzman (1979) model of sequential search with recall. There, the optimal solution is characterized by the reservation values of each option, defined as the value of the "sure thing" that would equate the cost and benefit of search. Shoppers search options in descending order of reservation values and learn the "true" or "realized" value of each option with a single sample. They stop when there are no more options with reservation values that exceed the best option on-hand. This solution has been popular within the marketing literature where there are many settings with multiple brands or firms (Chen and Yao 2017; Honka and Chintagunta 2016; Kim, Albuquerque, and Bronnenberg 2010). However, it does not allow for multiple *samples* of the same option.

To summarize, our model assumes that shoppers examine the options on the landing page based on their attractiveness, with four key distinctions from classical sequential search problems. First, our model assumes that shoppers sample from multiple options, and can revisit these options multiple times to gather more information. Second, shoppers have a single decision threshold for all options that determines when to end search on the landing page, rather than separate reservation utilities for each option that are computed before search begins. Third, there is path dependence such that both the physical position and perceived attractiveness of searchable options factors into the probability of those options being searched, and these values change with each fixation. Fourth, because of the nature of eye movements and clicks, shoppers only click on options they just searched (or physically nearby options), and so the last maximization step across all options assumed within optimal

sequential search models is potentially violated.

Based on these four distinctions, generalized optimal search solutions that allow for multiple searches with learning (Lippman and McCardle 1991; Koulayev 2014) and even very recent work that employs numerical approximations that allow for non-consecutive revisits (Chick and Frazier 2012; Ursu, Wang, and Chintagunta 2017) are also not mappable to the specific decisions of shoppers in our setting. Moreover, the purpose of our model is not to assume that shoppers are optimal, but rather to use a boundedly rational framework that allows us to assess how close to optimally shoppers were searching.

2.3. Model Specification and Estimation

Let $i = 1, ..., I$ denote shoppers and let $m = 1, ..., M_i$ denote shopper i's sequence of fixations. Each fixation m is the result of a two-stage decision process. In Stage 1, the shopper selects which AOI_j, for $j \in \{1, ..., J\}$, to fixate on with probability $P_{ijF}[m]$, which depends on both the attractiveness of the AOI, given by $A_{ij}[m]$, and the physical and mental effort of making the eye movement from the previously fixated AOI, given by $E_{ij}[m]$.

In Stage 2, the shopper decides whether to click on the currently fixated AOI, which depends on the attractiveness of the AOI relative to the decision threshold (note that effort is assumed to be zero because the AOI is currently fixated). The probability of clicking on the currently fixated AOI_j is given by $P_{ijC}[m]$, while the probability of continuing search on the landing page is given by $1 - P_{ijC}[m]$. Note that during Stage 2, individuals base their continue/click decisions on the AOI_j, that they chose to fixate on during Stage 1. Intuitively, it makes sense that shoppers would only click on an AOI they had just fixated on instead of clicking without looking. Also in our data we did in fact find that all shoppers clicked on a currently fixated AOI.[1]

Therefore, the probability $P_{ij}[m]$ of observing each fixation m is simply the product of the

---

[1]In some contexts consumers may be able to direct visual and motor attention to different screen areas and end up clicking on adjacent or recently fixated AOIs. Our model formulation could easily be relaxed to include such situations.

Stage 1 and Stage 2 probabilities (i.e., probability of fixating on a specific AOI multiplied by the conditional probability of clicking on the link within the AOI, given the fixation on the AOI). Since shoppers choose to continue search in Stage 2 of fixations $m = 1, ..., M_i - 1$, and only click in Stage 2 of the last fixation $m = M_i$, we can express the fixation probabilities as follows:

$$P_{ij}[m] = \begin{cases} P_{ijF}(1 - P_{ijC}[m]), & \text{if } m = 1, ..., M_i - 1 \\ \\ P_{ijF}P_{ijC}[m], & \text{if } m = M_i \end{cases} \tag{2.1}$$

The fixation events for all shoppers can be captured by the following likelihood function, where $\boldsymbol{\theta}_i = \{\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i, \rho_i, \lambda_{1i}, \lambda_{2i}, \lambda_{3i}, \lambda_{4i}\}$ (described in detail in the subsequent sections) represents the vector of estimated parameters for shopper i.

$$\mathcal{L}(\boldsymbol{\theta}_i) = \prod_{m=1}^{M_i} P_{ij}[m] \tag{2.2}$$

In Sections 2.3.1.-2.3.3., we describe the specification of the baseline model with a static decision threshold. In Sections 2.3.4. and 2.3.5., we describe alternative specifications of the model to account for the boundedly rational forward-looking and threshold-changing components. Section 2.3.6. lays out the estimation procedure.

2.3.1. The Baseline Model: AOI Attractiveness and Effort

Each fixation event m depends on AOI attractiveness, $A_{ij}[m]$, and the effort of making the eye movement, $E_{ij}[m]$. $A_{ij}[m]$ consist of a linear combination of S features and $E_{ij}[m]$ consist of a linear combination of T features. X_i refers to the $J \times S \times M_i$ matrix of attractiveness feature values and $\boldsymbol{\beta}_i$ is the corresponding vector of coefficients (recall that J is the number of AOIs and M_i is the number of events). Z_i is the $J \times T \times M_i$ matrix of effort feature values and γ_i is its vector of coefficients. Thus, the attractiveness $A_{ij}[m]$ and

16

effort $E_{ij}[m]$ associated with AOI_j are given by the following:

$$A_{ij}[m] = \sum_{s=1}^{S} \beta_{is} x_{ijs}[m] \tag{2.3}$$

$$E_{ij}[m] = \sum_{t=1}^{T} \gamma_{it} z_{ijs}[m] \tag{2.4}$$

Since shoppers presumably have no knowledge of the attractiveness values of the AOIs before they fixate on them, we assume that X_i represents the shopper's *perceived* feature values and start at the same neutral point for all AOIs at $m = 1$. The perceived feature values are then updated to approach their *true* values (i.e., true preference, true price, etc.) as information is accumulated over time via fixations on the AOIs. Note that true values might be larger or smaller than the neutral starting value. If the shopper fixates on AOI_j at fixation m, then the perceived feature value x_{ijs} for AOI_j is updated according to the following expression, where x_{ijs}^* represents the true value of that feature:

$$x_{ijs}[m+1] = (1 - \alpha_s) x_{ijs}[m] + \alpha_s x_{ijs}^* \tag{2.5}$$

The AOI features in the matrices X_i and Z_i that we included in the estimation of AOI attractiveness $A_{ij}[m]$ and effort $E_{ij}[m]$, respectively, are provided in the context of Experiment 1 in Section 2.4.3.

2.3.2. Stage 1: AOI Fixation Decision

In Stage 1, the shopper decides which AOI to fixate on. The utility of fixating on AOI_j is given by $U_{ij}[m] = A_{ij}[m] + E_{ij}[m] + \epsilon_{ijU}$, and thus depends on both attractiveness and effort (see Equations 2.3 and 2.4). Assuming the error term ϵ_{ijU} follows the type-1 extreme value distribution, the probability of fixating on AOI_j is modeled as a multinomial logit:

$$P_{ijF}[m] = \frac{e^{A_{ij}[m]+E_{ij}[m]}}{\sum_{k=1}^{J} e^{A_{ik}[m]+E_{ik}[m]}} \tag{2.6}$$

The perceived attractiveness feature values of the fixated AOI are then updated according to Equation 2.5 after Stage 1, but before Stage 2.

2.3.3. Stage 2: Click vs. Continue Search on the Landing Page

In Stage 2, the shopper decides whether to continue search with another fixation or click on the AOI that was fixated on in Stage 1. Clicking on an AOI ends search on the landing page and brings the shopper to another page in the store. The shopper's estimate of the value of clicking on AOI_j at event m is $V_{ij}[m] = A_{ij}[m+1] + \epsilon_{iV}$. (Note that the attractiveness used here has been updated following Stage 1.)

The shopper's decision threshold is $V_{iR}[m] = R_i[m] + \epsilon_{iV}$ and represents the perceived value of continuing search on the landing page. This value is subjective and no assumptions about optimality are made. We will test the optimality of the empirically estimated decision threshold via simulation in Section 2.4.6.

The probability $P_{ijC}[m]$ that the shopper will click on AOI_j (given that it is currently fixated on) is modeled as a binary logit that represents the likelihood that $V_{ij}[m] > V_{iR}[m]$, assuming the error term ϵ_{iV} follows the type-1 extreme value distribution:

$$P_{ijC}[m] = \frac{e^{A_{ij}[m+1]}}{e^{A_{ij}[m+1]} + e^{R_i[m]}} \tag{2.7}$$

For each shopper we can estimate a decision threshold that defines the stopping rule for information search. In the baseline model, the decision threshold is assumed to be a static heterogeneous parameter:

$$R_i[m] = \rho_i \tag{2.8}$$

2.3.4. Forward-Looking Extensions

Here we introduce two extensions of the baseline model that represent boundedly rational, one-step forward-looking dynamics. First, individuals may be making "forward-looking

fixations," such that when choosing which AOI to fixate on in Stage 1, they take into account the probability of clicking on the AOI in Stage 2, given by $P_{ijC}[m]$. In other words, the utility of fixating on AOI_j, given by $U_{ij}[m] = A_{ij}[m] + E_{ij}[m] + \epsilon_{ijU}$ in the baseline model, is weighted by the term $W_{ij}[m] = P_{ijC}[m]^{\lambda_{1i}}$ and Equation 2.6 becomes the following:

$$P_{ijF}[m] = \frac{e^{(A_{ij}[m]+E_{ij}[m])*W_{ij}[m]}}{\sum_k e^{(A_{ik}[m]+E_{ik}[m])*W_{ik}[m]}} \tag{2.9}$$

The parameter $\lambda_{1i} \in [0,1]$ indicates the degree to which individuals weight the Stage 1 utilities for each AOI by the Stage 2 probabilities of clicking on that AOI. At $\lambda_{1i} = 0$, all AOIs utilities are equally weighted, as in the baseline model. For small values of λ_{1i} near zero, the exponential function is steeply curved and equates the highest probabilities at a value near one and strongly penalizes the lowest probabilities, similar to a consideration set function. At $\lambda_{1i} = 1$ the click probabilities are directly used as weights.

Alternatively, individuals may be making "forward-looking clicks," such that when choosing whether to click on the currently fixated AOI in Stage 2, they take into account the expected attractiveness of the AOI they would click on if they were to make one more fixation. Equation 2.8 now becomes:

$$R_i[m] = \rho_i + \lambda_{2i}\left(\sum_{j\in\text{Seen}} P_{ijF}[m+1]A_{ij}[m+2]) + \sum_{k\notin\text{Seen}} P_{ikF}[m+1]A_{ik}[m+1]\right) \tag{2.10}$$

The terms in the parentheses represent the expected attractiveness of the clicked AOI if the individual were to make one more fixation at $m+1$ by taking the sum of the future attractiveness values of the AOIs, weighted by the probability of fixating on the AOIs, given by $P_{ijF}[m+1]$. Note that for each AOI_j that the shopper has already "Seen" (i.e., at least one fixation), the future attractiveness of the AOI, given by $A_{ij}[m+2]$, can be calculated by extrapolating the updating process described in Equation 5. However, we assume that the future attractiveness of each unseen AOI_k is *not* extrapolated and remains at $A_{ik}[m+1]$ (i.e., the neutral starting value) because the shopper has no information on

which to base an extrapolation. The parameter λ_{2i} indicates the degree to which individuals take into account the future value of making one more fixation at $m+1$ when making the click/continue decision at m.

Thus, with either model variation, individuals can range from being myopic to forward-looking in a boundedly rational way. With forward-looking fixations, the model allows the shopper to place additional value on AOIs that are likely to result in a click, over and above their feature-based attractiveness. With forward-looking clicks, individuals assume that the next fixation at $m+1$ will be the last fixation ending with a click. This approach is similar to those used by other models of choice behavior that assume consumers are boundedly rational and maximize expected utility one step into the future (De los Santos and Koulayev 2017; Gabaix 2006; Yang, Toubia, and de Jong 2015).

2.3.5. Threshold-Changing Extensions

Here we introduce two extensions of the baseline model that represent how individuals might change their decision threshold. First, we allow the decision threshold to change linearly with the number of fixations as an empirical test of the declining thresholds described in prior literature. For example, starting with a threshold that is "too high" and then declines to an "appropriate" level is consistent with explore-exploit strategies where shoppers expect to learn more about the values of the offered products and do not want to terminate search too early, as well as the meta-cognition that if one's threshold has not been reached, then perhaps one's expectations are too high (Dellaert and Haübl 2012; Huang and Hutchinson 2013; March 1991). Equation 2.8 now becomes the following:

$$R_i[m] = \rho_i + \lambda_{3i}m \tag{2.11}$$

Second, we allow the decision threshold to start at ρ_i at $m=1$ and approach the maximum attractiveness of all Seen AOIs, with a learning rate of $\lambda_{4i} \in [0,1]$. Thus, we test a parameterization of the bounce rules (Hey 1982, 1987) that are closely related to local information

influencing behavior (Häubl, Dellaert, and Donkers 2010), and Equation 2.8 becomes the following:

$$R_i[m] = (1 - \lambda_{4i})R_i[m-1] + \lambda_{4i}(\max_{j \in \text{Seen}} A_{ij}[m+1]) \tag{2.12}$$

Note that $R_i[m]$ is a mixture of the threshold value at the *previous* fixation $m-1$, given by $R_i[m-1]$, and the maximum AOI attractiveness at the *next* fixation $m+1$, given by $A_{ij}[m+1]$, because AOI attractiveness is updated after Stage 1, before the Stage 2 decision that involves the decision threshold.

2.3.6. Estimation

The main goal of our estimation procedure is to compare the performance of the baseline static decision threshold model to extensions that incorporate forward-looking and/or threshold-changing dynamics. Table 1 provides a summary of the parameters and utility specifications of the four model extensions relative to the baseline model.

Table 1: Summary of forward-looking and threshold-changing model extensions

Model	Add	Fixation Utility	Decision Threshold
Baseline	N/A	$A_{ij}[m] + E_{ij}[m]$	ρ_i
Forward-Looking			
Fixations	λ_{1i}	$(A_{ij}[m] + E_{ij}[m]) * P_{ijC}[m]^{\lambda_{1i}}$	ρ_i
Clicks	λ_{2i}	$A_{ij}[m] + E_{ij}[m]$	$\rho_i + \lambda_{2i}\left(\sum_{j \in \text{Seen}} P_{ijF}[m+1]A_{ij}[m+2]) + \sum_{k \notin \text{Seen}} P_{ikF}[m+1]A_{ik}[m+1]\right)$
Threshold-Changing			
Empirical Pattern	λ_{3i}	$A_{ij}[m] + E_{ij}[m]$	$\rho_i + \lambda_{3i}m$
Bounce Rule	λ_{4i}	$A_{ij}[m] + E_{ij}[m]$	$(1 - \lambda_{4i})R_i[m-1] + \lambda_{4i}(\max_{j \in \text{Seen}} A_{ij}[m])$

To capture the heterogeneity across shoppers, we estimated our model using a standard hierarchical Bayes procedure (Gelman et al. 2014). For the first-stage priors, we assumed

the set of parameters $\boldsymbol{\theta}_i = \{\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i, \rho_i, \lambda'_{1i}, \lambda_{2i}, \lambda_{3i}, \lambda'_{4i}\}$ followed a multivariate normal distribution with mean μ and precision Ω. (Note that λ'_{1i} and λ'_{4i} were logit-transformed versions of λ_{1i} and λ_{4i}, whose values were thereby constrained to fall between between 0 and 1.) For the second-stage priors, we assumed that μ followed a conjugate multivariate normal distribution with mean μ_0 (a vector of zeros) and precision Ω_0 (an identity matrix), and that Ω followed a conjugate Wishart distribution with degrees of freedom r (the number of parameters in θ_i plus 3) and inverse scale matrix Λ (where Λ^{-1} was an identity matrix). Because the attractiveness feature coefficients β_i and the attractiveness updating rates α_i are combined multiplicatively (see Equations 2.3 and 2.5), it is empirically difficult to estimate them simultaneously; therefore a grid search was used to obtain population level estimates of α (see Section 2.4.3. for details).

We simulated the parameters using a Markov chain Monte Carlo (MCMC) sampler in the programming language R. For each model, we ran three MCMC chains from different starting values for 100,000 iterations each. We used the first 90,000 iterations as burn-in and checked for convergence by determining that the Gelman-Rubin convergence statistic was less than 1.2 for all parameters (Gelman and Rubin 1992). After thinning the chains to reduce auto-correlation, we were left with 3,000 posterior samples for each parameter. To check for empirical identification, we estimated each version of our model on simulated data to ensure parameter recovery (see Appendix A.1. for details).

We compared the relative performance of the models by calculating the deviance information criterion (DIC), which is a Bayesian version of the Akaike information criterion (AIC). Like the AIC, the DIC includes a penalty for the number of parameters by using the deviance from the mean of the posterior samples, and may increase when parameters are added to the baseline model (Gelman et al. 2014).[2]

We also compared the in-sample model fit with posterior predictive checks by using the

---

[2]We use DIC $= -2\log p(y|\hat{\theta}) + 2p_{DIC}$ where the first term represents the likelihood of the data y given the posterior mean $\hat{\theta}$ and the second term $p_{DIC} = 2(\log p(y|\hat{\theta}) - \frac{1}{S}\sum_{s=1}^{S}\log p(y|\theta^s))$ represents the deviance of the posterior samples θ^s from the mean.

posterior samples to simulate each shopper's fixation pathways (3,000 simulations per shopper) and looked at how closely simulated descriptive statistics matched observed descriptive statistics. Specifically, we assess whether our model captures search length (as measured by the total number of fixations before a click), which AOIs are fixated on and their frequency, as well as which AOI is ultimately clicked on.

2.4. Experiment 1: Comparison of Baseline Model and Extensions

The purpose of Experiment 1 was to collect data for modeling the eye movements of shoppers on the landing page of an online store. We compare the fit of five models (i.e., the baseline model, two forward-looking extensions, and two threshold learning extensions; see Table 1) using the DIC and posterior predictive checks. All analyses for Experiment 1 were replicated with the data collected in Experiment 2.

2.4.1. Experimental Design

We conducted an incentive-compatible eye-tracking experiment in which female undergraduates at a large university were paid $10 to browse the online store of American Apparel, plus a 1 in 20 chance of winning a randomly chosen item from their cart up to a maximum price of $75.

When shopping for clothing, consumers rely heavily on visual features, so eye-tracking was well-suited for capturing the accumulation of product information in this setting. Participants in our study were informed that their eye movements and mouse activity would be unobtrusively recorded, and were calibrated to the eye-trackers using the standard procedures recommended for SensoMotoric Instruments eye-tracking systems.

Participants were told that they would be shopping for a new female wearable top, and randomly assigned either a utilitarian shopping goal (i.e., something to wear for work or class) or a hedonic shopping goal (i.e., something to wear for a party; see exact instructions in Appendix A.2.). The purpose of this manipulation was to induce exogenous variations in

search patterns and test whether or not these patterns could be recovered by our models. In particular, we expected that hedonic shoppers would be more stimulus-driven rather than goal-driven compared to utilitarian shoppers (Corbetta and Shulman 2002; Holbrook and Hirschman 1982) and thus explore the landing page more extensively.

Participants were then brought to the landing page of American Apparel, and given 5 minutes to shop. The virtual landing page (see Figure 2) was a slightly modified version of the actual store website's "New Arrivals" page and remained unchanged across the data collection period. The first row of products was visible upon arrival to the website, while the remaining rows could be seen by scrolling down with the computer mouse.

After clicking on a link on the landing page, shoppers could not return to the landing page, but were free to visit any other part of the website and could select items by adding them to their virtual carts. Eye-tracking data and mouse clicks were collected throughout the entire shopping trip. We restricted data collection to a 3-week period so that the product selection and prices would not change substantially. In this chapter, only data from the landing page are analyzed.

After 3 months, participants were contacted by email to complete a follow-up survey for an additional $10 and a chance to win a $100 American Apparel gift card. They rated how much they liked the 30 products from the landing page, the products from their virtual cart, and a random sample of 30-35 products from other participants' carts, presented in random order without price information. Measuring individual preferences for products either before or after the main experimental task is common within research employing psychophysical or neurophysiological methods in order to study individual-level preferences (Krajbich, Armel, and Rangel 2010; Reutskaja et al. 2011). Out of the 84 total participants, 8 were excluded from the analyses because of eye-tracker calibration issues, and 10 were excluded for not completing the follow-up survey, leaving 67 shoppers total.

Table 2 contains the descriptive statistics for several features of the shopping trip. The statistics from the landing page fixation pathways include the total number of fixations, the number of fixations on Product AOIs, the total number of unique AOIs seen, the lowest row of Product AOIs seen, and the percentage of shoppers who clicked on the Product AOI.

Table 2: Descriptive statistics of shopping trips, reporting means and standard deviations in parentheses.

Shopping Stage	Activity	Exp 1	Exp 2 Random	Exp 2 Worst-First	Exp 2 Best-First
Landing Page	Total fix	29.39 (24.69)	94.49 (82.80)	109.38 (59.83)	76.61 (57.03)
	Product AOI fix	23.09 (24.88)	88.61 (80.52)	104.18 (60.24)	74.65 (56.07)
	AOIs seen	10.06 (9.05)	22.39 (13.70)	28.31 (11.63)	20.13 (14.21)
	Lowest row	2.81 (2.17)	6.39 (3.79)	7.56 (2.91)	5.83 (3.82)
Exit Landing Page	% clicked on Product AOI	19.40%	92.68%	93.33%	89.13%
Final Shopping Cart	# of products	2.17 (1.18)	2.66 (1.93)	2.78 (2.15)	3.0 (1.52)
	Total value	$96.04 (62.59)	$55.89 (51.12)	$56.91 (52.80)	$55.51 (34.51)

In Experiment 1, shoppers fixated on between a quarter to a third of the Product AOIs and got about a third to halfway down the page. Most shoppers ended up clicking on the Category AOI rather than on one of the Product AOIs, and found on average 2 products to put in their shopping carts, valued at around $90 total. In addition, we found that shoppers who were randomly assigned to the hedonic shopping goal condition made significantly more fixations on the last three rows of product AOIs ($M = 8.39$, $SD = 11.60$) compared to shoppers assigned to the utilitarian goal condition ($M = 3.19$, $SD = 7.08$, $t(48.10) = -2.17$, $p = .04$), and we verify in Section 2.4.4 that our model is able to capture this difference.

Finally, although the focus of our model is capturing the fixations and clicks on the landing page, we find that landing page decisions are predictive of subsequent search. Specifically, shoppers who spent more time searching on the landing page subsequently explored less of

the rest of the store's website, with about one less unique webpage per 20 additional eye fixations ($t = -13.10, p < 0.01, R^2 = 0.11$), and put \$0.58 less in their shopping carts per additional fixation ($t = -1.95, p = 0.06, R^2 = 0.04$).

2.4.3. Baseline Model: Features of AOI Attractiveness and Effort

Table 3 summarizes the features we included in the calculation of AOI attractiveness $A_{ij}[m]$: product preferences, product prices, whether the product was out-of-budget, and an intercept for the Category AOI. The shopper's perceived value of each attractiveness feature s was initialized at a starting value and updated by learning rate α_s (see Equation 2.5).

Table 3: Summary of features comprising AOI attractiveness and effort

Feature Type	Name	Description	Range of Values
Attractiveness $A_{ij}[m]$	Preference	Self-reported product liking ratings	-3 to 3
	Price	Product prices	-2.50 to 2.30
	Out-of-Budget	Was the product in the given budget	0 or 1
	Category	AOI contained category links	0 or 1
Effort $E_{ij}[m]$	Scroll	Requires initial scrolling to be visible	-1 or +1
	Distance	Near vs. far fixations	-1 or +1
	Unseen	Has not been fixated on yet	-1 or +1

True product preferences were measured by using the 7-point liking ratings (1="Don't Like At All", 7 = "Like Very Much") from each shopper. These ratings were centered at 0 and the preference for each product started at 0 and was gradually updated to its true value (ranging from -3 to +3) as shoppers made more fixations to gather information about the product. On average, among all the products that shoppers rated in the follow-up survey, only 8% were liked more than the products in their own carts. This suggests that shoppers were pretty good at finding products they liked and lends validity to the liking rating as a measure of stable product preferences.

For products that were within the $75 budget, prices were mean-centered and scaled by dividing by 10, with 0 as the starting value. Although higher prices are typically viewed as undesirable, in our experiment shoppers were incentivized to select products that were close to $75, without going over, so we expect price to contribute positively to Product AOI attractiveness. (Note that in Experiment 2 we use an alternative incentive scheme that predicts a negative price coefficient.) We also included a variable that indicated whether products were out-of-budget, which began at 0 for all products and approached 1 for products that cost over $75. The out-of-budget feature captures non-linearities in price effects for products that participants would not have been able to purchase.

Since the Category AOI did not possess any inherent measurable features, we included a feature that started at 0 and approached 1 for fixations on the Category AOI. This Category feature is scaled by the coefficient estimate and is the only feature that contributes to the attractiveness of the Category AOI; all other feature values are fixed at 0 for the Category AOI. The Category feature is fixed at 0 for the Product AOIs.

The price, out-of-budget, and category features are objective and can be learned by the shopper "at-a-glance." Thus, we set the updating rates for these features to 1, so their perceived values are instantly updated to their true values after a single fixation. For the preference feature, we used a grid search of $\alpha_s \in \{.1, .3, .5, .7, .9\}$ to estimate the updating rate, since it is empirically difficult to disentangle low/high updating rates from weak/strong preferences. Based on this grid search, we selected $\alpha_s = .1$ as the updating rate for the preference feature (see Appendix A.3. for details; Yang, Toubia, and de Jong 2015 employ a similar approach). For example, suppose individual i fixates on AOI_j at event $m = 1$, and that AOI_j has a *true* preference value of $x^*_{ijs} = 2$. Then at the next fixation $m = 2$, the *perceived* preference value of AOI j is updated according to Equation 5: $x_{ijs}[2] = (1 - \alpha_s)x_{ijs}[1] + \alpha_s x^*_{ijs} = (1 - .1) * 0 + .1 * 3 = 0.3$.

Table 3 also summarizes the features that contributed to AOI effort $E_{ij}[m]$: Scroll, Distance, and Unseen. The Scroll feature separated AOIs that required the shopper to scroll down

with a computer mouse from those that were immediately available upon entering the landing page (including the first row of Product AOIs and the Category AOI). Distance indicated that the AOI was positioned non-adjacent to the previously fixated AOI, thus requiring "far" rather than "near" eye movements. Thus, Distance captured differences in fixation probabilities for AOIs that are in the more central versus more peripheral areas of the shopper's field of view, taking into account the changes in the screen display due to mouse scrolls.[3] Unseen indicated that the AOI had *not* been fixated on previously; new AOIs were expected to take more cognitive effort to fixate on. We used contrast coding (-1 and +1) to refer to the two levels of the effort variables.

2.4.4. Model Comparison: Goodness-of-fit and Posterior Predictive Checks

Table 4 panel A compares the fit of the baseline model and all model extensions for Experiment 1, based on the DIC and a series of posterior predictive checks. The DIC is an aggregate measure of model fit, with a better fit corresponding to a lower DIC. To generate posterior predictive checks, for each individual, we took each of their 3,000 posterior samples and simulated a fixation pathway ending in a click, and then looked at whether the characteristics of the simulated pathways matched those of the observed pathways in terms of the total number of fixations, which AOIs were fixated, and which AOI was clicked.

Total number of fixations and fixation share correlations describe the fixation patterns of shoppers. To compute the root mean squared error (RMSE) for the total number of fixations, we took the median number of fixations across each individual's simulated pathways, and compared these medians to the observed number of fixations. To compute the fixation share correlations, for each shopper, we computed the AOI fixation shares (i.e., percentage of times each AOI was fixated on) for each simulated pathway and calculated the correlation with the fixation shares of the observed pathway. We averaged across simulations within-shopper, and Table 4 panel A reports the average across shoppers.

---

[3]See Appendix A.4. for an alternative specification of the Distance variable based on Euclidean pixel distance, which did not perform as well as our binary Distance variable.

Table 4: Comparison of model fit statistics

A. Experiment 1

Model	DIC	# Fix (RMSE) Actual = 29.39	Fix Corr	Prod Click Ref = .69	AOI Click Ref = .66
Baseline	8803	26.24 (13.46)	.41	.6571	.5283
Forward-Looking					
Fixations	8778	25.66 (13.36)	.45	.7036	.5629
Clicks	8832	26.24 (13.96)	.42	.6495	.5197
Threshold-Changing					
Empirical Pattern	8736	29.16 (7.86)	.41	.6458	.5326
Bounce Rule	8780	26.16 (13.17)	.41	.6521	.5237
Full Model	8731	28.73 (7.30)	.44	.6739	.5442

B. Experiment 2

Model	DIC	# Fix (RMSE) Obs = 94.49	Fix Corr	Prod Click Ref = .86	AOI Click Ref = .06
Baseline	17486	78.61 (56.02)	.20	.8973	.1642
Forward-Looking					
Fixations	17387	84.51 (53.99)	.21	.9006	.1737
Clicks	17492	93.76 (57.62)	.21	.9038	.1764
Threshold-Changing					
Empirical Pattern	17465	87.27 (25.21)	.20	.9033	.1650
Bounce Rule	17486	77.98 (48.40)	.20	.9004	.1626
Full Model	17383	88.39 (33.83)	.21	.9077	.1724

We also look at how well the simulated pathways were able to predict the AOIs that shoppers ultimately clicked on. We calculated the percentage of times the simulations correctly predicted either a Product vs. Category click, as well as the percentage of times the simulations correctly predicted the exact AOI (across all 31 AOIs) that the shopper clicked on. If we simply treat the click on either a Product or Category AOI as a random draw with equal probabilities, then the average percentage of correct simulated clicks is well above chance (50%). Similarly, the average percentage of clicks out of all 31 AOIs is also well above chance (3.22%).

However, a more stringent reference point would be to assume a zero-order, multinomial stochastic process with the probabilities of each choice option estimated by the marginal frequencies observed in the data. This is done by by taking the sum of squared marginal frequencies of clicking on the AOIs. This gives us a reference of 69% for the Product vs. Category click, and a reference of 66% for the click out of all 31 AOIs. In Experiment 1, 80.6% of shoppers actually clicked on the Category AOI, so the click reference is very high and our model actually falls short of predicting the Category click. In Experiment 2, we see that our model is more suited to situations where there is a more even spread across the large number of Product AOI choices.

Table 4 panel A shows that between the two forward-looking model extensions, forward-looking fixations improves the DIC, the fixation correlation, as well as the percentage of correct clicks, while forward-looking clicks only slightly improves the fixation correlation. Both threshold-updating extensions improve the DIC and the RMSE of total fixations, with the empirical pattern resulting in the most improvement. Figure 3 panel A plots how the decision threshold changes as an empirical pattern versus the bounce rule. Both versions result in decreasing decision thresholds, with the bounce rule able to explain some, but not all, of the decrease demonstrated with the empirical pattern.

In the last row, we tested a "Full Model" with forward-looking fixations (λ_1) and empirical threshold-changing (λ_3), which resulted in the best fit across all models in terms of DIC,

Figure 3: Comparison of decision thresholds for threshold-changing extensions
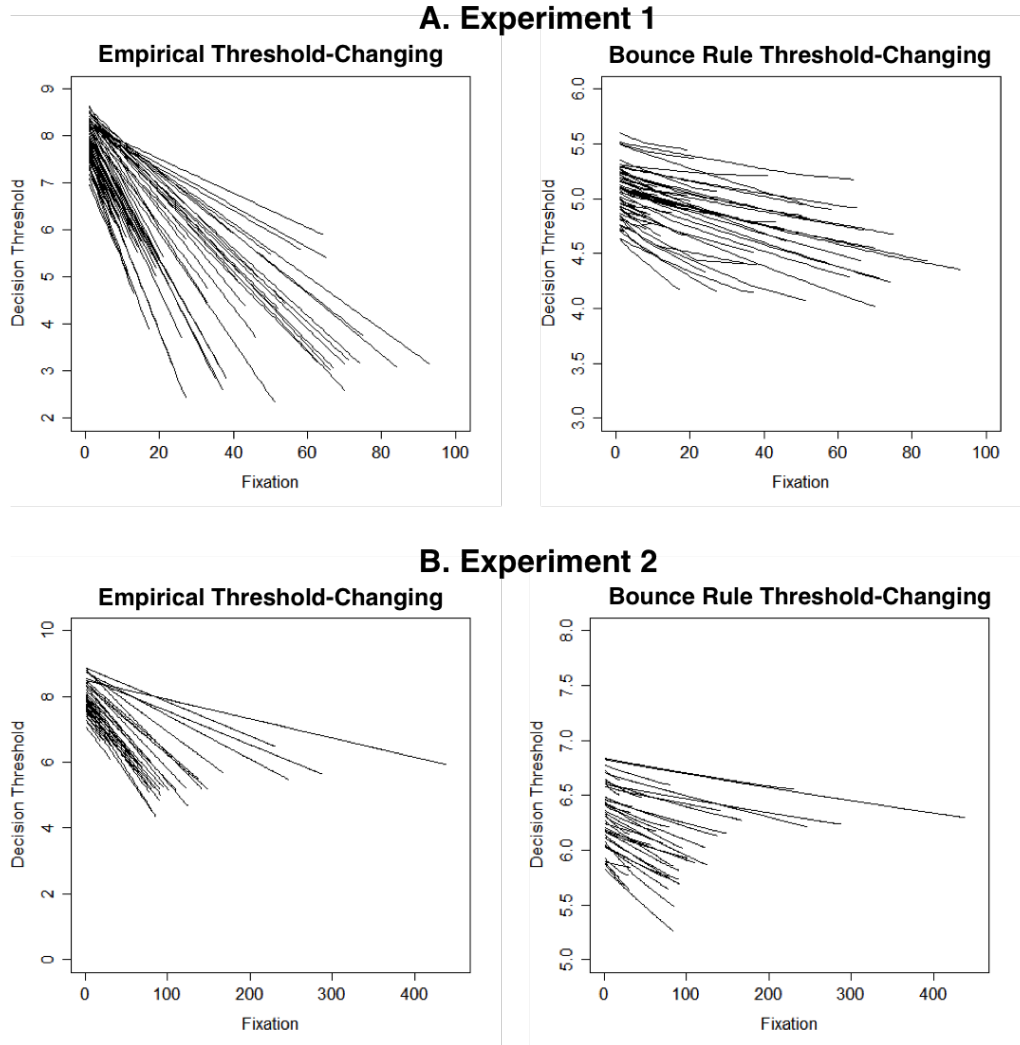
as well as improvements over the baseline model across all four posterior predictive checks. The forward-looking fixations component is built into the Stage 1 fixation decisions, so it makes sense that it improves predictions of which AOIs the shopper fixates on and which AOI she clicks on (since she is fixating on the correct AOIs). The empirical threshold-changing component is built into the Stage 2 continue/click decisions, so it makes sense that it improves predictions of when the shopper will end search with a click, in terms of the number of fixations.

In addition, Figure 4 panel A plots the observed versus simulated total number of fixations

31

for each individual, as well as the regression line through the points, which demonstrates that our model is able to capture the heterogeneity across shoppers in terms of search length, with most individuals lying close to the 45-degree reference line. We also overlaid the regression line for the baseline model to show that the static decision threshold model systematically underestimates the total number of fixations. This is corrected primarily by the empirical threshold-changing component of the model, which allows the decision threshold to start high so shoppers can gather information (and not end too soon, as with the static threshold in the baseline model), and gradually decrease as shoppers approach a click decision.

Figure 4: Observed versus simulated number of fixations for baseline vs. full model
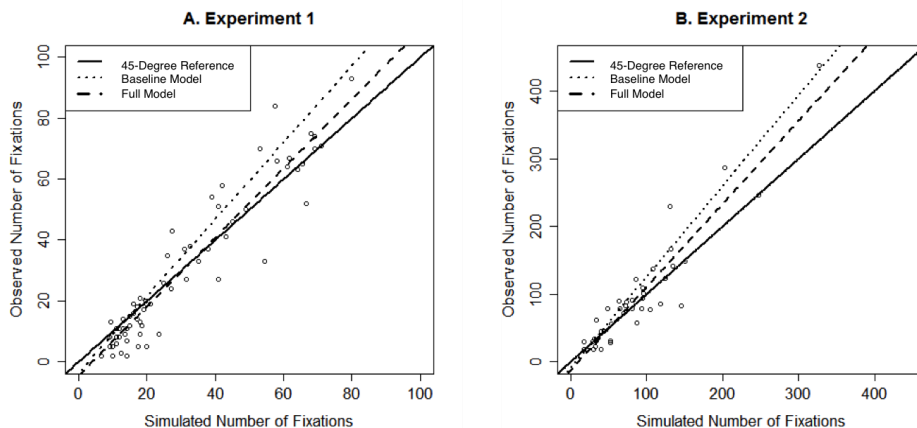
Moreover, our full model was also able to capture the differences in fixations on the bottom three rows that we observed between shoppers given a hedonic vs. utilitarian shopping goal. On average, our simulations resulted in 7.42 (SD = 10.21) fixations for hedonic shoppers and 4.39 (SD = 8.29) fixations for utilitarian shoppers, which closely matches the observed difference (although the simulated difference is not significant, likely due to shrinkage from the hierarchical Bayes estimation procedure). This demonstrates that our model is able to capture exogenously-induced differences in search length.

2.4.5. Parameter Estimation Results

Table 5 gives the parameter estimates of the full model, including the mean of the posterior samples μ, the 95% Credible Intervals (CI), and the mean of the posterior variance samples σ^2. The product preference, product price, and category features have positive coefficients, while the out-of-budget feature has a negative coefficient. The variables that represent search effort, including Scroll, Distance, and Unseen, have negative coefficients. Thus, the signs of all coefficients are as predicted, demonstrating that the model yields plausible results.

The forward-looking fixations term λ_1 captures how much shoppers weighed the probability of clicking on an AOI in Stage 2 when deciding where to fixate in Stage 1. The explore-exploit learning term λ_3 is negative and captures the decreasing decision threshold, as illustrated in Figure 3.

Table 5: Parameter estimation results for full model

Parameter	Experiment 1			Experiment 2		
	μ	95% CI	σ^2	μ	95% CI	σ^2
β_1 Preference	.98	[.45, 1.51]	1.69	.65	[.42, .87]	.43
β_2 Price	.43	[-.01, .51]	.54	-.11	[-.34, .12]	.46
β_3 Out-of-Budget	-.83	[-1.38, -.39]	1.25	N/A	N/A	N/A
β_4 Category	2.81	[2.14, 3.69]	3.65	1.94	[1.24, 2.68]	1.94
γ_1 Effort	-1.47	[-1.79, -1.17]	0.68	-3.75	[-4.60, -2.85]	2.36
γ_2 Distance	-1.74	[-2.40, -1.26]	2.12	-.88	[-1.45, -.40]	2.03
γ_3 Unseen	-.58	[-.93, -.23]	.86	-.84	[-1.44, -.32]	1.96
ρ Static Threshold	7.04	[5.81, 8.58]	1.42	7.42	[6.72, 8.19]	1.10
λ_1 Forward-Looking Fixations	.03	[.01, .06]	.000	.12	[.08, .16]	.001
λ_3 Explore/Exploit Updating	-.97	[-1.56, -.47]	.71	-.12	[-.20, -.04]	.01

Since we did not assume that shoppers were searching optimally, we were interested in assessing "how well" they searched, in terms of clicking on attractive AOIs. To do this, we independently varied the values of the constant term of the decision threshold ρ, the forward-looking fixations parameter λ_1, and the empirical threshold-changing parameter λ_3. We simulated fixation pathways for each shopper and looked at how the attractiveness of the clicked AOI and number of fixations changed with these alternative sets of parameter values.[4]

Figure 5 panel A plots the perceived and true attractiveness of the clicked AOI as open and closed circles, respectively, averaged across individuals. The dotted vertical line represents the point of 0 deviance form the original posterior parameter samples.

For all three parameters, we see that at 0 deviance, the attractiveness of the clicked AOIs are near their max values, indicating that shoppers seem to searching close to optimally, and there are not huge gains to changing the values of their decision thresholds. Very negative deviations of the constant term of the decision threshold ρ result in fewer fixations and clicking on AOIs that are low in attractiveness. Shoppers may click on slightly more attractive AOIs (and also shrink the difference between the perceived and true values) by either being more selective with a higher decision threshold ρ, or slowing the decrease of the decision threshold by increasing the empirical threshold-changing parameter λ_3. However, this comes at the cost of making more fixations.

Shoppers may also take the Stage 2 click probability more into consideration when making their fixation decisions in Stage 1 by increasing the forward-looking fixations term λ_1, which slightly increases the attractiveness of clicked AOIs up to a deviation of 0.1, at which point the attractiveness starts to decrease. One possible reason for this is that higher λ_1 values

---

[4]For λ_1, the deviance is on the logit-transformed parameter λ_1' in order to restrict λ_1 to values between 0 and 1. For large positive deviations of ρ and λ_3, we simulate "never-ending" search because the decision threshold is too large or increasing, so we artificially end a simulation if the number of fixations exceeds 500.

lead to more fixations on AOIs that are likely to be clicked on, and thus decreases the length of search and the possibility of finding more attractive AOIs.

Figure 5: Comparison of clicked AOI attractiveness with different parameter values

A. Experiment 1

B. Experiment 2
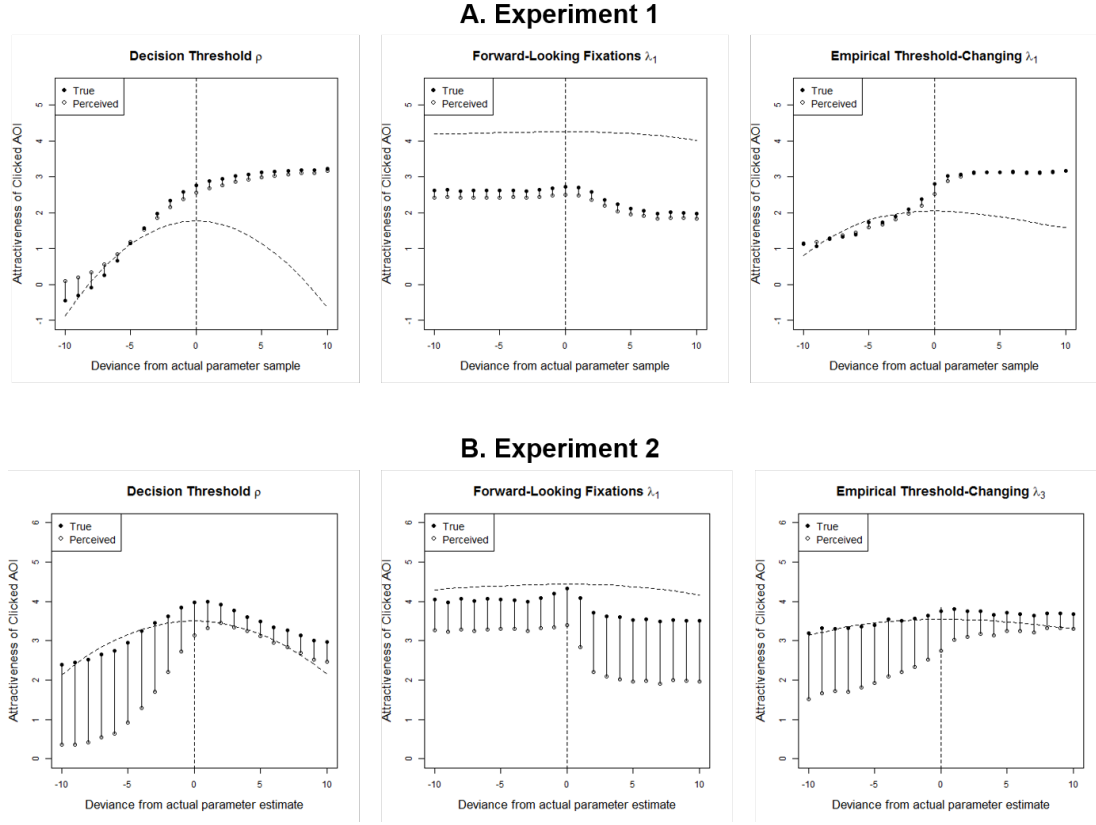
For each shopper, we rationalize their parameter estimates by solving for the "exchange rate" between attractiveness and total fixations for which the original estimated parameter (i.e., at zero deviance) is a local maximum for net attractiveness (i.e., attractiveness minus scaled fixations; see Appendix A.5. for details). The curved dotted lines in Figure 5 panel A represent the slope-adjusted difference between attractiveness and total fixations, which peaks at the dotted vertical line representing 0 deviance. For the decision threshold ρ, we obtained maxima for each shopper and an average exchange rate of 0.04, indicating that if each additional fixation is worth 0.04 attractiveness points, then net attractiveness is maximized at zero deviance. For the forward-looking fixations term λ_1, we obtained a

maxima for only 69% of shoppers with an average exchange rate of -0.06. For the empirical threshold-changing term λ_3, we obtained maxima for all shoppers, with an average exchange rate of 0.003.

2.4.7. Counterfactual Analysis: Preference-Based Product Ordering

Online retailers often provide shoppers with search refinement tools, such as the option to organize the products on a page by price or popularity. Retailers also use past browsing behavior to make recommendations to consumers or use morphing to customize the "look and feel" of a website (Hauser et al. 2009). De los Santos and Koulayev (2017) uses past observed search refinements to predict the performance of targeted hotel rankings in terms of increasing click-through rates (CTRs).

Similarly, we are interested in how search patterns would change if the products on the landing page had been ordered in a different way. Thus, we conducted a counterfactual analysis in which we simulated shoppers' patterns of fixations for conditions in which the 30 products on the page were ordered according to each individuals' self-reported liking ratings of the landing page products, either "worst-first" or "best-first".

Figure 6 panel A compares the results from simulations for the best-first and worst-first orderings to a random ordering. With the worst-first ordering, we predicted that shoppers would make more fixations ($M = 32.03$, $SD = 20.85$) compared to the best-first ordering ($M = 29.72$, $SD = 18.23$; paired-$t(68) = 2.56$, $p = 0.01$). However, we also predicted that shoppers in the worst-first ordering would click on less preferred AOIs ($M = 2.66$, $SD = 1.51$) compared to the best-first ordering ($M = 2.95$, $SD = 1.28$; paired-$t(68) = 2.31$, $p = 0.02$), based on the 7-point liking rating scales, despite the presence of the relatively attractive Category AOI in the same position for all orderings.

Figure 6: Counterfactual comparison of total number of fixations with best-first vs. worst-first product ordering
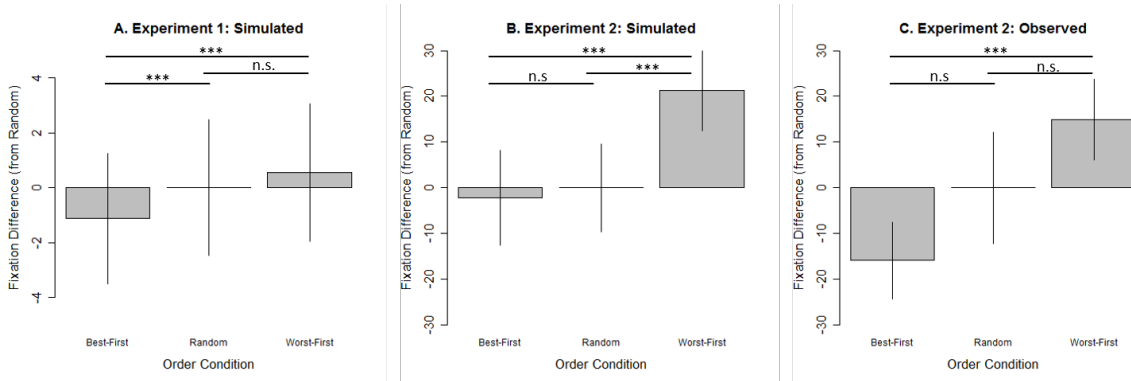
2.5. Experiment 2: Replication and Empirical Test of Product Ordering

The main purpose of Experiment 2 was to empirically test the predictions of our counterfactual simulations from Experiment 1 with worst-first and best-first product orderings. In addition, we verify that our results from Experiment 1 can be replicated in a different retail setting with a new sample of participants.

2.5.1. Experimental Design

We conducted a second eye-tracking experiment in which female undergraduates at a large university were paid $10 to browse the online store of the clothing retailer Forever 21. We required that all participants complete a pre-screening survey at least 5 days prior to the eye-tracking portion of the study. In this survey, participants were presented with the images of 60 products that were available for online purchase from Forever 21. Participants first rated their liking of each product on a 7-point Likert scale, and then ranked the products within each point rating to break ties, which gave us a full ranking of all 60 products from each individual. Among the 151 participants, 19 were excluded from the analysis because of eye-tracker calibration issues, which left 132 participants total for analysis.

During the eye-tracking experiment, participants were brought to a modified version of the landing page of Forever 21 and had 5 minutes to shop, again with the possibility that they

might win one of the items in their shopping cart up to $50. We lowered the "budget" from $75 in Experiment 1 to $50 in Experiment 2 because the Forever 21 products had lower prices. Participants were also told that if chosen for the lottery they would receive the chosen item plus the difference between the item price and the $50 budget in cash. Thus, in contrast to Experiment 1, we expect price to have a negative effect on AOI attractiveness.

The Forever 21 landing page consisted of 40 Product AOIs and 1 Category AOI, each with links to specific products or more general product categories (see Figure 7). The 40 products were a subset of the 60 products that they had rated and ranked in the pre-screening survey. Participants were again free to shop anywhere on the website once they left the landing page.

Figure 7: Layout of Forever 21 landing page for Experiment 2

Table 2 compares the descriptive statistics for shoppers in Experiments 1 and 2. We find that across all three product ordering conditions (random, worst-first, and best-first), shoppers in Experiment 2 searched longer on the landing page compared to shoppers in Experiment 1 by making more total fixations, and the vast majority clicked on one of the Product AOIs rather than on the Category AOI. This might have been because in Experiment 2, we used a different retailer (Forever21) that was more popular with participants in our sample, there were more products on the landing page, and shoppers also liked the set of products more, rating them more highly on the 7-point liking scales ($M = 3.8$, $SD = 0.8$) compared to the average ratings from Experiment 1 ($M = 3.4$, $SD = 0.8$; $t(134.9) = 3.2$, $p < 0.01$).

To further illustrate these differences, Figure 7 plots the average liking ratings across shoppers for the products ranked from lowest to highest in average liking rating. Note that there were 30 products in Experiment 1 and 40 products in Experiment 2. The distribution of liking ratings is generally higher among the Forever 21 products in Experiment 2 compared to the American Apparel products in Experiment 1, and this difference increases as the product rank decreases. This is consistent with the resulting greater number of total fixation on the landing page and greater number of clicks on Product AOIs in Experiment 2. These large differences in shopping behavior between the two experiments enhances the value of Experiment 2 as a test of external validity for our decision threshold model. Our model is able to capture the fixations and clicks on landing pages where most shoppers prefer (and end up clicking on) the Category AOI, as in Experiment 1, as well as those where most shoppers actually prefer the Product AOIs, as in Experiment 2.

In Experiment 2, we also replicated the finding from Experiment 1 that search patterns on the shopping page are predictive of downstream behaviors. Specifically, shoppers in Experiment 2 visited one less unique webpage per 55 additional fixations on the landing page ($t = -3.22, p < 0.01, R^2 = 0.07$), and put \$0.18 less in their shopping carts per additional fixation ($t = -3.27, p < 0.01, R^2 = 0.08$).

Figure 8: Comparison of average liking ratings for Experiments 1 and 2
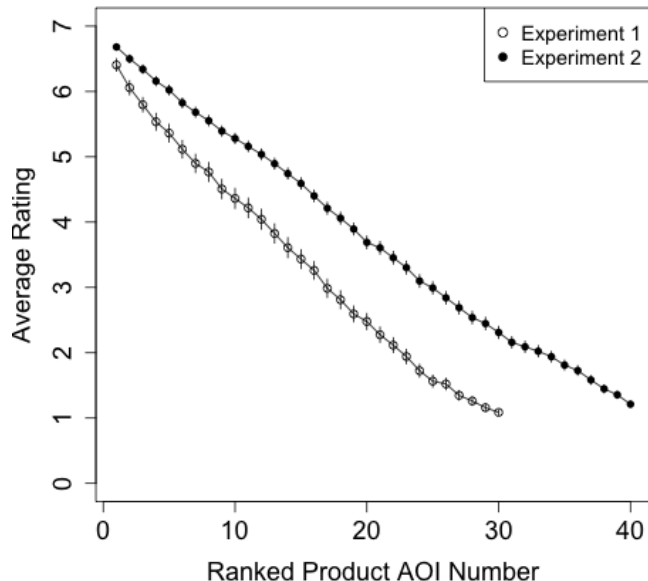
To replicate the model comparison and estimation from Experiment 1, we used the data from the 41 shoppers in the random ordering condition in Experiment 2 (we obtained similar results using data from the best-first and worst-first conditions; see Appendix A.6. for details). We again estimated the different variations to the model to compare to the baseline model, with results shown in Table 4 panel B. In general, the pattern of results is very similar to those observed for Experiment 1. Among the forward-looking extensions, we see that forward-looking fixations improves the model fit in terms of DIC and the percentage of correct clicks, while forward-looking clicks does not improve the DIC but does improve the percentage of correct clicks as well. Among the threshold-changing extensions, the empirical pattern improves the DIC and the RMSE of the total number of fixations, while the bounce rule does not improve the DIC but does improve the RMSE of the total number of fixations. Again, from Figure 3 panel B we see that both types of threshold-changing result in decreasing decision thresholds, and from Figure 4 panel B we see that the full model improves the slight overestimation of the number of fixations over the baseline model.

Table 5 shows the parameter estimation results for both experiments. We included the same variables for AOI attractiveness and effort as we did for the Experiment 1 data, except for the out-of-budget parameter because in Experiment 2 all the products presented on the landing page were within the $50 budget. We see that for most of the parameters, the sign of the posterior sample mean μ is the same across both experiments, with the exception of the price parameter, which is positive in Experiment 1. This difference was expected because we told participants that they could keep the difference between the product's price and the maximum budget price if they were selected for the lottery. However, in both experiments the price parameter was not significant (i.e., 0 is contained in the 95% Credible Interval), suggesting an effect of "playing with house money." Again, λ_1 reflects the degree to which participants were one-step forward-looking when making fixation decisions, while λ_3 is the rate at which the decision threshold decreases, which is considerably smaller for Experiment 2 compared to Experiment 1- consistent with the longer search length among shoppers in Experiment 2.

Again, we varied the parameter values of the static decision threshold ρ, the forward-looking fixations term λ_1, and the explore-exploit learning term λ_3. Figure 5 panel B shows that we again find that shoppers are searching close to optimally and there are not huge gains in AOI attractiveness to be had from deviating away from the estimated parameter values. Again, we are able to calculate the exchange rate of clicked AOI attractiveness and fixations. For the decision threshold ρ, we calculated an exchange rate for all shoppers, with an average of 0.009. For the forward-looking fixations term λ_1, we calculated an exchange rate for 68% of shoppers, with an average of -0.003. For the empirical threshold-changing term λ_3, we calculated an exchange rate for all shoppers, with an average of 0.002.

2.5.3. Empirical Verification of Counterfactual on Preference-Based Ordering

We also simulated fixation pathways with best-first and worst-first product orderings, based on each individual's self-reported product liking ratings in the random ordering condition. Figure 6 panel B shows that we again find that the worst-first ordering is predicted to

result in more total fixations ($M = 137.26, SD = 82.08$) compared to the best-first ordering ($M = 112.63, SD = 90.32; t(40) = 3.25, p < 0.01$). We also predict the worst-first ordering to result in lower clicked AOI attractiveness ($M = 3.37, SD = 3.49$) compared to the best-first ordering ($M = 3.93, SD = 1.49; t(40) = 3.14, p < 0.01$).

Importantly, in Experiment 2, we also have data from individuals assigned to conditions where we experimentally manipulated whether they saw the landing page products in the best-first or worst-first orderings, based on their individual self-reported product ratings. Figure 6 panel C shows the results of the experiment. Consistent with the counterfactual predictions of both Experiments 1 and 2, we find that shoppers in the worst-first condition made more fixations ($M = 109.38, SD = 59.83$) compared to shoppers in the best-first condition ($M = 78.61, SD = 57.03; t(88.56) = 2.51, p = 0.01$). This empirical verification lends credence to the assumptions of our decision threshold model of how shoppers make fixation and click decisions on the landing page.

2.6. Discussion

We contribute to the growing literature that treats the eye fixations of shoppers as decisions determined by the shopper's internal valuations of the product options as well as the physical positioning of the options within the shopping environment. Using a sequential sampling framework, we build and empirically test a search model for how online shoppers visually acquire product information on a store's landing page leading up to the decision to click on a link that brings them to a new part of the website. We estimate a heterogeneous decision threshold, and also find that shoppers exhibit both forward-looking and threshold-changing behaviors when making fixation and click decisions on the landing page.

In this work, we focus on shopper decisions on the landing page, and also provide preliminary evidence that early search may be predictive of more downstream behaviors. Specifically, we found that more search on the landing page is negatively correlated with the depth of search in the remainder of the shopping trip, as measured by the number of unique

web pages visited, as well as the dollar value of the final shopping cart. In addition, we find that an experimentally manipulated worst-first ordering resulted in fewer total pages visited compared to the best-first ordering ($t(85.2) = 2.88$, $p < 0.01$). Thus, our findings have implications for how retailers may personalize a virtual store landing page to increase traffic through a store. Although online retailers may not always have available detailed self-reported individual product preferences, other methods may be used to infer preferences, such as information from prior store visits or collaborative filtering methods.

One challenge in future research will be in formulating models for the complex decision processes within an entire shopping trip, or possibly across multiple shopping trips within a customer's "lifetime". On the landing page, the decision processes of shoppers can be reduced to a sequence of "stay" versus "leave" choices. But as shoppers visit different virtual or physical areas of a store, the set of possible actions grows very large and the search stopping rules become more complicated. For example, for any given area of the store, we must take into account the qualitative differences between ending search to visit a new location versus to revisit a prior location, which adds further complexity to incorporating dynamics like time considerations or forward-looking computations into models of shoppers' split-second decisions.

Trying to model the behavior of shoppers as they diverge from a common landing page and continue on their own search paths presents an opportunity to test different theories about how shoppers process and retain product information, as well as how they manage the physical and cognitive effort involved in search. For example, shoppers may begin to employ more forward-looking strategies in later stages of the shopping trip as they familiarize themselves with the shopping environment, as suggested by Gopalakrishnan, Iyengar, and Meyer (2014). However, note that the type of micro-level modeling that we employed in this chapter quickly becomes intractable for "real" stores because of the number of products and store locations. Thus, the current data could be modeled at a more macro-level or we can conduct future experiments within more limited shopping environments.

CHAPTER 3 : Using Variety-Seeking Preferences to Predict New Product

Adoption within Online Video Games

3.1. Introduction

As consumers spend time engaging in experiential products, they may become satiated on
certain product attributes and exhibit preferences for new experiences (variety-seeking),
or they may become hooked on certain familiar features and prefer consistency (inertia).
It is useful for firms to identify when consumers are variety-seeking or inertial when pre-
dicting purchase patterns and offering recommendations for future consumption occasions.
Although there are multiple definitions of variety-seeking versus inertial behavior, in the
present research we are focusing specifically on how consumers choose between different
options across multiple consumption occasions within the same product category.

Researchers have proposed several reasons for why consumers exhibit variety-seeking behav-
iors, including response to external factors such as price promotions, satiation on product
attributes over time, and exploration of different options to reduce future preference un-
certainty (Kahn 1995). We propose that consumers may also become variety-seeking in
response to consumption outcomes, which we define to be measures of the quality of the
consumer's experience for a specific consumption occasion.

Our first main contribution in this chapter is to quantify variety-seeking preferences not
just by the incidence of switching between products, but also by the perceptual *distance*
between products, as measured by the relative number of shared and unshared attributes.
There is a wealth of theories and empirical evidence for reward signaling and reinforcement
learning behaviors predicting that humans and animals are drawn towards positive stimuli
and away from negative stimuli (Schultz 2015), which corresponds to the prediction that
positive consumption outcomes lead to inertial preferences (i.e., repeated selection of the
same or similar choices). However, in this chapter we are also interested in testing whether
or not negative consumption outcomes lead to variety-seeking (i.e., subsequent selection of

conceptually distant choices). Therefore, we build a model framework that flexibly captures variety-seeking preferences as the coefficient on the distance from previously chosen options, and these preferences may vary over time in response to consumption outcomes. Furthermore, we allow these effects to be heterogeneous across individual consumers.

For example, Pandora is a music streaming platform that uses algorithms to predict customer preferences based on past behavior. Users can specify music artists or songs and the platform will create a radio station playlist based on this input. For each song that is played, the user can give feedback in the form of a thumbs up or a thumbs down, and Pandora will incorporate this feedback into future song selections. So if the user gave a song a thumbs up, Pandora might recommend songs by similar artists, but if the user gave a song a thumbs down, Pandora might try to recommend songs by more dissimilar artists. For Pandora, the similarity between songs is based on the firm's proprietary database known as the Music Genome Project, which consists of genre and subgenre categorizations of songs that are manually coded by human listeners.

To test the effects of consumption outcomes on variety-seeking preferences, we use data on individual player behavior in an online video game. Across 30 to 40-minute rounds of play, players choose which map they want to play on and experience consumption outcomes that can be measured by their performance or points earned during the round. A unique feature of our context is that the firm released five sets of expansion packs containing new maps at various dates following the initial release of the base game. Thus, our second main contribution is using our parameterization of variety-seeking in response to consumption outcomes to help improve out-of-sample predictions of whether or not players will adopt the new maps following their release. Although we focus specifically on the context of online video games, our findings can be applied to the broader set of experiential products, including watching movies and dining at restaurants.

3.2. Literature

In this chapter, we are specifically focused on decisions that consumers make when choosing between different options across consumption occasions within the same product category. Variety-seeking is defined to be when consumers switch frequently between very different options, while inertia (or reinforcement) is defined to be when consumers repeatedly choose the same or similar options across multiple consumption occasions. We review the main reasons that researchers have identified for why consumers exhibit variety-seeking or inertial behavior, the behavioral literature that predicts the effects of consumption outcomes, models that have been developed to capture variety-seeking, and models that use different variables (e.g., marketing mix, individual or household differences, cross-category learning) to predict new product adoption.

3.2.1. Why Are Consumers Variety-Seeking or Inertial?

The concepts of inertia or reinforcement behaviors and variety-seeking were developed separately before researchers began to think of them as two ends of the same continuum, so much of the research that provides explanations for these behaviors focuses on one extreme. Early models used time-lagged variables to capture inertial choices and attributed them to "brand loyalty" (Jacoby and Kyner 1973; Guadagni and Little 1983). Researchers explored alternative explanations for inertial behaviors such as state dependence and habit persistence, which can be disentangled using more sophisticated utility models (Erdem and Keane 1996; Seetharaman 2004).

On the other hand, Kahn (1995) summarizes the three main reasons for why consumers may be variety-seeking: (1) external situations, (2) satiation, and (3) future preference uncertainty. External situations include marketing decisions made by firms. For example, different firms may promote in alternating weeks (Kahn and Raju 1991) or engage in price discrimination (Shaffer and Zhang 2000), driving consumers to switch between brands. Satiation is a well-studied phenomenon in both behavioral and quantitative research. Sa-

tiation may occur on brands or attributes and lead consumers to seek out products with new features (McAlister 1982; Inman 2001). Finally, forward-looking consumers may use variety-seeking as a way to resolve future preference uncertainty and learn about unknown choices (Walsh 1995; Erdem 1996).

3.2.2. Effects of Consumption Outcomes on Variety-Seeking Preferences

In this chapter, we are interested in how consumers vary their variety-seeking preferences in response to consumption outcomes. Consumption outcomes are observable for a variety of experiential products. These include star ratings for movies or TV shows on Netflix, star ratings for restaurants on Yelp, thumbs up or down for videos on YouTube, and a player's score on a video game. We are going to focus on the context of video games, which has been mostly unexplored within the marketing literature. Player scores allow for a clean, relatively objective, and continuous measure of each player's consumption outcomes.

In our setting, players are able to select the specific map environment they want to play on during each round of gameplay, or specify specific attributes they prefer and let the firm's platform choose a corresponding map. Each map possesses a set of attributes, and thus we are able to calculate a perceptual map distance between any pair of maps based on their shared versus unshared attributes (Blin and Dodson 1980; Hutchinson 1986, 1989). This is contrast to the approach of Kitahira (1990), which uses individual choices in a logit framework to construct a perceptual mapping.

In our model, an individual's degree of variety-seeking (vs. inertia) is measured by his or her preferences for perceptually close vs. distant map choices in the subsequent round of gameplay. To test how these variety-seeking preferences change with consumption outcomes, we allow the coefficient on distance to vary in response to the performance measure from the previous round (i.e., points earned, kill count, etc.).

Since video games are designed to be immersive and stimulating, we assume that positive and negative outcomes correspond to positive and negative affect, respectively. Researchers

in psychology and consumer behavior have long been interested in the effects of emotions on people's choice behaviors, but there is some disagreement on how positive and negative affect influences variety-seeking. For example, positive affect has been shown to increase variety-seeking behavior among enjoyable products, as long as they don't have any negative features (Kahn and Isen 1993). Positive moods seem to drive people to seek out more stimulation, but this pattern might break down at very extreme positive moods (Roehm and Roehm 2005). Other research shows that differentiation of positive and negative emotions of the experience slows the satiation process due to cognitive appraisal, and so focusing on negative emotions may result in more enjoyment of repeated experiences (Poor, Duhachek, and Krishnan 2012). These findings would predict that in our model, the coefficient on distance (for the *next* map choice) should vary positively with consumption outcomes; in other words, better outcomes should lead to variety-seeking while negative outcomes should lead to inertia.

In contrast, the reinforcement learning literature (Schultz 2015) and related work suggests that encountering high value rewards will intensify motivational states towards the same reward source (Berridge 2012) and that positive rewards may even "whet" the reward appetite (Wadhwa, Shiv, and Nowlis 2008). In our context, a video game player may experience a hot streak and expect continued positive outcomes from playing within the same or similar map environments. Research that has focused on negative outcomes has shown that helplessness and sadness result in people wanting to change their current state (Keltner, Ellsworth, and Edwards 1993; Lazarus 1991), and they may express this desire through their consumption choices (Lerner et al. 2004). A player may feel sad or frustrated after a tough loss, and thus seek a change of scene in the next round. Thus, these findings predict that the coefficient on distance should vary negatively with consumption outcomes.

3.2.3. Dynamic Discrete Choice Models

In the classic models of variety-seeking, the underlying assumption is that the consumer is making choices between options following a first-order Markov process (Jeuland 1979;

Givon 1984; Kahn, Kalwani, and Morrison 1986). The key feature is that there is an explicit variety-seeking parameter that can be estimated for each individual consumer. Brand choices are formulated as a standard logit model, but the first-order Markov property allows the probability of choice to depend on the brand that was chosen previously. The individual-specific variety-seeking parameter determines whether repeat choices or brand switching is more likely between subsequent consumption occasions.

There have been various extensions to this classic model to take into account variation across brands, consumers, and time. The brand choice probabilities can be revised to include brand-specific marketing variables (Seetharaman and Chintagunta 1998). The variety-seeking parameter can also vary within shoppers by assuming they come from a flexible distribution. For example, the Beta distribution allows for a bimodal pattern that can account for shoppers switching between inertial and variety-seeking states (Trivedi, Bass, and Rao 1994). Heterogeneity across individuals can be modeled as individuals receiving information that arrives according to a Poisson timing function (Roy, Chintagunta, and Haldar 1996). In our current model, we will demonstrate the advantages of attribute-based variety-seeking using a continuous distance between options (Chintagunta 1998). We also allow variety-seeking preferences to change over time based on previously experienced consumption outcomes.

3.2.4. Predicting New Product Adoption

We use our model of dynamic variety-seeking preferences to improve predictions of new product adoption. There is a rich literature on modeling new product diffusion and examining the characteristics of consumers and retailers that drive new product adoption and acceptance (Bass 1969; Bronnenberg and Mela 2004; Chandrasekaran and Tellis 2008; Gielens and Steenkamp 2007). Social networks often facilitate diffusion of new products (Berger et al. 2010), for example within networks of physicians adopting new prescription medications (Iyengar, Van den Bulte, and Valente 2011). New product adoption has also been studied as a result of attribute-based preference learning (Chintagunta, Jiang, and Jin

2009), as well as cross-category learning (Shridhar, Bezawada, and Trivedi 2012).

In our context, the game platform released five new expansion packs, at various dates following the release of the original base set of map options. Each expansion pack contained a set of new maps with both shared and unshared attributes with the previously available maps, and therefore we can calculate the distance between any pair of old and new maps. Our baseline model is a discrete choice model that allows map preferences to depend only on the attributes. We use the data of players before the release of the expansion packs as the "training" data to calculate the coefficients on map attributes. We then use these coefficients to make out-of-sample predictions of whether players will adopt the maps in the expansion packs following their release by choosing them to play on, as well as other choice sequence characteristics such as overall map choice shares and the degree of map switching.

3.3. Model Specification

The degree of variety-seeking may vary across product categories (Kahn, Kalwani, and Morrison 1986) or across individuals (Givon 1984), but the model we develop is more appropriate for capturing how the degree of variety-seeking varies *within* individual consumers. Our model captures individual choices across multiple consumption occasions. First, we outline the baseline model where choice depends on a linear combination of the attributes of each option within a logit framework. Then, we lay out the specification of an alternative benchmark model with state dependence and a model with a variety-seeking component that varies with time and consumption outcomes from the previous choice.

3.3.1. Baseline Model: Attribute-Based Map Choices

In the baseline model, the probability of a consumer choosing option $j \in \{1, ..., J\}$ at consumption occasion t depends on the individual's intrinsic preferences for option j's vector of attributes $\boldsymbol{x}_j$, which is given by the attribute coefficients $\boldsymbol{\beta}_i$. Without any time varying

effects, the probability of each choice j is formulated by a standard logit:

$$P_{ij}[t] = \frac{\exp\{\boldsymbol{\beta}_i^\top \boldsymbol{x}_j\}}{\sum_{k=1}^{J} \exp\{\boldsymbol{\beta}_i^\top \boldsymbol{x}_k\}} \tag{3.1}$$

In our context, the options are the different maps that individual players can choose from in each round t, and the attributes are the different binary features of the maps (e.g., whether or not the landscape is a forest, whether or not players can use a specific weapon, etc.). Each individual's likelihood function is the product of map choice probabilities, with parameters $\boldsymbol{\theta}_i = \{\boldsymbol{\beta}_i, \alpha_i, \boldsymbol{\gamma}_i\}$ described in detail in the following sections:

$$\mathcal{L}(\boldsymbol{\theta}_i) = \prod_{t=1}^{T_i} P_{ij}[t] \tag{3.2}$$

3.3.2. State Dependence

To test whether map switching across rounds can be captured by state dependence, we include a parameter α_i as the coefficient of a dummy variable for whether the map choice j at the current round t is the same map as the map j' chosen during the previous round $t-1$. Thus, the probability of choosing map j at round t becomes the following:

$$P_{ij}[t] = \frac{\exp\{\boldsymbol{\beta}_i^\top \boldsymbol{x}_j + \alpha_i \mathbb{1}\{j = j'\}\}}{\sum_{k=1}^{J} \exp\{\boldsymbol{\beta}_i^\top \boldsymbol{x}_k + \alpha_i \mathbb{1}\{k = j'\}\}} \tag{3.3}$$

3.3.3. Dynamic Variety-Seeking Parameter

To capture how consumers respond to consumption outcomes over time, we specify choice preferences to be first-order Markov across rounds. So the probability of selecting map j depends on the map j' that was selected in the previous round $t-1$ through a variety-seeking parameter $\delta_i \in [-\infty, +\infty]$ that determines how much weight is given to the perceptual distance $d_{jj'}$ between map j and the previously chosen map j'. Thus, the probability of

choosing map j during round t is conditional on j':

$$P_{i,j|j'}[t] = \frac{\exp\{\boldsymbol{\beta}_i \boldsymbol{x}_j + \delta_i[t]d_{jj'}\}}{\sum_{k=1}^{J}\exp\{\boldsymbol{\beta}_i \boldsymbol{x}_k + \delta_i[t]d_{kj'}\}} \tag{3.4}$$

In our context, the distance between any two map options is determined by the relative number of shared and unshared binary attributes between the maps. Note that we allow the variety-seeking parameter δ to vary across rounds t.

To illustrate how $\delta_i[t]$ comes into play, we note that the probability of repeating option j' (in other words, if $j = j'$) is different from the probability of switching from option j' to *different* option j (in other words, if $j \neq j'$), since the distance between a map and itself is 0. For example, the probability of repeating map j during round t is given by the following probability:

$$P_{i,j|j}[t] = \frac{\exp\{\boldsymbol{\beta}\boldsymbol{x}_j\}}{\sum_k \exp\{\boldsymbol{\beta}x_k + \delta[t]d_{kj}\}} \tag{3.5}$$

If $\delta_i[t] \geq 0$, then the player has a positive weight on map distance and can be considered variety-seeking during round t, and so the probability of switching maps (relative to the previous map j') is higher. On the other hand, if $\delta_i[t] \leq 0$, then the individual has a negative weight on map distance and can be considered inertial at round t, and so the probability of staying with the previous map is higher.

These probabilities also depend on the distance between the maps. So a variety-seeking player is also more likely to switch to maps that are farther away, while an inertial player is more likely to choose maps that are closer to the previously chosen map.

We allow each individual's variety-seeking parameter $\delta_i[t]$ to vary across rounds t as a function of the number of rounds that have passed so far and the consumption outcome experienced at the previous round $t - 1$. For example, in our context, the consumption

52

outcome could be the points earned during the round or the number of kills.

$$\delta_i[t] = \gamma_{i0} + \gamma_{i1}t + \gamma_{i2}\text{Outcome}[t\text{-}1] \tag{3.6}$$

3.4. Data

Our dataset was awarded through the Wharton Customer Analytics Initiative (WCAI) from a large video game developer. We have data on the activity of 1,309 frequent players of an online multiplayer first-person shooter video game. Players engage in campaigns averaging 20 minutes in length in two competing teams. We focus on the rounds played on the firm's public servers and exclude the rounds played on player's private servers. There are on average 20 players involved in each round, and people rarely play with the same player twice. We have each player's daily rounds played across two years, starting from the game's release.

Figure 9: Outline of a player's sequence of actions for each round
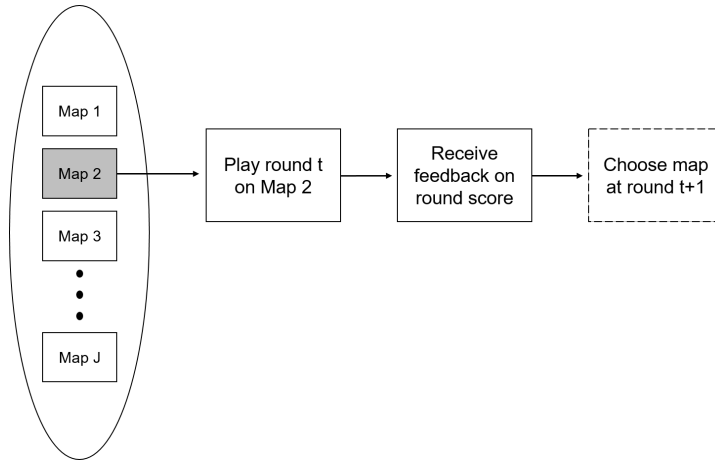
Figure 9 illustrates a player's actions before and after each round of play. During each round, players are presented with a set of J maps (which may vary over time with the release of expansion packs). A map is basically a game environment with a set of attributes or features.The player chooses a map and is dropped into a server by the firm's matching algorithm with other players to play a round. After the round, the player is shown his/her

individual round outcomes, which might include the number of kills, number of deaths, individual points earned for completing certain tasks in the round, etc. In the next round, the player will again have the opportunity to select a map.

On the first day of the game's release, there were 9 "base" maps available for players to choose from. At various dates in the following two-year period, the firm released 5 expansion packs, each containing 4 maps, for a total of 29 maps. One main objective of our model is to use our parameterization of dynamic variety-seeking preferences during the base map period to predict whether players adopt the new maps upon their release in the expansion packs. Therefore we restrict our analyses to players who had played at least 10 rounds on at least 2 different maps before the release of the first expansion pack, which resulted in 569 players for analysis. In other words, we excluded players with insufficient training data.

Figure 10 shows the frequency of rounds played across the days following the base map release date. Expansion packs were released 50 days, 226 days, 316 days, 400 days, and 498 days after the game's release, as indicated by the black bars. We see that after the release of each expansion pack, there is an increase in player activity.

Figure 11 shows the round frequency across maps, averaged across players. Maps 1 through 9 belong to the base map set, and were by far the most popular in terms of round frequency. Rounds played on the expansion pack maps, consisting of Maps 10 through 29, account for 24% of the rounds played.

Figure 12 shows the percentage of players who played at least one round for each map choice. We see that nearly all players in our sample played at least once on each map in the base set (Maps 1 through 9), with lower percentages for the different map expansion packs (Maps 10 through 29, released in groups of 4). Whether or not players play a round on a map is our measure of new product adoption, which we will use the data on rounds

Figure 10: Distribution of rounds over time since game release
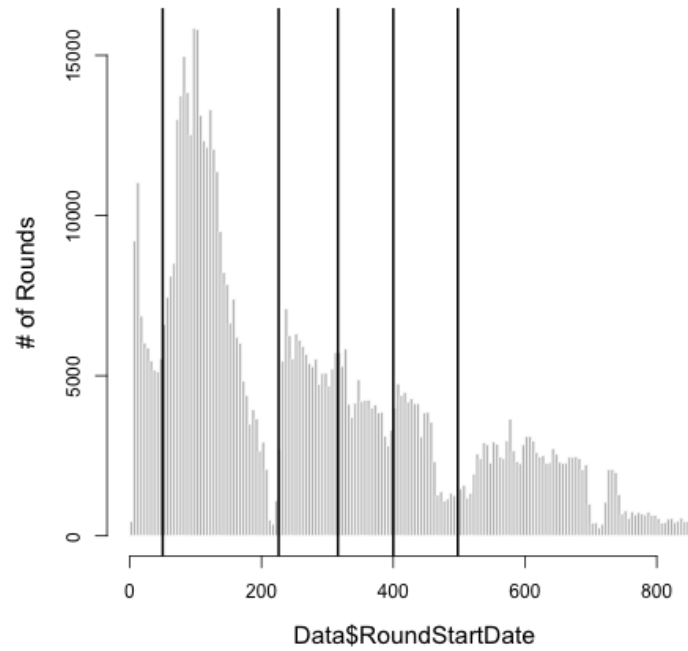
Figure 11: Outline of a player's sequence of actions for each round of play.
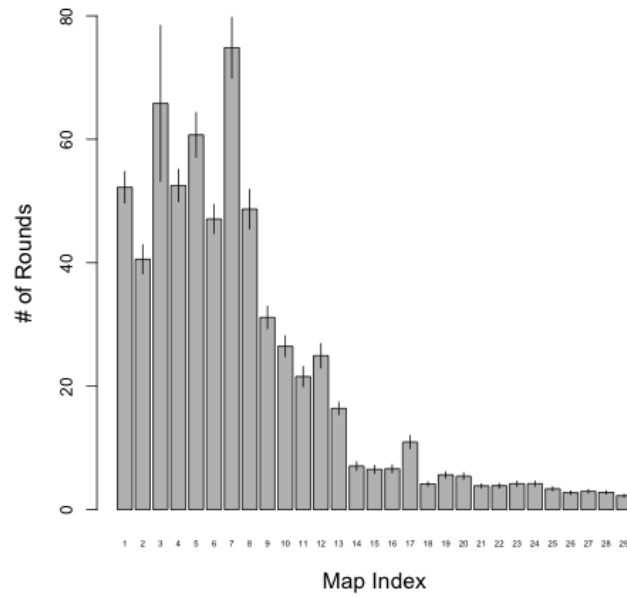
preceding the expansion pack release to predict.

Figure 13 shows the distribution of the number of total rounds across players. The 569 players in our sample played an average of 639 total rounds (SD = 782) across the data collection period, with high variation across players.

Figure 12: Percentage of players who played at least one round for each map
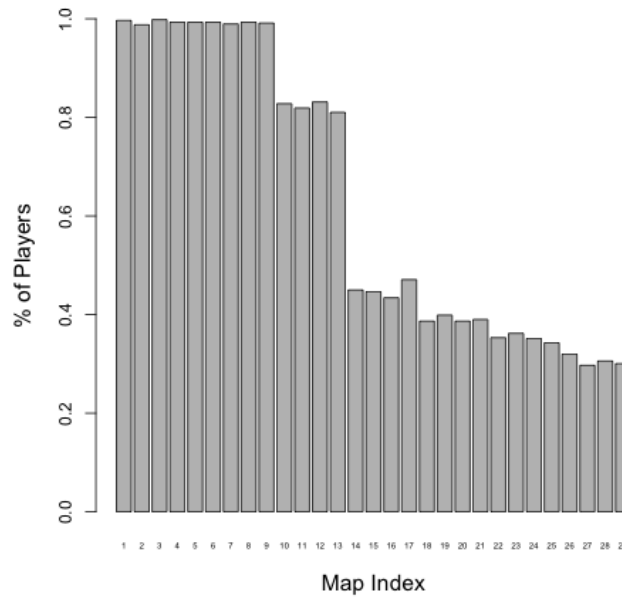
3.5. Descriptive Analysis

To examine the variety-seeking preferences of players in an exploratory way, we first look at how often players simply switch maps between rounds. For each player in our sample, we calculated the percentage of rounds where they switched to a different map from the previous round. We see that the average switching rate is quite high at 83%, so it would appear that players are very variety-seeking because they are switching between maps very often, as shown in Figure 14.

Figure 13: Distribution of total rounds across players
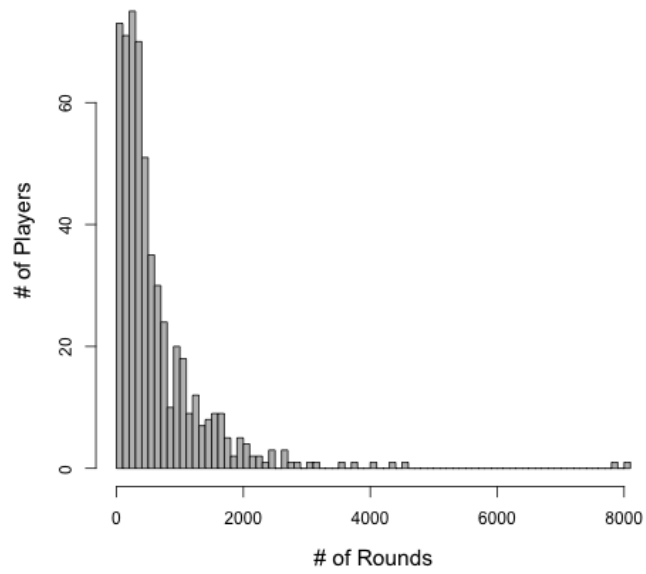
Figure 14: Distribution of map switching rates across players
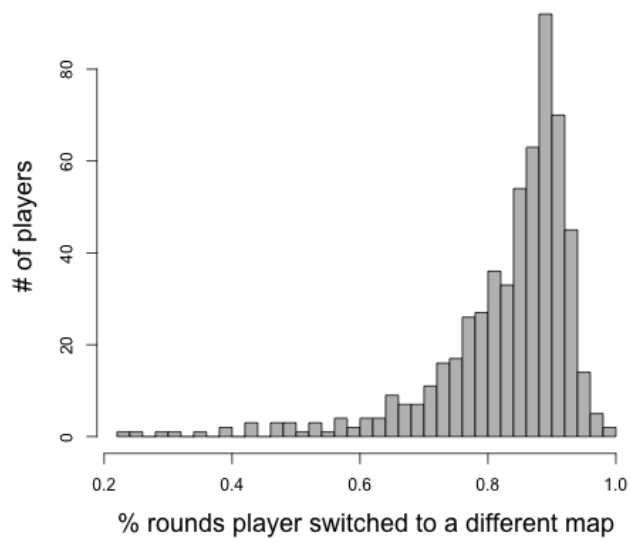
However, one important extension of the classic brand switching variety-seeking model is the use of attribute-based variety-seeking. For example, we might conclude that a player is very variety-seeking because they are switching maps nearly every round, but they may actually be switching only occasionally between clusters of similar maps.

In order to quantify the similarity between any two maps, we created a measure of map distance. We have 29 maps total, and they possess 14 different binary attributes. To calculate the distance between any two maps, we simply take the correlations across the maps and subtract them from 1. The map distance can fall anywhere between 0 and 2. The distance between any map and itself is 0. The range of map distances in our dataset falls between 0 and 1.4, as shown in Figure 15.

Figure 15: Distribution of map distances
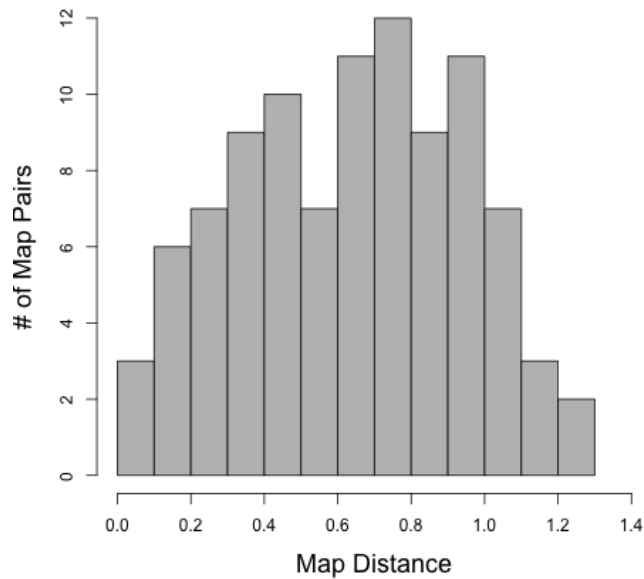
If we rank the maps from most played to least played for each player, and look at the cumulative percentage of rounds played on these maps, we see that most players spend their time playing on just a few maps (see Figure 16). The top 10 favorite maps make

58

up on average about 90% of a player's rounds. If we plot the average distance from the favorite map (see Figure 17), then we see that the most commonly played maps are also more similar to each other.

Figure 16: Cumulative frequency of maps played, ranked from most to least played
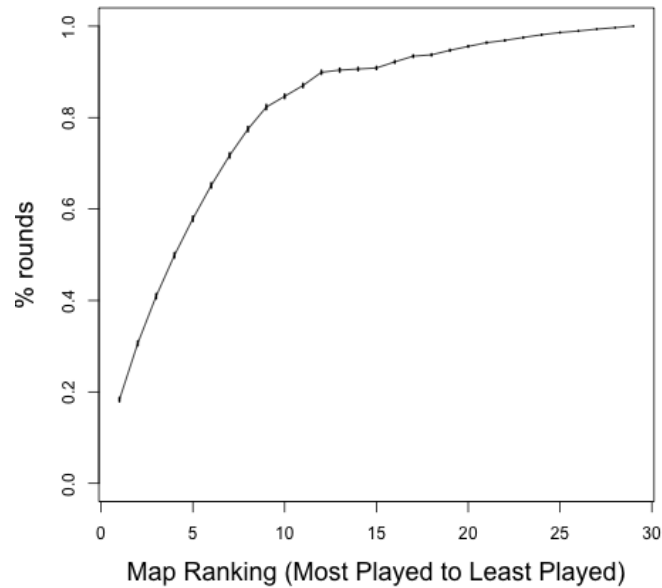
Over time, the average rate of map switching across the entire player sample seems to decrease slightly over time (see Figure 18). There are two explanations for this pattern. One is that at the individual level, players are first exploring the different map options, and eventually settle on playing a smaller set of their favorite maps. The second explanation is that this pattern arises from the heterogeneity across players. It is possible that players who start playing the game early on are more variety-seeking, while players who enter later are less variety-seeking and bring down the average switching rate. However, because we restricted our sample to players who played at least 10 rounds before the release of the first expansion pack, we can rule out the second explanation.

After each round, the player experiences a set of individual consumption outcomes. These

Figure 17: Distance from favorite maps, with maps ranked from most to least played
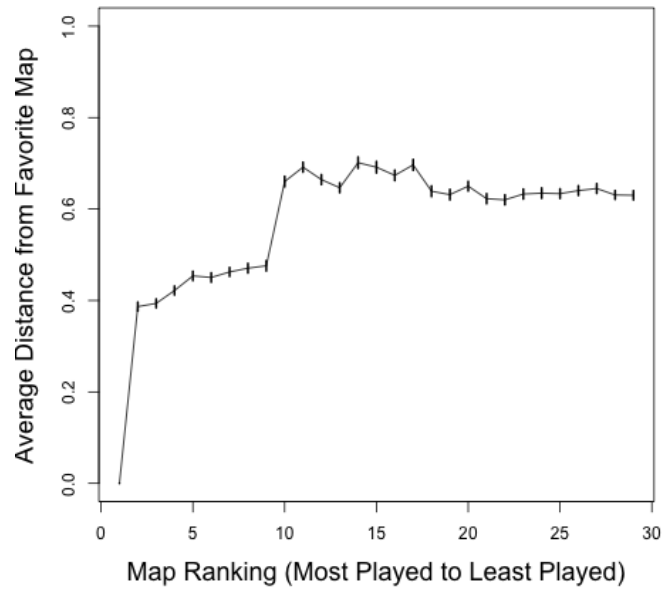
Figure 18: Players switch between maps less over time
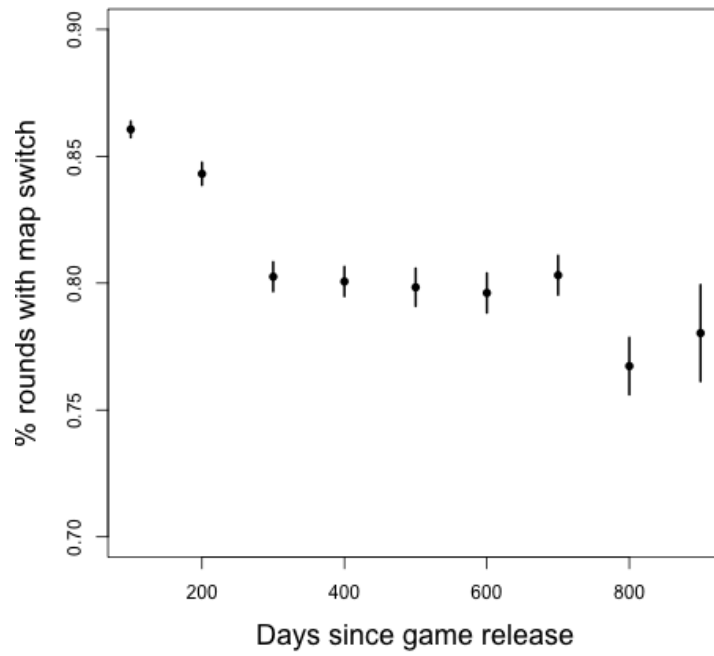
include Total Points, Combat Points, Kill Count, Death Count, and net kills (Kills - Deaths). These variables are all pretty highly correlated (see Table 6), so we use Kills, which is generally the primary and most salient objective of the game.

Table 6: Correlation between consumption outcomes

	Total Points	Combat Points	Kill Count	Death Count	Kills - Deaths
Total Points	1				
Combat Points	.79	1			
Kill Count	.73	.84	1		
Death Count	.47	.59	.66	1	
Kills - Deaths	.44	.44	.57	-.24	1

3.6. Model Estimation Results

We estimate the model using maximum likelihood estimation for each individual, using the rounds played on the base maps *before* the release of any expansion packs. We are able to implement this because of the large number of rounds for each player. However, in future work, we plan to use a hierarchical Bayes estimation procedure, which will allow information pooling across players and improve the robustness of the results. First, we compare the model fit across three model variations: (1) a baseline model with choices based only on map attributes, (2) a model with state dependence, and (3) a model with our dynamic variety-seeking parameterization.

Each player in our sample had a sufficient amount of training data, so we use their individual-level parameter estimates to predict their out-of-sample map choices upon the release of each of the five expansion packs.

3.6.1. Comparison of Model Fit Statistics

Table 7 reports the in-sample model fit statistics, with the likelihoods summed across individuals to compute the -2LL, AIC, and BIC. In addition to these fit statistics, we were also

interested in looking at whether our model was able to capture the patterns in the sequence of map choices, specifically in terms of the map choice shares and how often players switched between maps.

For each player, we simulated 1,000 map choice sequences using their estimated parameters, conditioning on the observed number of choices, for the period before the first expansion pack release. For each simulated sequence, we computed the correlation between the simulated map choice shares (across the 9 base maps) and the observed choice shares. We also computed the mean map switching rate, as well as the root mean squared error of the map switching rate between the simulated and observed sequences (the observed average switching rate is 83.72%).

Table 7: Comparison of fit statistics across models

Model	-2LL	AIC	BIC	Choice Corr	Switch Rate (RMSE)
Attribute-Based Choice	248389.6	254079.6	260559.4	.53	.8481 (.090)
State Dependence	246022.0	252850	260625.7	.56	.8370 (.050)
Variety-Seeking	244654.9	250344.9	256824.7	.53	.8393 (.051)

We see that adding state dependence and the variety-seeking parameters improves the in-sample model fit statistics. Interestingly, adding state dependence improves model fit in terms of the map choice share correlation and the switching rate, but adding variety-seeking does not. Both state dependence and variety-seeking slightly improve the RMSE of the map switching rate.

However, our main objective is to assess whether the variety-seeking parameter is able to improve out-of-sample predictions of whether players adopted the new maps upon their release across the five expansion packs. Therefore, we simulated map choice sequences using each individual's parameter estimates during the periods following the release of the 5 expansion packs, conditioning on the number of rounds they played between releases. We

then computed the percentage of times the model was able to correctly predict whether or not players actually played on one of the new maps.

Figure 19 compares the percentage of correct adoption predictions for each of the five expansion packs across the three model variations. We see that the state-dependence model does not improve predictions over the attribute-based model. However, the variety-seeking model does significantly improve the percentage of correct adoption predictions across all 5 expansion packs. This suggests that the firm could potentially use individual players' past variety-seeking preferences, as exhibited through their sequence of map choices, to predict how receptive they would be to adopting new content (i.e., maps), given that the firm knows the perceptual relationship between the old and new options.

Figure 19: Expansion pack map adoption predictions across models
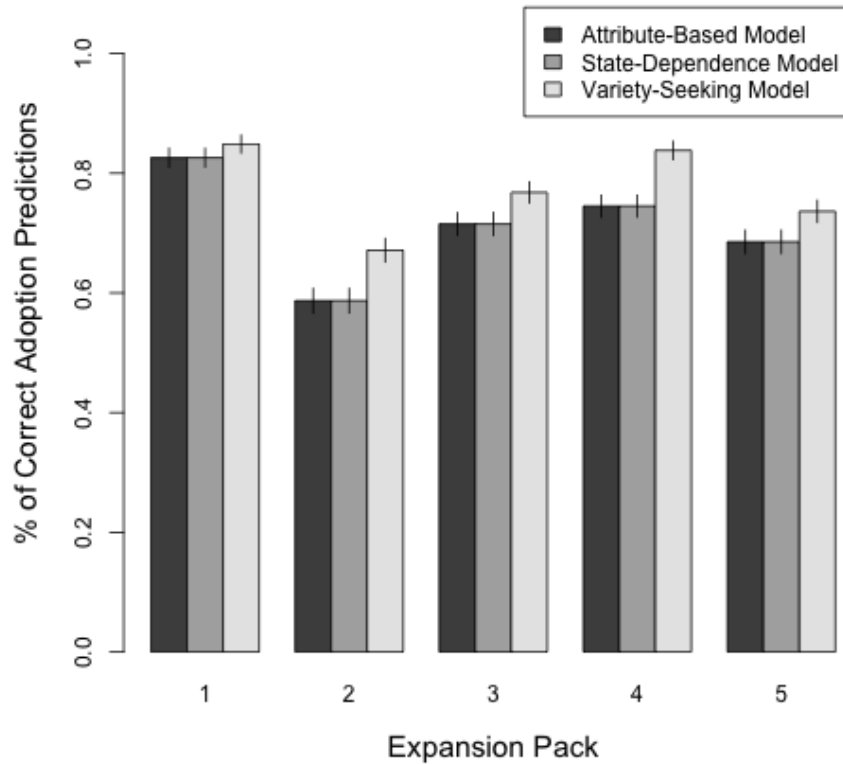
3.6.2. State Dependence and Variety-Seeking Parameters

In this section we take a closer look at the heterogeneity across players for the coefficient on state dependence, as well as the coefficients on the variety-seeking variables, including the intercept, the round number, and the consumption outcome. There is a lot of heterogeneity across all variables. On average, the state dependence variable is negative, which captures the frequent switching between maps. The intercept on variety-seeking is positive on average, which captures the overall propensity to also switch to more distant options. The coefficient on round is negative, which reflects the general trend that players become less variety-seeking over time (consistent with Figure 18). The coefficient on round outcome is evenly spread around 0, which suggests that we may need to consider whether individuals have different reference points for "good" versus "bad" outcomes.

Figure 20: Distribution of individual parameter estimates for state dependence and variety-seeking
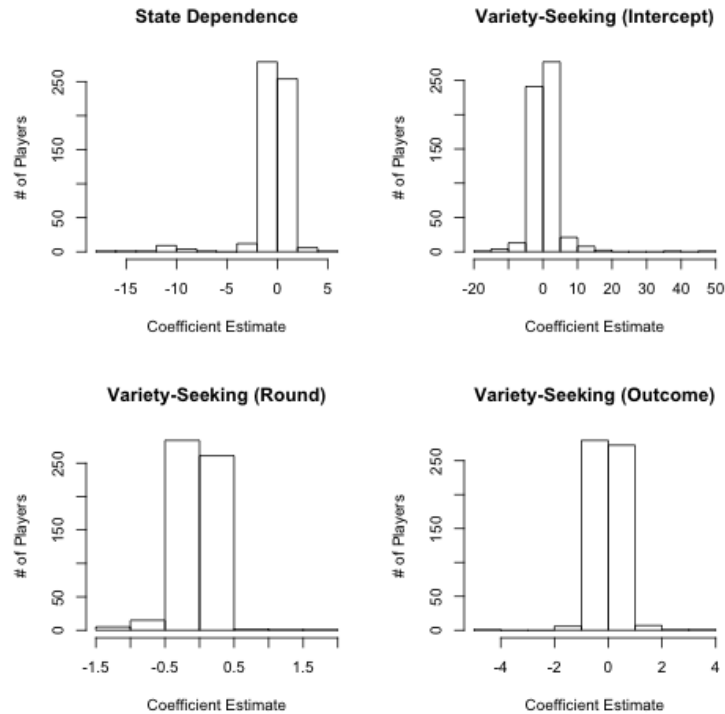
We can also examine how the variety-seeking parameter changes across rounds. Figure 21 plots $\delta_i[t]$ from Equation 3.6 for a sample player. We see that the player's variety-seeking intercept is negative, so the player's overall prefers closer maps. Across rounds, the overall trend is a decrease in variety-seeking preferences, but there are also round-to-round spikes that reflects the reaction to the round outcomes (i.e., kill count). This suggests that an alternative model specification might involve allowing round outcomes to have more long-term or "smoother" effects on variety-seeking.

Figure 21: Variety-seeking parameter δ_i across rounds for one sample individual
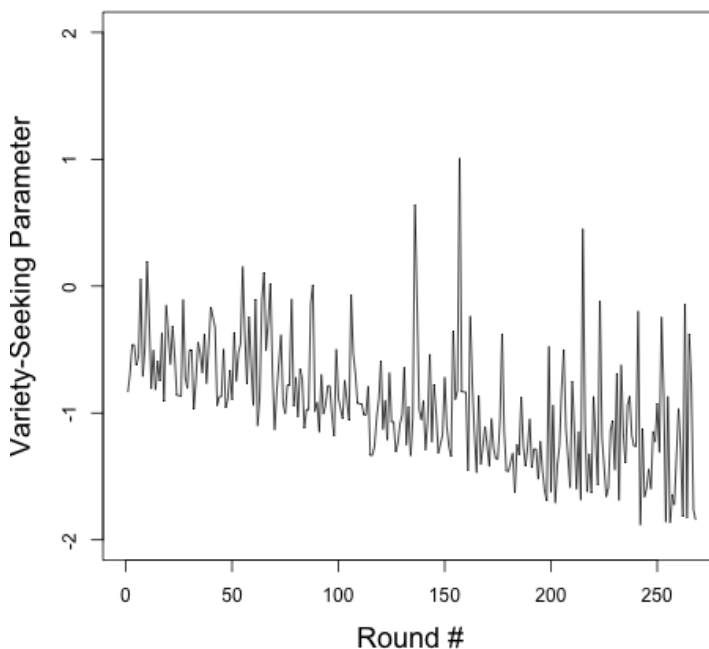
3.7. Discussion

We built a descriptive model that allows for an individual-specific variety-seeking parameter to vary over time. Variety-seeking depends explicitly on time and consumption outcomes. Specifically, we show that incorporating dynamic variety-seeking preferences into the model improves out-of-sample predictions of whether or not consumers adopt new options in the future. Although our analyses were conducted within the context of player map choices

within an online video game, our model can be extended to other domains of experiential products, including TV shows and restaurants.

There are several extensions that could be made to our current model. Currently we are assuming that all attributes contribute equally to the distance between maps, but this assumption could be relaxed and we could potentially estimate a different weight for each attribute. Since our data begins at the game's release, we could also incorporate how players are learning map attributes or map distances over time. In addition, we are assuming that players are simply maximizing of the utility for each map, but they may actually be trying to fill some goal or quota of gameplay, which would change the way we formulate their utility function. Finally, we may consider that players have different reference points for their consumption outcomes that could vary across maps or even be learned over time as players gain experience and improve their skills within the game.

For managers, this research will provide a method for determining individual customer preferences and how these preferences change either over time or based on consumption outcomes, as well as when customers might be more susceptible to the release of new products and whether they should be novel or similar to existing products. For example, within our specific context, the firm could determine in real-time whether players are becoming bored or frustrated with the current playing experience and prefer a change of scene, or prefer to continue with the same experience. This provides an opportunity to enhance the firm's current matching algorithm by suggesting that the player's variety-seeking preferences should be taken into account. In the long run, the firm can also use information about players' variety-seeking preferences to determine the timing of new content release, as well as the design of the new content relative to the existing content.

CHAPTER 4 : Testing Theories of Goal Progress and Knowledge Accumulation in
Online Learning

4.1. Introduction

Providers of online content cover a wide spectrum of content types, ranging from entertainment to education. For example, video streaming services like Netflix and Hulu offer content for entertainment, while websites like Coursera, edX, and Khan Academy offer educational content. YouTube offers a mix of both; for example, music videos and comedic skits could be categorized as entertainment, while how-to videos and TED talks could be categorized as educational.

Most of the prior work studying online content consumption has looked at entertainment platforms (e.g., Nelson, Meyvis, and Galak 2009; Hui, Meyvis, and Assael 2014; Schweidel and Moe 2016). In contrast, we specifically focus on the setting of online education, which has experienced rapid growth and change within the last decade, with firms offering both free and monetized content, as well as partnering with universities to offer specializations and degrees. The setting of online education has unique features such as scheduled content tied to learning assessments that make it an appropriate testing ground for theories of how consumers progress towards their goals and accumulate knowledge.

Using a unique clickstream dataset from Coursera, we examine the behavior of two groups of paying users engaged in two business courses offered on the platform between 2015 and 2017: Introduction to Marketing and Introduction to Operations Management, each consisting of four weeks of lecture videos and quizzes. We observe individual-level lecture video and quiz consumption across both courses, as well as outcome variables such as quiz scores, which learners accumulate to obtain a certificate for passing a given course. For the first group of users, course content was released *sequentially* with a few videos and lectures made available each week, akin to scheduled TV programming. For the second group of users, course content was released *simultaneously* with all videos and lectures available from the

start, akin to Netflix's full season release style. This data provides a rich, real-life testing ground for recent behavioral theories as described in the next section.

4.1.1. Research Contributions

We make four main research contributions. First, we build a mathematical model of choice that captures individual decisions about which course to consume, whether the content is a lecture or a quiz, and when to take breaks of different lengths. Thus, our model integrates choice and discrete timing components. The parameters of our model can be mapped to specific theories from consumer psychology, which provides an underlying behavioral interpretation for the patterns we find in the data. There are three key components in our model: the contemporaneous utility of consumption as individuals make goal progress (Kivetz, Urminsky, and Zhang 2006), the utility gained from accumulating knowledge through lectures in order to pass quizzes (Alba and Hutchinson 1987; Camerer and Ho 1999), and "boundedly rational" forward-looking decisions (Camerer, Ho, and Chong 2004; Gabaix et al. 2006). These components are particularly well-suited to capturing consumption within the context of online education, but individuals may exhibit these behaviors within other media domains as well. For example, consumers may be motivated to continue watching a TV show because they have the goal to get to the season finale, or to accumulate knowledge of the show's characters within a complex storyline.

Second, we demonstrate that our model is able to capture patterns in the data that we characterize as "binge learning." In this chapter, we are the first to distinguish between two types of binging: "temporal binging," which depends on the *amount* of content that is consumed in a single sitting, and "content binging," which depends on the *type* of content that is consumed. In our context, temporal bingers consume multiple lectures and quizzes without taking a break, while content bingers switch infrequently between courses, and we observe that many learners in our dataset exhibit both types of binging. Similarly, Schweidel and Moe (2016) modeled in a reduced-form way how viewers of TV shows on Hulu.com choose to continue the viewing session, whether the next episode viewed is from the same

show or a different show, and the time elapsed between sessions, as a function of the number of episodes viewed of various shows so far and individual-level traits. This approach allowed them to distinguish between bingers and non-bingers, which was related to individual-level advertising response. In contrast to this approach, in our model, binge behaviors (both temporal binging and content binging) are not explicit choices that individuals are making or explicit parameters, but are instead an endogenous outcome of consumer decision processes. This is similar to the approach of Erdem and Keane (1996) who demonstrate that variety-seeking may be a result of consumers reducing uncertainty about brands.

Third, firms like Coursera have some control of more upstream marketing policies. Thus, we test how the timing of content release affects consumption patterns and knowledge accumulation. To do this, we use data on individual learner behavior from both before and after a natural experiment policy change that shifted content release from weekly sequential release to on-demand simultaneous release for the courses that we examine. For each set of data, we use the parameter estimates of our model to conduct counterfactual simulations in order to predict consumption patterns in the alternative scenario (i.e., we use the sequential data to predict the simultaneous data, and the simultaneous data to predict the sequential data). This has implications for how firms should time product release, either using the traditional scheduling of TV networks or the on-demand streaming of online sites.

Fourth, we examine whether temporal and content binging patterns are predictive of more downstream behaviors that are of interest to both instructors and firms. In particular, prior research has found that the clumpiness of purchase decisions is predictive of customer lifetime value (Zhang, Bradlow, and Small 2014), satiation and variety-seeking are incorporated into models to improve predictions of brand choice (McAlister 1982; Chintagunta 1998), and binge-watching TV shows has been related to advertising response (Schweidel and Moe 2016). In our data, we find that temporal and content binging are predictive of both within-course and cross-course downstream behaviors. For example, binging in earlier weeks within a course is predictive of binging in later weeks, course completion, and

grades. Binging within Marketing and Operations is also predictive of engagement in later courses such as Accounting and Finance, in terms of total consumption, course completion, and whether individuals pay for the course certificate. These findings have implications for new product launch, cross-selling, bundling, monetization of subscription services, and CLV (Kumar et al. 2008; Fader, Hardie, and Lee 2005; Zhang, Bradlow, and Small 2015).

The remainder of this chapter is outlined as follows. In Section 4.2., we describe a clickstream dataset obtained from Coursera where we observe the lecture and quiz consumption patterns of students across two business courses offered on the platform (Marketing and Operations). We also present an overview of our model of individual consumption decisions of lectures and quizzes within the two courses. In Section 4.3., we provide descriptive evidence of binge consumption within the individuals observed in our dataset; importantly, we distinguish between content binging and temporal binging. We also examine how the propensity to engage in the two courses changes as individuals progress through the content. Section 4.4. lays out our formal model, including the parameterization of each component of the choice utilities. Section 4.5 presents our estimation procedure. Section 4.6 describes the results from the parameter estimation of the model for each dataset and the behavioral implications of the parameter values.

In Section 4.7., we conduct counterfactual simulations to see how consumption patterns might change with different content release schedules, and also test our predictions using observed data. In Section 4.8., we look at whether the patterns observed in Marketing and Operations are predictive of within-course behaviors, as well as cross-course behaviors, which is akin to assessing the value of this data for different product launches. Section 4.9. concludes with directions for future research in the area of online consumption.

4.2. Data and Model Overview

We examine the behavior of individuals engaged in the introductory Marketing and Operations courses offered on the Coursera platform through Wharton Online. We describe

our sequential data sample for individuals who took the courses in 2015, as well as our simultaneous data sample for individuals who took the courses in 2016 and 2017 following the platform's policy change regarding the timing of content release. We then give an overview of our model where individuals make two-stage consumption decisions, with the first stage involving the choice of which course to take (with the outside option being to take a break) and the second stage involving the choice between watching lectures or taking quizzes within a course (or breaks of different lengths).

4.2.1. Sequential Data Sample

During the sequential weekly release period in 2015, both Marketing and Operations were offered multiple times throughout the year, which we refer to as "sections". Each section of each course spanned 5 weeks, with 4-7 hours of lecture videos and quizzes made available at the beginning of the week for the first 4 of the 5 weeks.

In order to be observed in the data, individuals had to be registered for a section of a course. Focusing on a single section of Marketing and Operations, in which both courses were held during the same four weeks (June 1st to June 29th), we observe 61,661 individuals registered for Marketing and 46,388 individuals registered for Operations. Because we are interested in modeling how learners switch between courses, we then filtered our sample by looking at individuals who had registered for *both* Marketing and Operations, resulting in 13,690 unique individuals. However, individuals could register without actually consuming any of the content, so we further filtered them by whether they had visited at least one quiz or lecture within each course, resulting in 2,943 individuals.

Finally, we focus on registered users who had also paid for both courses, giving us a final sample of 467 individuals. Thus, we are focusing on the most committed students who have consumed most or all of the course, which allows us to assess whether our model can capture their temporal and content binging patterns.[1]

---

[1]Our model can be applied to a broader set of online courses with similar lecture and quiz structures, and although here we condition on payment, our model can also be extended to make predictions about if

71

For each individual in our final sample, we obtained their clickstream sequence recorded while they were on the Coursera website. We aggregated this data to the URL level, focusing on "submissions," which were defined to be when individuals neared the end a lecture video (80%) or submitted a quiz score for grading.[2] Thus, at each observation we know whether the individual is at the URL of a lecture or a quiz,[3] which lecture or quiz they were looking at in particular, and the timestamp of the URL submission.

Lectures consisted of 2-30 minute videos, while quizzes consisted of multiple-choice questions. The Marketing course consisted of 34 lectures and 4 quizzes, while the Operations course consisted of 26 lectures and 5 quizzes. For all individuals in our sample, we used the first time they engaged in either course as their start date, and looked at URL submissions within a 5-week period following this individualized start date (i.e., effectively "left-justifying" the data). We observed a mean of 64.0 URL submissions per individual (SD = 39.2).

Each week, new lecture and quiz materials were released, and individuals could revisit material from previous weeks. Figure 22 shows the weekly release schedule of lecture and quiz content for both Marketing and Operations (see Appendix A.7. for a list of lecture/quiz names and video run times). We refer to this type of weekly release as sequential, in contrast to simultaneous release where all content is made available from day one.

We abstract away from the specific lecture or quiz number in this research because binging, under our definition, does not depend on the specific unit consumed, but rather on how much in total the individual has engaged in the lectures and quizzes, as well as the knowledge that results from content consumption.

---

and when individuals pay. However, we leave these possibilities for future research.

[2] We removed redundant observations where individuals submitted the same lecture multiple times within a 1-hour cutoff period.

[3] We ignore URLs of pages on the website that did not offer content (e.g., FAQs, course announcements, forums, etc.), as well as optional in-lecture quizzes that were presented as one-question answers required to move through the lecture.

Figure 22: Sequential release of content for Marketing and Operations
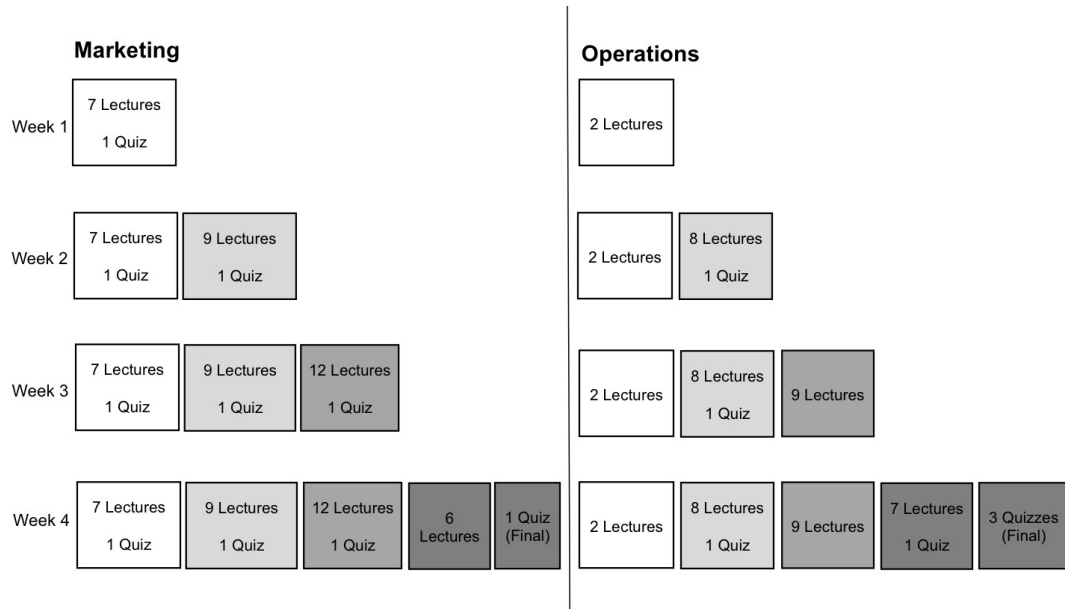
Figure 23 plots the density of URL submissions over time for Marketing and Operations, separated by the week that the content was made available. At the beginning of each week, there is a spike in activity when new content is released, which drops off until the end of the week, after which there is a second spike in activity, which can be interpreted either as individuals cramming for quizzes or increasing engagement due to the new content. Note that individuals can't (and don't) engage in content that has not been made available yet.

Figure 24 plots the sequence of URL submissions for a single individual (each submission here is referred to as an "event" j), separated by Marketing and Operations. The lectures and quizzes are numbered from 1 to 38 for Marketing and from 1 to 31 for Operations, in the order of their appearance on the website. In this example, at Event 1, the individual chose to engage in Marketing Lecture/Quiz 1 (see left hand side of Figure 24), and continued to engage in Marketing lectures and quizzes until Event 9. At Event 9, she switched over to Operations (see right hand side of Figure 24) and engaged in Operations Lecture/Quiz 1, and so on and so forth. Overall, her visits seem to be relative evenly distributed across the different lectures and quizzes. In fact, the individual mostly seems to progress system-

atically through the material by watching lectures and taking quizzes one-by-one in order, with few "skips" ahead or back to previous content.

Figure 23: Density of engagement in content separated by week.
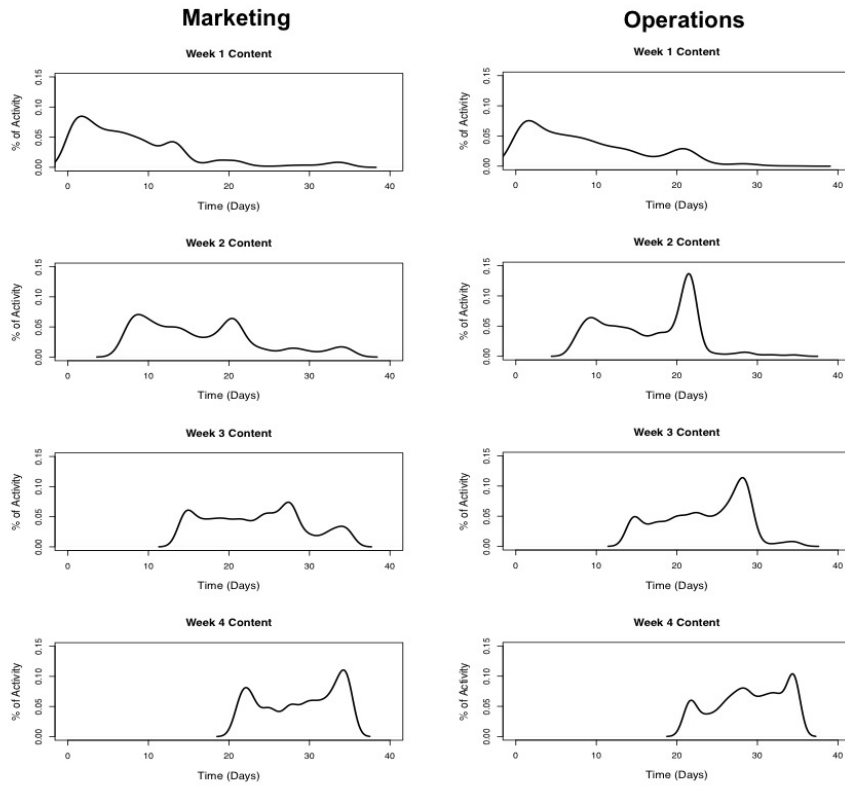
Figure 24: Sequence of lecture and quiz choices at each event j for a sample individual.
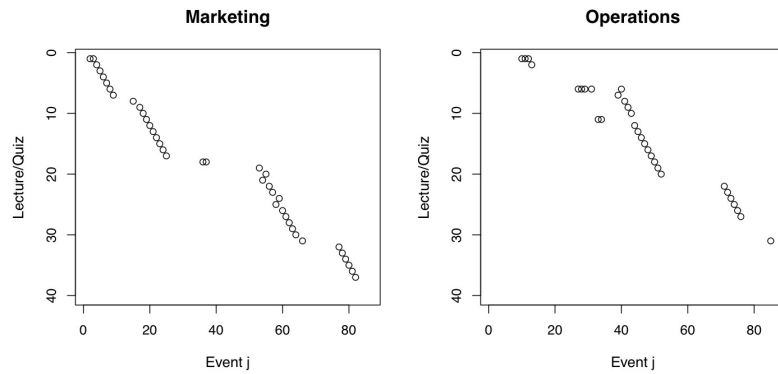
Figure 25 plots bars for each lecture and quiz representing the percentage of individuals who visited that particular lecture or quiz at least once. The shades of grey indicate the week that the content was released, while the stars indicate the bars that represent quizzes. We make three main observations. First, the quizzes were visited by a larger percentage of individuals than the lectures, which is consistent with the fact that only passing the quizzes was required for passing the course. Second, the percentages exhibit a decreasing trend, especially in Operations, which is consistent with the general observation within online courses that there is attrition over time. It is also possible that individuals who had passed "enough" of the course (i.e., obtained at least 80% of the quiz points) simply felt they didn't need to continue consuming content during the later stages of the course. Third, we notice that there are not glaring differences in the percentages within a particular week (besides the greater percentages for quizzes vs. lectures), which further suggests that there might not be significant gains, as previously mentioned, to modeling choices at the specific lecture/quiz level.

Figure 25: Percentage of individuals who viewed each lecture and quiz at least once
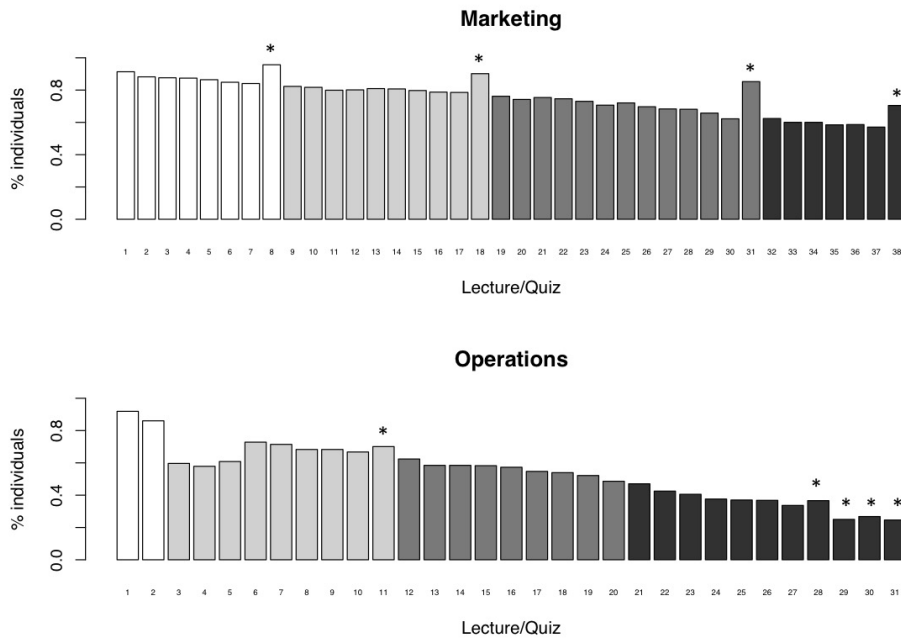
4.2.2. Simultaneous Data Sample

During the simultaneous release period beginning in 2016, all the content from both Marketing and Operations were available on-demand (i.e., at any time) to registered individuals. We started with 7,198 individuals who had registered and paid for Marketing and 2,828 individuals who had registered and paid for Operations. We then filtered our sample by looking at individuals who had registered and paid for *both* Marketing and Operations, resulting in 1,736 unique individuals.

Unlike during the sequential release period, there were no sessions or explicit start course dates. In order to create a matching sample of users with the simultaneous release data, we again assumed that each individual's start time was the time at which they first engaged in a lecture or quiz from either course. We then filtered our sample by selecting for individual who had started the *other* course within two weeks following the initial start time. This gave us enough overlap to examine and model switching behaviors between the two courses, and matched the characteristics of individuals in the sequential data. This resulted in a final sample of 448 individuals for analysis.

In addition, just like with the sequential data, we use URL submissions within the 5-week period following the initial start date. This resulted in a mean of 92.4 (SD = 56.8) URL submissions across the individuals in our simultaneous data sample.

4.2.3. Choice Model Overview

We assume that at each event j individuals are making a two-stage decision, as illustrated in Figure 26. In stage 1, the individual decides whether she wants to engage in Marketing, engage in Operations, or take a Break. In stage 2, the choice set for the individual is conditional on what she chose in stage 1. If she decided to engage in either Marketing or Operations in stage 1, then in stage 2 she chooses to either consume a quiz or a lecture. If she decided to take a Break in stage 1, then in stage 2 she chooses among 3 "ranges" of break lengths: (1) Short: 1 hour to 12 hours, (2) Medium: 12 hours to 36 hours, and (3)

Long: 36 hours to 5 weeks. In summary, there are 7 choice options: Marketing Quiz (MQ), Marketing Lecture (ML), Operations Quiz (OQ), Operations Lecture (OL), Short Break (B1), Medium Break (B2), and Long Break (B3).[4]

Figure 26: Two-stage decision process
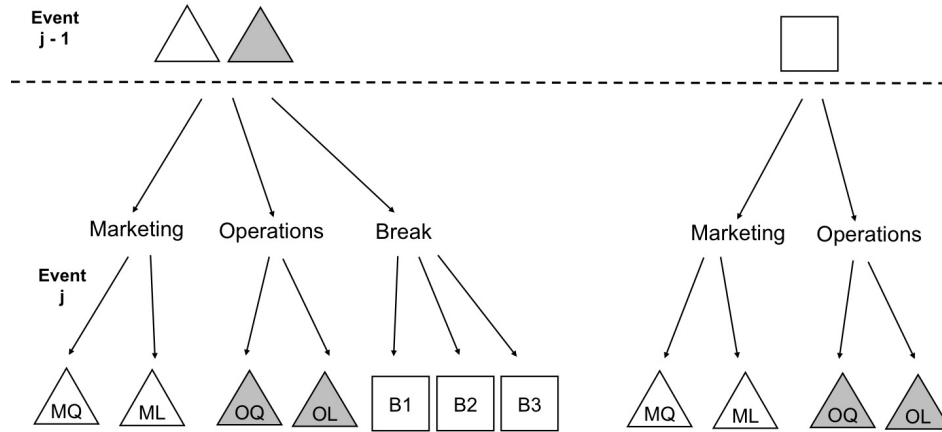
Note that Figure 26 depicts the available choices at event j as being conditional on the choices made at event $j - 1$. If either Marketing or Operations quizzes or lectures were chosen at event $j - 1$ (as represented by the white and grey triangles), then at event j, the individual can choose between Marketing, Operations, and Break. However, if a Break of any length was chosen at event $j - 1$ (as represented by the white square), then at event j, the individual can choose between *only* Marketing and Operations. This prevents individuals from taking multiple breaks in a row, which we have no way of identifying in the data.

Figure 27 shows the frequency of each choice across the entire course for both the sequential and simultaneous data. In general, individuals chose Marketing more often than Operations,

---

[4]The shortest Marketing lecture was 2:09, while the longest Marketing lecture was 19:57. The shortest Operations lecture was 6:22, while the longest Operations lecture was 26:16. So we coded all breaks that existed at least twice as long as the longest lecture. The cutoff locations and robustness checks are further discussed in Appendix A.8.

which is consistent with the fact that Marketing has more lectures and quizzes, as well as less attrition over time. Individuals also chose lectures more often than quizzes, consistent with the fact that there are simply more lectures than quizzes, and lectures are needed to accumulate knowledge to pass the quizzes. The Short Break (B1) was the most common among the three break lengths, followed by the Long Break (B3) and Medium Break (B2). In addition, individuals in the simultaneous dataset had on average more submissions or visits to the lectures and quizzes of both courses compare to individuals in the sequential dataset, but nearly the same number of breaks of different lengths.

Figure 27: Number of times each choices was made, averaged across individuals.
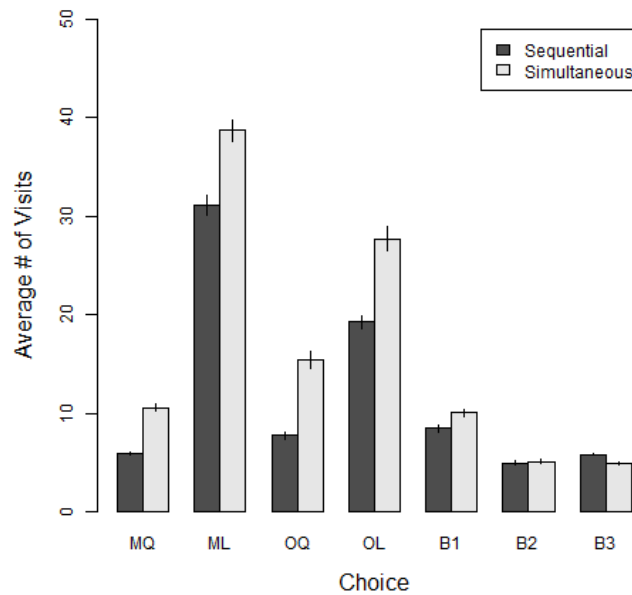
We assume three restrictions on the choice set that impact our model and the likelihood function. First, if no quizzes were available in a particular week (i.e., Operations during week 1) or all quiz attempts were "used up" (i.e., Marketing quizzes were limited to 3 attempts each), then these options were not available in stage 2 of the event.[5] Second,

---

[5]In other words, we use dynamic choice sets, noting that in our case the choice set is observed unlike in work that models latent consideration sets (Ben-Akiva and Boccara 1995).

observed activity of individuals is restricted to the 5 weeks following their individualized start times, and so longer break length options are unavailable when there is not enough time remaining. For example, on the last day it would not be possible to choose the Long Break of 36 hours to 5 weeks. Third, as mentioned earlier, individuals can't take two consecutive breaks, as that would be categorized as a longer break (compare the left and right hand sides of Figure 26).

4.2.4. Model Overview: Stage 1 Choices

Figure 28 depicts the constructs that contribute to the utilities of the options available in stage 1 of each individual's decision process (i.e., $u_M[j]$, $u_O[j]$, and $u_B[j]$). In stage 1 of each event j, individuals choose between Marketing, Operations, and Break.

Figure 28: Stage 1 choice at event j.
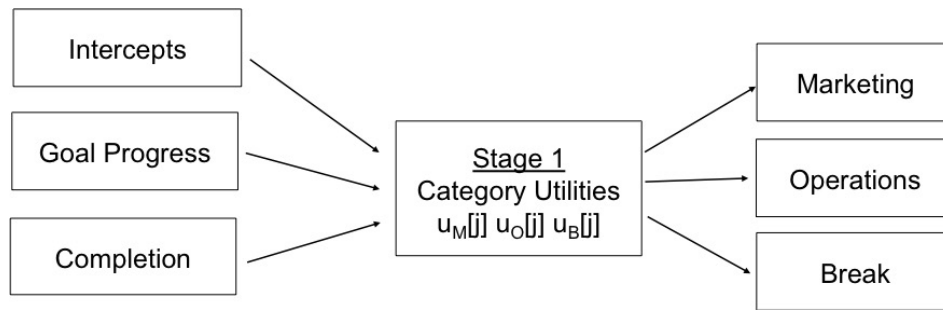
First, we include intercepts to capture the baseline propensity of individuals to engage in Marketing, Operations, and Break. In Figure 27, we saw that Marketing overall has a higher choice share compared to Operations, which can be accounted for by the relative values of the intercepts for Marketing and Operations. In addition, these intercepts can account for the individual-level propensity to consume course content rather than take a break.

Second, we allow the utility of Marketing and Operations to depend on goal progress, as measured by the percentage of available lectures and quizzes within a course that the individual has visited at least once so far. Within the context of online learning, it seems particularly appropriate to incorporate how the utility of consuming course content is in-

fluenced by individuals' goals to complete the course. Consumers may experience utility from consuming course content, and exhibit patterns of increasing motivation for consumption as they approach completion of the course, which is consistent with the goal-gradient hypothesis that predicts that individuals are more motivated to engage in a task as they approach an end goal (Kivetz, Urminsky, and Zhang 2006). Thus, one reason why individuals may be temporal bingers is that the Marketing and Operations choices become more attractive (relative to taking a Break) as they approach the goal of visiting each lecture and quiz at least once within a given course. Similarly, individuals may be content binging because Marketing, for example, becomes more attractive (relative to Operations and Break) as individuals approach the goal within Marketing, with the same intuition applying to Operations.

On the other hand, consumers may also exhibit patterns of decreasing motivation because of satiation or hedonic adaptation, which has been observed in prior work for both pleasant and unpleasant experiences (Inman 2001; McAlister 1982; Nelson and Meyvis 2008; Nelson, Meyvis, and Galak 2009). Our model flexibly allows for both increasing and decreasing motivation to consume course content over time, as well as non-monotonic patterns. For example, the "stuck in the middle" effect is an extension of the goal-gradient theory that suggests a consumer's reference point for progress may shift from the beginning to the end state, resulting in a dip in motivation in the middle of the course (Bonezzi, Brendl, and De Angelis 2011). Alternatively, individuals may exhibit an inverse U-shaped utility because early on, when they first committed to taking the course, they may have overestimated their future slack for time resources (Zauberman and Lynch 2005), resulting in a peak and then dip in utility with progress.

Third, we allow our model to capture what might happen to the utilities of the courses once individuals actually reach "completion" of their goal. In extant research that documents goal progress, the task is either "terminal," that is the researcher stops observing the individual's consumption patterns after task completion, or the task "resets" so the individual

starts again with a new goal (e.g., a customer starts a second coffee loyalty stamp card after completing the first; see Kivetz, Urminsky, and Zhang 2006). In our data, individuals are allowed to revisit any available lectures and quizzes throughout the 5 weeks of the course, so before new content is released to "reset" the goal, we would actually expect individuals to experience a "crash" in their likelihood to consume course content once they have visited all available lectures or quizzes within a particular course. This "completion effect" might contribute to content binging: once individuals complete their goal, they may then switch entirely to consuming content from the other course until new content is made available in a subsequent week.

4.2.5. Model Overview: Stage 2 Choices

In stage 2 of each event j, individuals choose among a set of options, with the available choices conditional on the choice made in stage 1. Figure 29 outlines the constructs that determine the utilities of Marketing Lectures and Quizzes, given that Marketing was chosen in stage 1. These constructs also contribute to the utilities of Operations Lectures and Quizzes, given the choice of Operations in stage 1. The utilities of Short, Medium, and Long Breaks are determined by the intercepts only (under the assumption of a myopic individual with no forward-looking steps, and is generalized under our model of boundedly rational forward-looking behavior).

When taking online courses, an important purpose for consumers is to gain expertise or accumulate knowledge (Alba and Hutchinson 1987; Camerer and Ho 1999) as they engage in the course material. Some individuals may be purely interested in gaining knowledge by watching lectures (although for some individuals, watching lectures may actually be an indication of a *lack* of knowledge), while others may be more interested in passing the course to earn a certificate, which may be done by passing the quizzes. So we propose that individuals' desire to take quizzes vs. watch lectures may be determined by their accumulated knowledge, as measured by how many lectures they have watched so far or beliefs about their quiz-taking abilities, which are updated in a Bayesian way as they get feedback from

taking quizzes. The details of the updating equations are presented in Section 4.4.6.

Figure 29: Stage 2 choice at event j, conditional on Marketing being chosen in Stage 1.
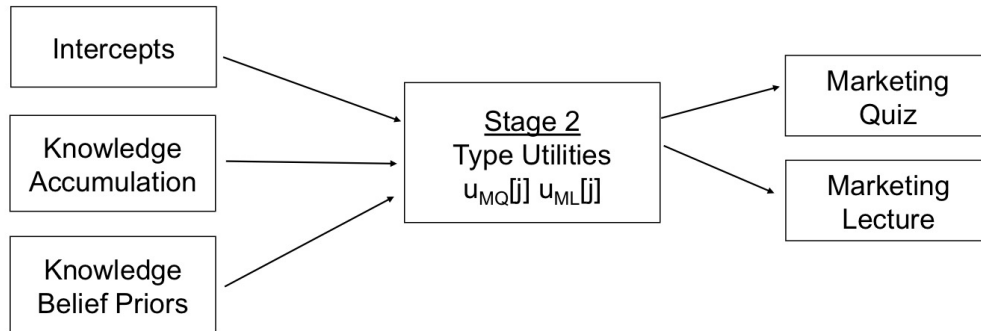
4.2.6. Model Overview: Boundedly Rational Consumers

Rather than assuming that individuals are solving a fully forward-looking utility maximization problem, we assume that they are forward-looking in a boundedly rational way, such that they may take into account the discounted expected utilities from one or two decisions into the future. As depicted in Figure 30, our model allows individuals to make choices as if they were capable of thinking at most one or two stages ahead. This "forward-lookingness" can be extended farther into the future (i.e., three stages ahead, four stages ahead, etc.). The computation of utilities farther and farther into the future quickly becomes very intensive, and classic experimental work has demonstrated that most people are limited to one or two "thinking steps" ahead (Camerer, Ho, and Chong 2004) and that limited (versus infinite) time-horizon models are more accurate in capturing real behaviors (Gabaix et al., 2006).

In the myopic version of our model, we assume that individuals are maximizing the utility

Figure 30: Boundedly rational forward-looking consumers during events j and $j + 1$.
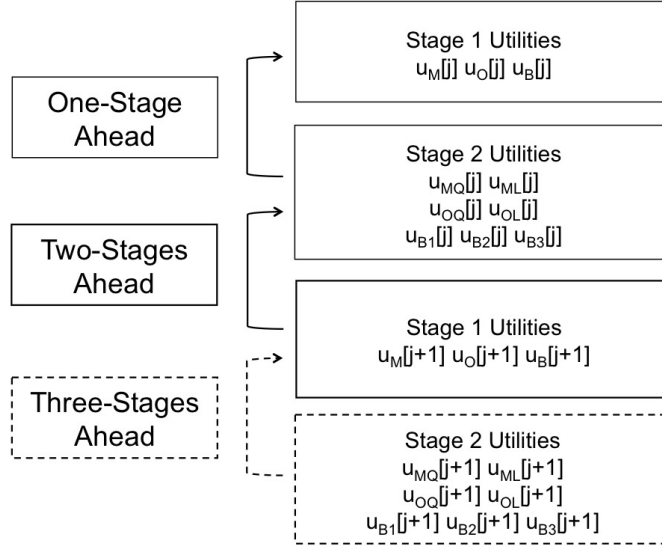
of their choices at each stage without consideration of how the current choice may impact the utilities of choices at future stages. In other words, decisions made in either stage 1 or stage 2 of a particular event j does not depend on the expected utilities of choices in future stages or events.

Allowing individuals to be forward-looking in a boundedly rational way (i.e., thinking one or two steps ahead) means that individuals may take into account the utility of choices in future stages when making choices at the current stage. Consumers may be thinking "One-Stage Ahead" such that during stage 1 of event j, they take into account the expected utilities of the choices that could be made in stage 2 of event j, but they also assume that at stage 2 of event j they will be myopic. Consumers may also be thinking "Two-Stages Ahead" such that during stage 1 of event j, they take into account the expected utilities of the stage 2 choices, but now also assume that at stage 2 they will take into account the expected utilities of the stage 1 choices at the *next* event $j + 1$ instead of being myopic. For example, when an individual in our dataset is making a decision about whether to watch another lecture video or to take a break, she may take into consideration the greater

expected future utility she can obtain from getting closer to course completion by taking a quiz in the next decision period.

We are particularly interested in quantifying how the expected utilities from the next stage impact the decisions in the current stage, which we accomplish by estimating a coefficient (or "discount factor") on the expected utilities, as described in the subsequent sections. Thinking one or two stages ahead may contribute to temporal binging; for example, if individuals anticipate greater utilities from stage 2 (or even stage 1 of the next event) if they engage in Marketing or Operations in stage 1 of the current event, then they may be less likely to choose Break in stage 1. In the Estimation section, we will present results comparing the myopic vs. One-Stage Ahead models. We note that our model can be generalized to individuals thinking more stages ahead, as shown in Figure 30.

In summary, the stage 1 utilities of Marketing, Operations, and Break depend on the following constructs: Intercepts, Goal Progress, Completion, and Time (see Figure 28). The stage 2 utilities of Quiz vs. Lecture (given that either Marketing or Operations was chosen in stage 1) or Short vs. Medium vs. Long Break (given that a Break was chosen in stage 1) depend on the Intercepts, Knowledge Accumulation, and Knowledge Belief Priors (see Figure 29). Moreover, individuals may be thinking One-Stage or Two-Stages Ahead when incorporating the utilities from future stages when making decisions in the current stage (see Figure 30).

4.3. Descriptive Analysis

Before we introduce our formal model, we present results from exploratory analyses of the data to provide descriptive evidence of how the utilities for the Marketing and Operations courses might vary with goal progress, as well as descriptive evidence of binge consumption, specifically distinguishing between temporal binging and content binging.

First, to demonstrate that our Goal Progress construct can be disentangled from time

Figure 31: Stage 1 choice shares between Marketing and Operations vary with goal progress.
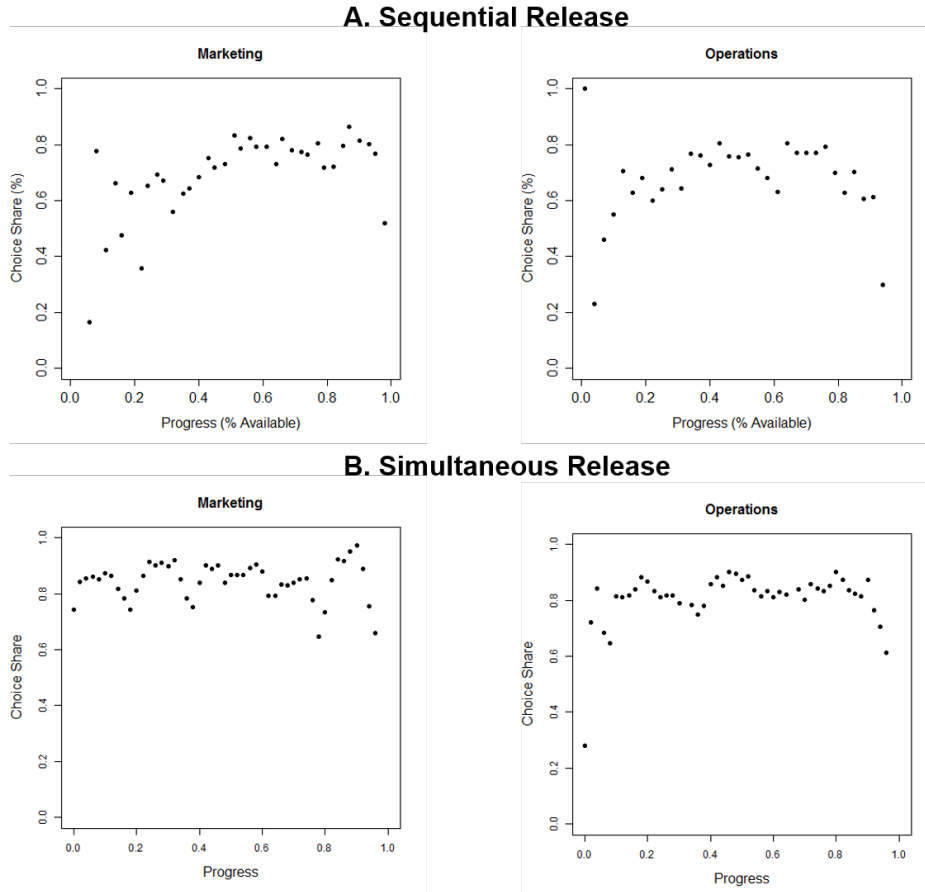
and/or state dependence, we look at how the stage 1 choices shares for Marketing and Operations (vs. Break) vary with "Progress," as measured by the percentage of available lectures and quizzes in each course. Figure 31 plots these choice shares for both the sequential and simultaneous data in Panels A and B, respectively. We see that for both datasets, the choice shares for Marketing and Operations both exhibit non-monotonic patterns. Importantly, a simple state dependence model would not be able to explain these systematic changes in choices shares. We also see that there is a large drop-off in choice shares when Progress reaches 1, which illustrates the importance of the Completion construct, which captures the drop-off engagement when an individual has submitted all content at least once. We verify this descriptive evidence by looking at the parameter estimates of Goal Progress and Completion in the next section.

Next, using metrics of temporal and content binging, we can test whether an individual is observed to be a temporal binger, a content binger, both, or neither. These metrics can then be used to test whether our model is able to capture the temporal and content binging patterns of the observed data, using posterior predictive checks (Gelman et al. 2014), as described in Sections 4.5 and 4.6.

Binging has been described as a phenomenon where consumers actually derive increasing returns to consumption, which is akin to the concepts of "fluency" in learning and judgment (Whittlesea and Leboe 2000; Greifeneder, Bless, and Pham 2011) and "flow" in video games (Chou and Ting 2003), as well as addiction in extreme cases (Becker and Murphy 1988; Gordon and Sun 2015). The way we define binge behavior in this chapter is agnostic to the underlying causes of this behavior, and we distinguish between temporal binging, which depends on the amount of content that is consumed in a single sitting, and content binging, which depends on the type of content that is consumed.

Temporal and content binging are both related to patterns of consumer behavior that have been modeled in other marketing settings. Temporal binging is analogous to the "clumpiness" of a series of purchase incidences, which takes into account the total number of events and inter-event times (Zhang, Bradlow, and Small 2014). On the other hand, content binging would mean that the individual consumes lectures from the same course in succession with few switches between courses, disregarding how spread out the lectures and quizzes are across time. The distinction between content binging vs. content savoring can be compared to inertial vs. variety-seeking behavior within the brand choice literature (Givon 1984; Kahn, Kalwani, and Morrison 1986; Chintagunta 1998).

In order to assess each individual's degree of temporal binging, we can look at the average length of consecutive quiz/lecture events ("runs") with no breaks. Longer runs correspond to more temporal binging. To assess the degree of content binging, we can look at the percentage of times individuals *did not* switch between Marketing and Operations content, given the opportunity. Larger no-switch percentages correspond to more content binging.

Figure 32 plots the average run lengths against the average percentage of non-switches for learners in the sequential dataset (one dot for each person in our sample), which shows that there is a positive relationship between the two binging metrics (r = 0.43, t(465) = 0.47, p < 0.001) such that individuals who are temporal bingers are also likely to be content bingers. However, note that this positive relationship is mainly driven by the long tails of each distribution.

Figure 32: Distribution of average run length and percentage of non-switches
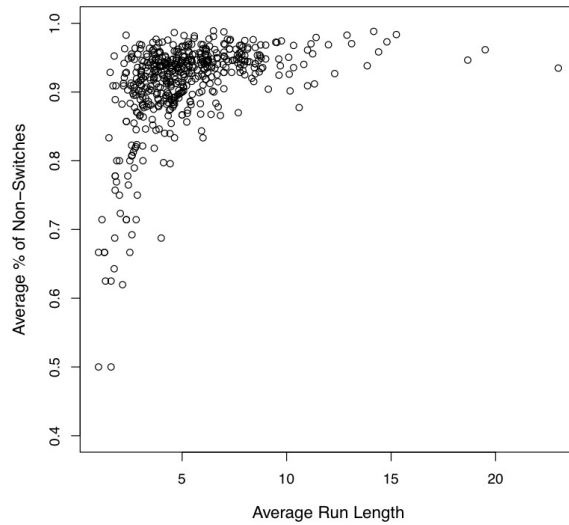
In order to determine whether or not individuals were statistically significant temporal bingers or content bingers (or both), we first created a null distribution of run lengths and non-switch percentages. We took each individual's sequence of stage 1 event choices (Marketing, Operations, or Break) and randomly permuted the sequence 10000 times, with two constraints: (1) breaks could not occur consecutively (as mentioned in the choice overview), and (2) events were permuted only within weeks to account for the sequential release of content. We excluded the 9% of individuals in our sample for whom there were no such permutations.

87

To determine whether or not individuals were temporal bingers, we calculated the percentage of random permutations with (strictly) longer average runs compared to the individual's actual sequence. To determine whether individuals were content bingers, we calculate the percentage of random permutations with (strictly) more Marketing/Operations switches compared to the individuals' actual sequence of events. These percentages represent the likelihood of observing a random sequence of choices with more binging behavior compared to the individual's actual sequence, which can be interpreted as a p-value. Figure 33 shows the distributions of the logit-transformed p-values.

73% of individuals had a p-value of 0 for temporal binging, that is, there were actually no permutations that had longer runs compared to their observed data, while 77% of individuals had a p-value of 0 for content binging, so none of their permutations had fewer switches. Since the logit-transform of 0 is negative infinity, we represent these individuals by the spike at -10 in each plot in Figure 33. We see that nearly all values in both plots fall below -3, which corresponds to a p-value of 0.05. This implies that there is evidence that most individuals were both temporal bingers *and* content bingers, according to our metrics. Specifically, we find that 82% of the individuals in our sample were statistically significant temporal bingers, 86% were statistically significant content bingers, and 69% were both.

4.4. Model and Notation

Our model captures how individuals choose to watch lectures, take quizzes, and take breaks as a series of discrete events. First, we outline the notation for the two-stage decision process that occurs at each event, and how individuals maximize over the utilities of the options at each stage. Next, we describe the parameterization of each construct that contributes to the utilities of the options.

Figure 33: Distributions of logit-transformed p-values for temporal and content binging
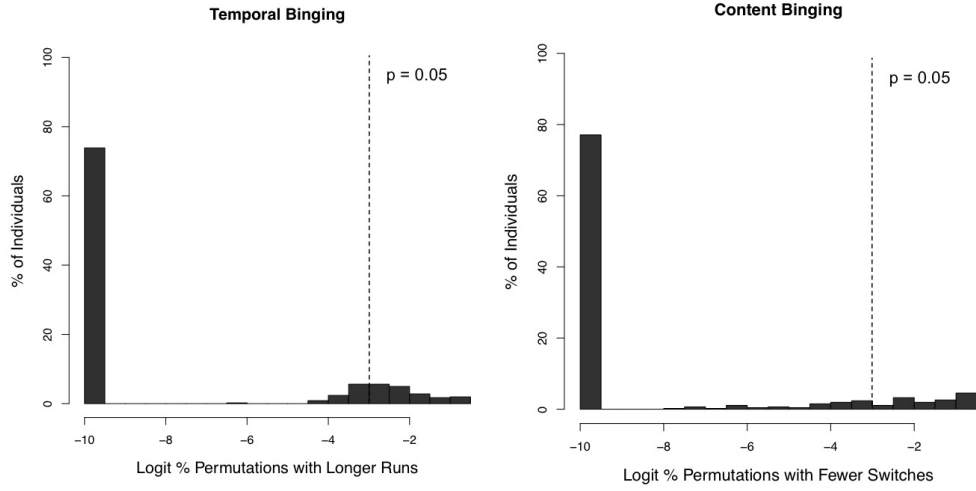
4.4.1. Two-Stage Decision Process

For each individual $i = 1, ..., I$ we observe a sequence of events $j = 1, ..., J_i$. At each event j, the individual makes a two-stage decision (see Figure 27) that ultimately results in choosing one of 7 options: Marketing Quiz (MQ), Marketing Lecture (ML), Operations Quiz (OQ), Operations Lecture (OL), Short Break (B1), Medium Break (B2), and Long Break (B3). Note that although we treat these choice events as discrete, each event j is observed to occur at a continuous calendar time $t[j]$, which we will take into account when simulating choice pathways from the estimated parameters to form posterior predictive checks.

Let $S_1[j]$ represent the choice made in stage 1 of event j between Marketing, Operations, and Break. Let $S_2[j]$ represent the choice made in stage 2 of event j between Marketing Quiz/Lecture, Operations Quiz/Lecture, or Short/Medium/Long Break. The following expression gives the likelihood of an individual's sequence of observed choices, given the individual's parameters $\theta_i = \{\beta_i, \alpha_i, \delta_i, \pi_i, \eta_i, \gamma_i\}$. The product across all individuals then

results in the full likelihood expression.

$$\mathcal{L}(\boldsymbol{\theta}_i) = \prod_{j=1}^{J_i} P(S_2[j]|S_1[j]) \times P(S_1[j]) \tag{4.1}$$

Table 8 outlines the parameters of the model. Note that although we allow all parameters to be heterogeneous across individuals in a Bayesian fashion (see the Estimation section for details), for ease of exposition we suppress the individual-level subscripts on the parameters and variables in the remainder of this section, except where noted.

Table 8: Summary of model parameters

Construct	Parameter	Variable	Description
Intercepts	β_{0M}, β_{0O}	–	Stage 1
	α_{0M}	–	Stage 2 (Marketing Quiz)
	α_{0O}	–	Stage 2 (Operations Quiz)
	δ_1	–	Stage 2 (Short Break)
	δ_2	–	Stage 2 (Medium Break)
Goal Progress	β_{1M}, β_{1O}	$G_M[j]$ or $G_O[j]$	Linear
	β_{2M}, β_{2O}	$G_M[j]^2$ or $G_O[j]^2$	Quadratic
Completion	β_{3M}, β_{3O}	$\mathbb{1}(G_{MQ}[j]=1)$ or $\mathbb{1}(G_{OQ}[j]=1)$	All available Quizzes visited
	β_{4M}, β_{4O}	$\mathbb{1}(G_{ML}[j]=1)$ or $\mathbb{1}(G_{OL}[j]=1)$	All available Lectures visited
Knowledge	α_{1M}, α_{1O}	$C_{ML}[j]$ or $C_{OL}[j]$	Consecutive Lectures
	α_{2M}, α_{2O}	$\pi_M[j]$ or $\pi_O[j]$	Quiz-Abilities
	π_{0M}, π_{0O}	–	Initial beliefs (Mean)
	η_{0M}, η_{0O}	–	Initial beliefs (Precision)
Forward-Looking	γ	E	One-Stage Ahead
	ρ	F	Two-Stages Ahead

In stage 1 of event j, the individual is maximizing over the utilities of engaging in Marketing ($u_M[j]$), engaging in Operations ($u_O[j]$), and taking a Break ($u_B[j]$). The utility of Marketing in stage 1 of event j can be represented by the following equation (as depicted

in Figure 28).

$$\text{Stage 1: } u_M[j] = \text{Intercepts}_M[j] + \text{Goal Progress}_M[j]$$
$$+ \text{Completion}_M[j] + \text{Time}_M[j] + \epsilon_M[j] \tag{4.2}$$

If the individual chose Marketing in stage 1 of event j, then in stage 2 she maximizes over the utilities of taking a Marketing quiz ($u_{MQ}[j]$) versus watching a Marketing lecture ($u_{ML}[j]$), and likewise if she chose Operations or Break in stage 1 instead. The utility of a Marketing Quiz in stage 2 of event j can be represented by the following equation (as depicted in Figure 29). Note that the Knowledge construct represents both the knowledge accumulated and the priors on knowledge beliefs.

$$\text{Stage 2: } u_{MQ}[j] = \text{Intercepts}_{MQ}[j] + \text{Knowledge}_{MQ}[j] + \epsilon_{MQ}[j] \tag{4.3}$$

Assuming that the error terms follow a type-1 extreme value distribution, the choice probabilities in stage 1 and stage 2 can each be formulated as a multinomial logit between the utilities of the available options.

We describe our model as being "one-stage" or "two-stages" forward-looking, without loss of generality, as more stages can be added (as depicted in Figure 30). Let γ represent the "One-Stage Ahead" term, or how much the individual weighs the expected utilities from stage 2 in the utilities of stage 1 during event j. We assume that when making a decision at stage 1 of event j, individuals are able to infer the maximum of the expected utilities of stage 2. For example, if the individual chose Marketing at stage 1, then the maximum expected utility during stage 2 would be given by the following expression:

$$\text{E} = \log\left(e^{u_{\text{MQ}}[j]} + e^{u_{\text{ML}}[j]}\right) \tag{4.4}$$

Thus, if individuals are thinking One-Stage Ahead, then γE is added to the utilities of the

stage 1 options.[6] In this way, our model is similar to a nested logit in that the 7 options (MQ, ML, OQ, OL, B1, B2, B3) are grouped into 3 "nests" (Marketing, Operations, Break), and so the probability of choosing one of the 7 options is the probability of the nest multiplied by the probability of choosing the option, conditional on the nest (see Equation 4.1). Just like the nest choice in a nested logit, the stage 1 choice in our model is an intermediate outcome, while the stage 2 choice (i.e., within a nest) is the end outcome of the consumer's decision process. The One-Step Ahead coefficient γ in our model is analogous to the correlation in unobserved factors within nests in a nested logit, while the expected maximum of the stage 2 utilities, E, is analogous to the inclusive value term.

Let ρ represent the "Two-Stages Ahead" parameter, or how much the individual weighs the expected utility of her choices in stage 1 of the next event $j + 1$, when making a choice at stage 2 of the current event j. Let $u_M[j + 1]$, $u_O[j + 1]$, and $u_B[j + 1]$ represent the utility of choosing either Marketing, Operations, or Break at the next event $j + 1$. Note that these utilities will vary depending on the stage 2 decision at event j. The following expressions give the expected maximum of the stage 1 utilities at event $j + 1$:

$$F = \log\left(e^{u_M[j+1]} + e^{u_O[j+1]} + e^{u_B[j+1]}\right) \tag{4.6}$$

By adding this Two-Stage Ahead ρF to the utilities of the choices in stage 2 of event j, we can allow the model to account for individuals thinking two stages ahead, such that they take into account the utilities of stage 1 of event $j + 1$ when making a decision at stage 2 of event j, and this propagates up to the utilities of stage 1 of event j through the One-Stage Ahead term γE. We next describe each of the model components in turn.

---

[6]Alternatively, the One-Stage Ahead term can be formulated as the expected mean utility rather than the expected maximum utility, in which case the expressions for E would be the following:

$$E = \frac{e^{u_{MQ}[j]} u_{MQ}[j] + e^{u_{ML}[j]} u_{ML}[j]}{e^{u_{MQ}[j]} + e^{u_{ML}[j]}} \tag{4.5}$$

4.4.2. Intercepts

The intercepts β_{0M} and β_{0O} represent the baseline utility of choosing Marketing and Operations, respectively, in stage 1. The intercept for Break is set to 0 for identification. The intercepts α_{0M} and α_{0O} represent the inherent utility of choosing a Quiz in stage 2, with the utility of Lectures set to 0. The intercepts δ_1 and δ_2 represent the baseline utility of choosing a Short or Medium Break, with the utility of a Long Break set to 0.

4.4.3. Goal Progress

Let $G_M[j]$ and $G_O[j]$ represent the proportion of all available quizzes and lectures in Marketing and Operations, respectively, that the individual has visited at least once by event j. In the sequential data the number of available lectures and quizzes changes each week in both Marketing and Operations, while in the simultaneous data the number of available lectures and quizzes remains the same across all weeks.

We define each individual's "goal" to be to visit all available lectures and quizzes. (Note that this is a general goal and is inclusive of more specific goals, such as passing all quizzes in order to obtain a certificate.) To capture the effect of Goal Progress, we estimate the coefficients $\{\beta_{1M}, \beta_{1M}\}$ and $\{\beta_{2O}, \beta_{2O}\}$ on the proportion and squared proportion of visited quizzes and lectures for each course. The following expression gives the Goal Progress effect for Marketing at event j.

$$\text{Goal Progress}_M[j] = \beta_{1M} G_M[j] + \beta_{2M} G_M[j]^2 \tag{4.7}$$

Depending on the shape of the quadratic function, individuals may become more or less motivated to engage in content from a course as they approach completion, or there may be a non-monotonic pattern. Figure 34 illustrates the values of the linear and quadratic coefficients that would result in different curves for how utility changes relative to progress.

Figure 34: Illustration of the Goal Progress construct

A. Goal Progress Curves

B. Corresponding Linear and Quadratic Values
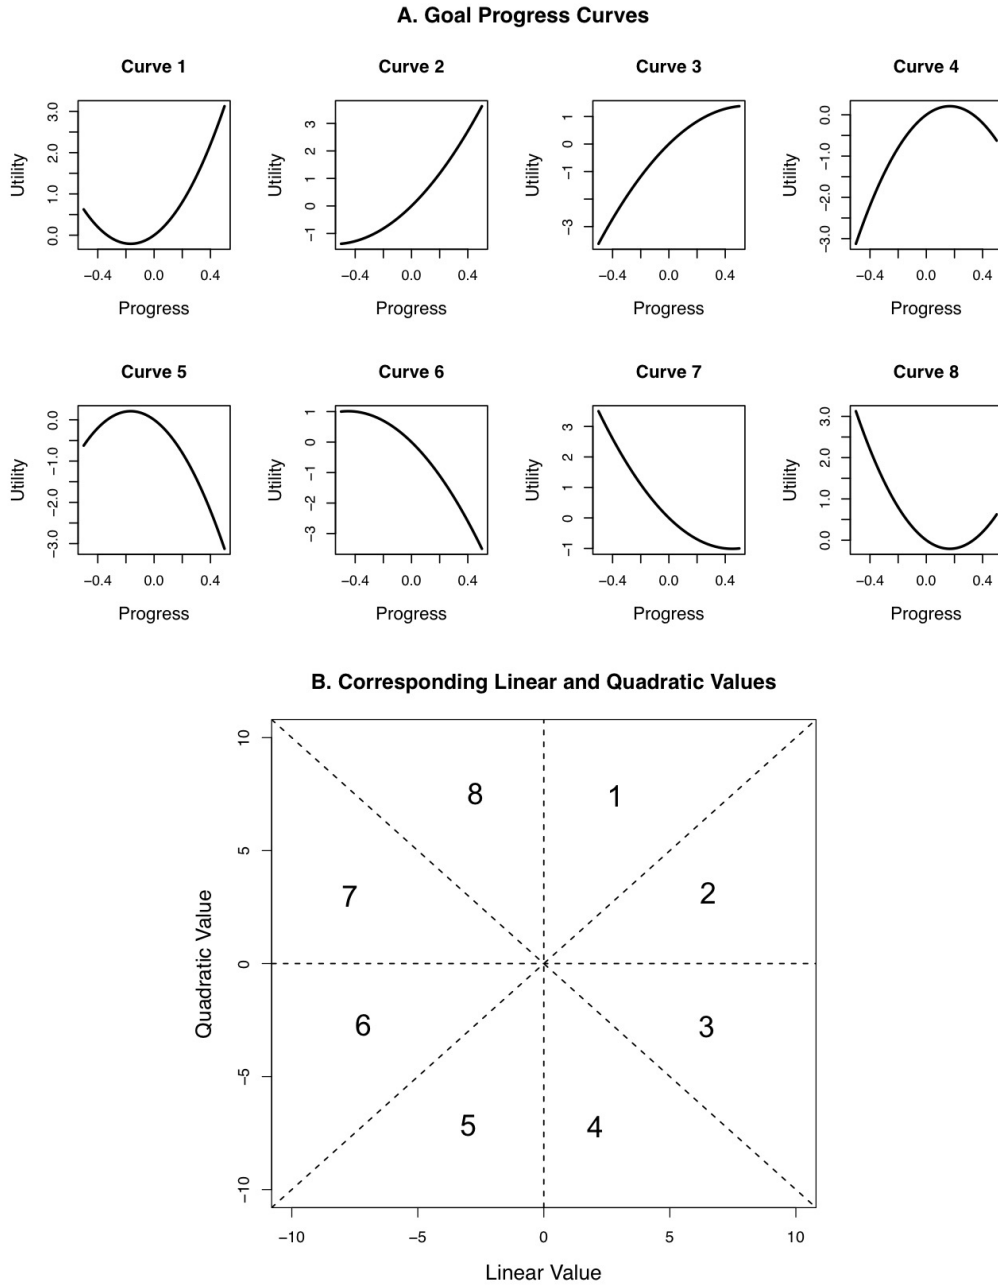
Figure 34 panel A shows 8 possible shapes of how the utility of a course option changes with goal progress, which represents the percentage of quizzes and lectures visited at least once so far (centered at 0 to range between -0.5 and 0.5). The curves may be monotonically increasing or decreasing, either exponentially or logarithmically, or be non-monotonic. Curves

2 and 3, for example, illustrate the goal gradient effect (Kivetz, Urminsky, and Zhang 2006) where utility is monotonically increasing with progress. Curves 6 and 7 illustrate satiation or fatigue (Inman 2001; McAlister 1982). Curves 1 and 8 illustrate the "stuck in the middle" effect as individuals switch from monitoring their progress relative to the initial state to the end state (Bonezzi, Brendl, and De Angelis 2011). Curves 4 and 5 illustrate mispredictions in future time slack (Zauberman and Lynch 2005).

Figure 34 panel B shows the corresponding values of the linear (β_{1M} and β_{1O}) and quadratic (β_{2M} and β_{2O}) coefficients that would result in each of the 8 Goal Progress curves. Using each individual's parameter estimates, we can classify individuals in our sample as different types of Goal Progress learners, similar to how Gilbride and Allenby (2002) classify individuals by the type of screening rules they use during choice.

4.4.4. Completion

As demonstrated in the descriptive analyses, we also expect individuals to experience a "crash" in their likelihood to consume course content once they have visited all available lectures ($G_{ML} = 1$ or $G_{OL} = 1$) or quizzes ($G_{MQ} = 1$ or $G_{OQ} = 1$) within a particular course at least once, in other words when their Progress equals 1. Since the Goal Progress component of our model is not able to capture this "completion effect," we use indicator variables, with coefficients $\{\beta_{3M}, \beta_{4M}\}$ for Marketing and $\{\beta_{3O}, \beta_{4O}\}$ for Operations. Depending on the sign of the coefficients, "completing" all quizzes or all lectures (i.e., visiting them all at least once) may either result in an increase or decrease in utility for the respective course. The following expression gives the Completion construct for Marketing at event j.

$$\text{Completion}_M[j] = \beta_{3M}\mathbb{1}(G_{MQ}[j] = 1) + \beta_{4M}\mathbb{1}(G_{ML}[j] = 1) \tag{4.8}$$

Note that for the sequential data, after new content is released each week, the Goal Progress values are reset, so individuals may vary between states of completion and non-completion. For the simultaneous data, since all the content is available from the start, once an individual

has completed all the quizzes or lectures in a course, she remains in the completion state for that subset of options.

4.4.5. Knowledge Accumulation

If individuals chose either Marketing or Operations in stage 1, then in stage 2 they choose between a Marketing Lecture/Quiz or Operations Lecture/Quiz. We allow the utility of taking a Quiz to depend on the accumulation of knowledge in the course, which may come from watching lectures or taking quizzes.

Let $C_{ML}[j]$ and $C_{OL}[j]$ represent the number of consecutive lectures that the individual visited in Marketing and Operations, respectively, by event j, with α_{2M} and α_{2O} as the coefficients. The intuition here is that each quiz corresponds to a certain subset of lectures, and so individuals may choose to watch these lectures leading up to taking the corresponding quiz. In future work, we plan to use content analysis to map out the relationships between specific lectures and quizzes. For now, we use the number of immediately preceding consecutive lectures within a course as a proxy.

Let α_{1M} and α_{1O} be the coefficients on the individual's beliefs about her quiz-taking abilities π_M in Marketing and π_O in Operations. After taking a quiz, the individual Bayesianly updates her beliefs about her quiz-taking abilities with a "signal" (à la Erdem and Keane 1996) that has a mean of the observed quiz score $x[j] \in [0,1]$ and precision ϕ (set to 1 for identification). For example, if the individual took a Marketing quiz at j, then she updates the mean π_M and precision η_M of her quiz-taking abilities for Marketing in the following way, with a parallel updating process for Operations quizzes.

$$\pi_{\mathrm{M}}[j+1] = \frac{\eta_{\mathrm{M}}[j]}{\eta_{\mathrm{M}}[j] + \phi}\pi_{\mathrm{M}}[j] + \frac{\phi}{\eta_{\mathrm{M}}[j] + \phi}x[j] \tag{4.9}$$

$$\eta_{\mathrm{M}}[j+1] = \eta_{\mathrm{M}}[j] + \phi \tag{4.10}$$

The following gives the expression for Knowledge within Marketing, which is an empirically

determined weighted average of knowledge accumulated from watching lectures and getting feedback from quiz scores:

$$\text{Knowledge}_{MQ}[j] = \alpha_{1M}C_{ML}[j] + \alpha_{2M}\pi_M[j] \qquad (4.11)$$

4.4.6. Knowledge Belief Priors

In order to compute the beliefs about quiz-taking abilities in Equations 4.10 and 4.11, individuals must have some initial beliefs at event $j = 1$, characterized by the prior means $\pi_M[1]$ or $\pi_O[1]$ and precisions $\eta_M[1]$ or $\eta_O[1]$.

First, we can assume that for all individuals the prior mean is 0.5 (i.e., 50% correct on quizzes) and the prior precision is 1 (i.e., equivalent to having taken one quiz, since we set the quiz precision to $\phi = 1$).

Alternatively, we can estimate the prior beliefs, with means $\pi_M[1] = \pi_{0M}$ or $\pi_O[1] = \pi_{0O}$ and precisions $\eta_M[1] = \eta_{0M}$ or $\eta_O[1] = \eta_{0O}$ heterogeneously across individuals. This would capture the differences across individuals in terms of how much the quiz score signals affect their beliefs about their quiz-taking abilities, and thus their probabilities of choosing quizzes over lectures. For example, an individual with a low prior precision may greatly increase her probability of taking a quiz after receiving good feedback, but an individual with high prior precision would not be affected as much.

4.5. Estimation and Model Comparison

We describe our estimation approach, which involves using a hierarchical Bayes procedure to account for the heterogeneity across individuals in the parameter estimates. We also describe our procedure for nested model comparison to demonstrate that each construct built into the utilities contributes to explaining the data in some way. This involves simulating data from the individual-level parameter samples to form posterior predictive checks.

4.5.1. Estimation Procedure

We use hierarchical Bayes estimation to account for unobserved heterogeneity across individuals (Gelman et al. 2014). We assume that each individual's vector of parameters θ_i follows a multivariate normal prior distribution with $\theta_i \sim \text{MVN}(\mu, \Omega)$. Let the mean μ follow a conjugate multivariate normal distribution with $\mu \sim \text{MVN}(\mu_0, \Omega_0)$, where the mean μ_0 is a vector of zeros and the precision Ω_0 is an identity matrix. Then let Ω^{-1} follow a conjugate Wishart distribution with ρ degrees of freedom, which is set to the number of parameters in θ_i plus 3 to make it proper, and an inverse scale matrix R, with the inverse R^{-1} as an identity matrix.

We estimated the parameters using a Markov chain Monte Carlo (MCMC) sampler in the programming language R. For each model, we ran three MCMC chains from different starting values for 3,000 iterations each. We used the first 2,000 iterations as burn-in and we checked for convergence by determining that the Gelman-Rubin convergence statistic was less than 1.2 for all parameters (Gelman and Rubin 1992). After thinning the chains to reduce auto-correlation, we were left with 300 posterior samples for each parameter.

To ensure parameter recovery, we picked a set of means and a covariance matrix with "reasonable" values to form a multivariate normal distribution, from which we drew parameter values and simulated data for 500 individuals (comparable to our observed sample size of 467 in the sequential data and 448 in the simultaneous data). We then estimated the model using this simulated data to determine that we could recover the true parameter values (see results in Appendix A.9.).

4.5.2. Model Comparison

We describe the estimation results for different versions of our model for both the sequential and simultaneous datasets. We start with an intercepts-only model, and then add each construct in succession (i.e., in a series of nested models) to determine whether or not they improve the in-sample model fit. For the sequential data, we estimated a model for each

additional construct. For the simultaneous data, because we did not have a comparable measure of quiz scores, we did not estimate the models for the Knowledge Accumulation and One-Step Forward constructs. However, we do plan to include the simultaneous data quiz scores and additional models in future work.

To compare the fit across the model variations, we calculated the Deviance Information Criterion (DIC) and also simulated data utilizing the MCMC draws to determine whether the models could recover the patterns in the observed data according to a series of posterior predictive checks. We then present and interpret the parameter estimates for the full model, which was also the "winning model" based on the DIC.

Table 9: Models fit statistics for different nested versions of the model

	Sequential			Simultaneous		
Model	DIC	Run Length	% Switch	DIC	Run Length	% Switch
Observed	N/A	3.45	0.13	N/A	5.01	0.07
Intercepts	839598	3.91 (.94)	.43 (.32)	1145069	5.03 (1.85)	.36 (.32)
Goal Progress	837357	3.89 (1.01)	.39 (.29)	1137386	5.13 (1.93)	.30 (.27)
Completion	833122	3.36 (1.19)	.35 (.24)	1132653	4.96 (1.62)	.28 (.24)
Knowledge	831548	3.52 (1.20)	.34 (.25)	N/A	N/A	N/A
One-Stage Ahead	830934	3.46 (1.41)	.34 (.24)	N/A	N/A	N/A

To determine whether adding the different constructs improves the fit of the model, compared to a baseline "Intercepts-only" model, we calculated the DIC for each version of the model, as shown in Table 9. We see that for the sequential data, adding the parameters for the Goal Progress, Completion, Knowledge Accumulation, and One-Step Ahead constructs improves the in-sample model fit in terms of decreasing the DIC. For the simultaneous data, adding the parameters for the Goal Progress and Completion constructs also improves the model fit. Note that the fit statistics in Table 9 are for nested models, for example the "Completion" model includes the constructs of Intercepts, Goal Progress, *and* Completion.

Based on the DIC, the winning model for both datasets is the full model with all the constructs (up to One-Stage Ahead for the sequential data and up to Completion for the simultaneous data). In addition to the model fit, we are also interested in how well our model could capture the specific patterns within each individual's sequence of choices, in particular the metrics for temporal binging (i.e., average run length) and content binging (% of non-switches).

To construct a series of posterior predictive checks, we took each individual's set of 300 posterior samples and simulated a sequence of choices, resulting in 300 simulations for each of the 476 individuals in the sequential data and each of the 448 individuals in the simultaneous data. For each individual MCMC sample, we start at $t_j = 0$ for the first event $j = 1$, with the initial conditions being that the individual has not yet consumed any course content. We then simulate their choices until they reach the end of 5 weeks.

Although we treat individual choices as discrete in the model estimation, individuals are also moving through continuous calendar time (for example, for the sequential data, the calendar time determines which lectures and quizzes were available in the Goal Progress construct). This means that each lecture, quiz, and break must take up some amount of calendar time. Although we do not explicitly model the amount of time spent on each of these event choices (besides the break length ranges), in the simulation we can "predict" how much time the individual would have spent by using the distribution of event lengths in the observed data, either via random sampling from the empirical distribution or using the parameter estimates of a regression on relevant variables such as lagged event lengths, week indicators, course progress, etc. This is consistent with the literature that uses rational expectations to generate predictive distributions of endogenous variables (i.e., price, Muth 1961).

Similarly, for the sequential data, we can use a regression to predict quiz scores for when individuals take either Marketing or Operations quizzes, which allows for updating of individual beliefs about quiz-taking abilities in each course. Thus, each time individuals made

100

a choice in our simulation, we were able to assign the choice an event length, which allowed the individual to move forward in calendar time until they reached the end of 5 weeks, as well as a quiz score if the individual took a quiz, which allowed her to update her quiz-taking abilities. (See Appendix A.10. for further details).

After obtaining the simulated choice sequences for each individual sample, we can examine how closely the simulated patterns fit the observed data. First, we look at whether the models can capture the temporal and content binging patterns of the data, as proxied by the average course choice run length (without breaks) and the percentage of switching between Marketing and Operations, respectively. Thus, in addition to the DIC values, Table 9 also compares the observed and simulated average run length and percent switches for each nested version of the model.

We see that the Intercepts model is able to accurately simulate the average run length, but greatly overestimates the percentage of choices for which individuals switch between courses. The addition of the various constructs reduces the switching percentage, bringing it closer to the observed pattern of content binging. Although the simulated percentage is still greater than the observed percentage, we note that there is no state dependence parameter in this version of the model to provide additional short-term "stickiness" in the cross-course decisions akin to brand loyalty (Guadagni and Little 1983), so the patterns of content binging are being captured purely by the theory-driven constructs in our model.

Since we estimated our model using hierarchical Bayes methods, we can also look at whether our model is able to capture the heterogeneity in consumption patterns across individuals. For both sequential and simultaneous datasets, Figure 35 plots the observed vs. simulated values for the run length and % switches. Each point in each plot represents the mean across simulations for one individual, and the solid line represents the 45-degree line. We see that for both sequential and simultaneous datasets, the models are able to simulate the heterogeneity across individuals in terms of run length, with most individuals falling on or near the 45-degree line. However, the models overestimate the percentage of switching,

which could be corrected for by the inclusion of a state dependence parameter.

Figure 35: Individual-level posterior predictive checks for observed vs. simulated data
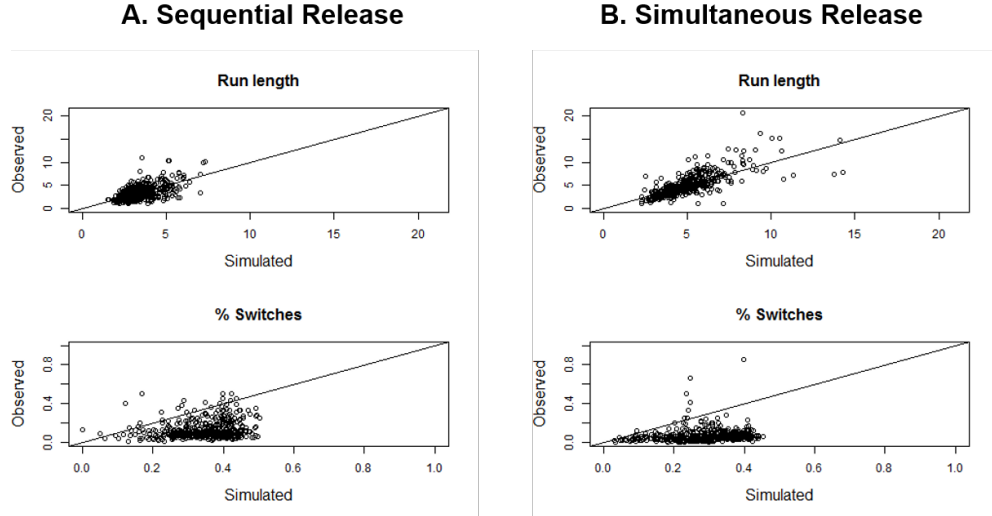
4.6. Parameter Estimation Results

We present the estimated parameters for the full 2-stage model with all constructs. We summarize the posterior distributions for the elements of μ and the diagonal elements of Ω, with each individual's parameters $\theta_i \sim MVN(\mu, \Omega)$.

First, we discuss the parameter estimation results for the constructs that determine the utilities in Stage 1, including the Intercepts, Goal Progress, Completion, and One-Stage Ahead. Then we discuss the constructs that determine the utilities in Stage 2, including the Intercepts and Knowledge Accumulation. Note that in the 2-stage model, for each construct we are able to estimate a set of parameters for both Marketing and Operations.

4.6.1. Stage 1 Constructs

Table 10 shows the results for the estimation of the parameters in Stage 1, while the results for Stage 2 are presented in the next subsection in Table 11. In both tables, we present the mean of the posterior draws of μ, the 95% credible interval of the posterior draws of μ, as

well as the mean of the posterior draws of σ^2, i.e. the diagonal elements of the variance-covariance matrix Ω.

Table 10: Summary of estimated parameters in stage 1

Construct	Parameter	Sequential			Simultaneous		
		μ	95% CI	σ^2	μ	95% CI	σ^2
Intercepts	β_{0M}	.46	[.41, .52]	0.26	.97	[.91, 1.03]	.29
	β_{0O}	.05	[-.03, .12]	0.23	1.49	[1.43, 1.57]	.25
Goal Progress	β_{1M}	.22	[.11, .32]	0.42	.44	[.32, .56]	.40
	β_{2M}	1.06	[.77, 1.31]	1.82	1.74	[1.37, 2.20]	5.31
	β_{1O}	-.03	[-.20, 0.14]	.90	.52	[.36, .71]	.48
	β_{2O}	-3.22	[-3.55, -2.89]	3.36	-5.49	[-5.79, -5.03]	8.98
Completion	β_{3M}	-1.99	[-2.14, -1.84]	1.38	-.33	[-.44, -.21]	.86
	β_{4M}	-.31	[-.40, -.21]	0.34	-1.36	[-1.53, -1.20]	1.58
	β_{3O}	-.46	[-.57, -.36]	.51	-.47	[-.58, -.35]	.70
	β_{4O}	.28	[.18, .37]	.38	-.65	[-.81, -.52]	1.45
One-Stage Ahead	γ	.74	[.65, .82]	0.29	N/A	N/A	N/A

The Intercept parameter estimates capture the relative choice shares of Marketing and Operations (vs. Breaks) for each dataset. The parameters values of the Goal Progress construct allow us to determine which of the 8 curves depicted in Figure 34 describes the change in utility with respect to goal progress for the population as a whole, as well as for each individual. The negative parameter estimates in the Completion construct indicate that the utility for courses drops off significantly after visiting all quizzes or lectures at least once. This effect is especially strong for Marketing quizzes ($\beta_{3M} = -1.99$) in the sequential data and for Marketing lectures in the simultaneous data ($\beta_{3M} = -1.36$). For the sequential data, the One-Stage Ahead construct has a mean of $\gamma = .74$. This indicates that individuals are taking into account a fraction of the utility that could be obtained in stage 2 of the

decision process when they make a choice during stage 1. In future work, we plan to obtain the analogous estimate for the One-Stage Ahead construct for the simultaneous data.

In Figure 36, we plot the posterior means of the linear and quadratic terms within the Goal Progress construct for each individual, with separate parameters for Marketing and Operations. This way, we can categorize each individual as a specific type of learner based on how the utilities of the courses change with progress.

In Marketing, 76% of sequential learners and 79% of simultaneous learners have parameters that classify them as Curve 1 (in Figure 34), which indicates a non-monotonic U-shaped function where utility first decreases and then increases greatly with progress, corresponding to the stuck in the middle effect. In Operations, 90% of sequential learners and 95% of simultaneous learners are classified as Curve 4, which indicates a non-monotonic inverse-U-shaped function where utility first increases and then decreases with progress, corresponding to slack theory. Interestingly, we find that these effects are more extreme (i.e., the parameters are farther from zero) and more varied for the simultaneous data compared to the sequential data.

One potential reason for why Marketing and Operations have different Goal Progress effects is that Operations has more quizzes at the end, including a cumulative final. Thus, the workload towards the ends of the Operations course, in terms of lectures and quizzes, is much higher compared to Marketing. Thus, one interpretation for the inverse U-shaped function for utility with respect to progress is that individuals underestimate the effort needed for Operations at the beginning, but motivation starts to drop off near the end as they run low on time or other resources.

In contrast, individuals in Marketing exhibit a U-shaped function for utility with respect to progress, indicating that they start strong and end strong, with a lull in the middle. The decrease at the beginning may be due to the general trend that learners tend to start the Marketing course first, and then wait a few days or weeks before starting the Operations

course. The Operations course may take away from time and resources that could be spent on Marketing, resulting in the lull in the middle, but as individuals become less motivated to engage in Operations due to its difficulty, they may switch back to the relatively easy Marketing course and exhibit a goal gradient increase in utility as they approach completion.

Figure 36: Goal Progress individual-level parameter estimates for Marketing and Operations

A. Sequential Release
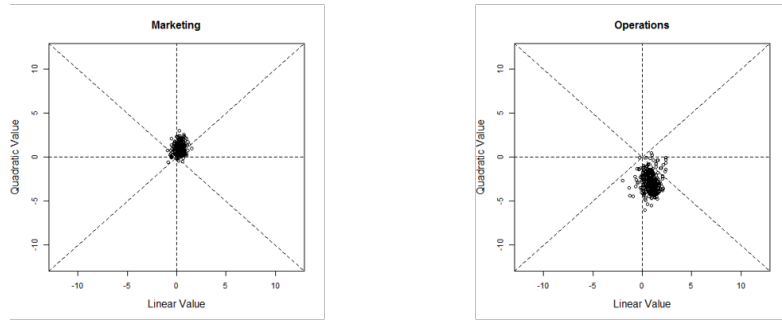
B. Simultaneous Release
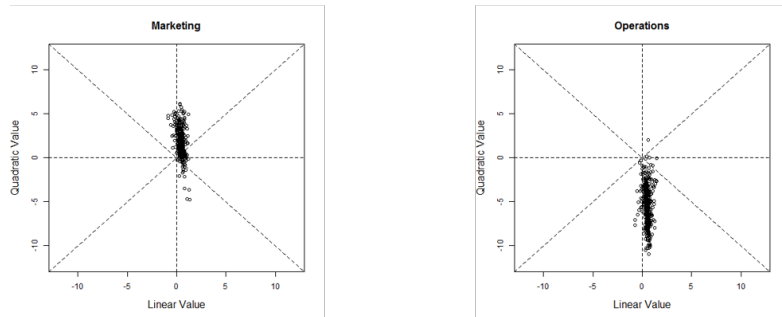
4.6.2. Stage 2 Constructs

Table 11 shows the estimates for the parameters in the stage 2 constructs. Note that we estimated the Intercepts and the Knowledge Accumulation construct parameters for the sequential data, but only the Intercepts for the simultaneous dataset. In future work, we plan to transform or normalize the quiz score data from the simultaneous dataset into values that are comparable to those in the sequential dataset in order to estimate the Knowledge

Accumulation construct parameters for the simultaneous learners.

For the Knowledge Accumulation construct, we set the initial means of the quiz-taking abilities to $\pi_{0M} = \pi_{0O} = 0.5$ and the precision to $\eta_{0M} = \eta_{0O} = 1$. This can be interpreted as individuals having a "weak" prior belief that they have average quiz-taking abilities for the course. In future work, we can estimate $\{\pi_{0M}, \pi_{0O}\}$ and $\{\eta_{0M}, \eta_{0O}\}$ heterogeneously to determine whether there are differences in initial beliefs across individuals and across courses.

Table 11: Summary of estimated parameters in stage 2 construct

Construct	Parameter	Sequential			Simultaneous		
		μ	95% CI	σ^2	μ	95% CI	σ^2
Intercepts	α_{0M}	-2.03	[-2.25, -1.82]	0.75	-1.32	[-1.40, -1.25]	.43
	α_{0O}	-3.15	[-3.41, -2.95]	1.11	-1.04	[-1.19, -.93]	1.26
	δ_1	.16	[.07, .23]	0.47	.54	[.46, .62]	.45
	δ_2	-.30	[-0.38, -.22]	0.34	-.07	[-.15, .02]	.30
Knowledge	α_{1M}	.18	[-.04, .41]	0.68	N/A	N/A	N/A
	α_{2M}	.14	[.10, .17]	0.10	N/A	N/A	N/A
	α_{1O}	2.07	[1.82, 2.35]	1.20	N/A	N/A	N/A
	α_{2O}	0.18	[.14, .21]	0.10	N/A	N/A	N/A

For both the sequential and simultaneous data, the intercepts for quizzes in both Marketing (α_{0M}) and Operations (α_{0O}) are negative, which accounts for the lower choice shares for quizzes in stage 2 of both courses. The intercepts for the different break lengths reflect the higher frequency of short breaks (δ_1) and lower frequency of medium breaks (δ_2), compared to long breaks.

Within the Knowledge construct, depending on the sign of the coefficient, different variables can be interpreted as contributing to either knowledge accumulation or decay. We see that

beliefs about quiz-taking abilities have a positive effect on the utility of taking a quiz in both courses (α_{1M} and α_{1O}). Since individuals are allowed to attempt quizzes more than once, one interpretation is that when individuals receive high quiz scores, their beliefs about their quiz-taking abilities increase, and so they become more motivated to take more quizzes (assuming that there are still available quizzes to take).

Watching more consecutive lectures has a positive effect in both Marketing (α_{2M}) and Operations (α_{2O}). These results are akin to the traditional view that watching more lectures can lead to knowledge accumulation, and therefore a greater likelihood to take quizzes. In future work, we can explore other factors that might lead to knowledge accumulation or decay, such as taking breaks of different lengths or engaging in other courses.

4.7. Counterfactual Simulations

In our analyses, we were able to estimate our model on two datasets that occurred before and after the Coursera platform switched the Marketing and Operations course content from weekly sequential release to on-demand simultaneous release. It is of interest to the firm and the academic institution offering the content how this change to the courses affects individuals consumption patterns. Thus, we conducted a counterfactual simulation where we changed the release structure of the lectures and quizzes in each course, and examine how the distribution of course content visited over time changes.

Figure 37 depicts the cumulative percentage of total visits to Marketing and Operations, averaged across individuals, for the sequential versus simultaneous datasets. The start date for each course is individual-specific and course-specific. We see that for both courses, the visits in the simultaneous data is more "front-loaded" compared to the visits in the sequential data, resulting in a more concave shape in the cumulative distribution.

First, we use the parameter estimates from the sequential data to perform out-of-sample predictions of the simultaneous data. In the sequential data, the lectures and quizzes in

Figure 37: Comparison of observed sequential vs. simultaneous cumulative distribution of visits
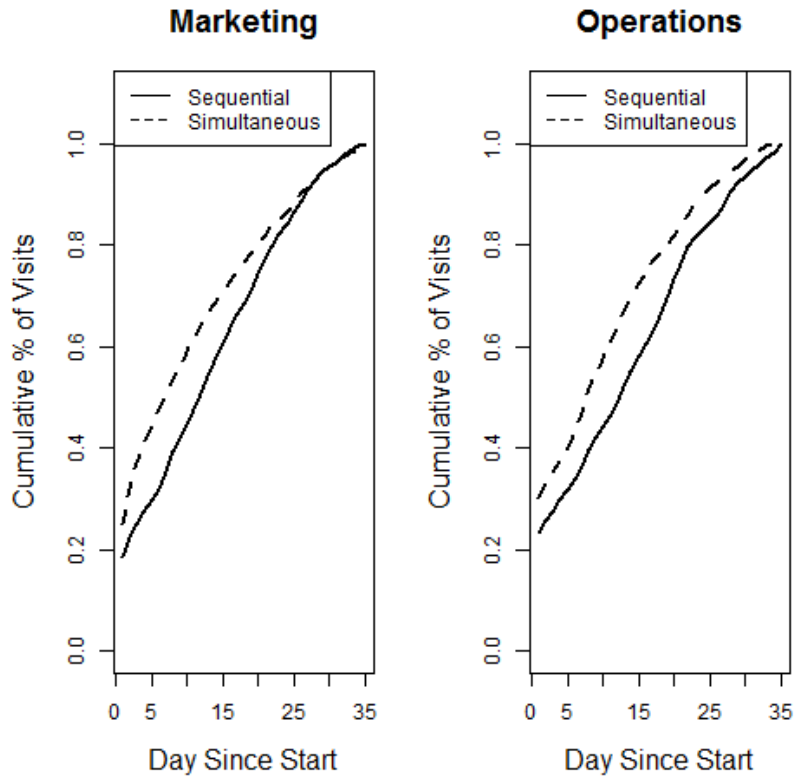
each course were released on a weekly basis, so in the counterfactual scenario, we simulated *simultaneous* release by having all lectures and quizzes available from the moment the individual starts consuming content. The main source of variation comes from the Goal Progress construct, since now instead of having an exogenous goal reset each week, there is one over-arching goal to complete all the content.

Figure 38 compares the simulated data using the parameter estimates from the sequential data. Thus, in this figure the solid sequential curve represents an *in-sample* fit, while the dotted simultaneous curve represents the counterfactual *out-of-sample* fit. We see that using sequential data to predict the simultaneous data is able to capture the contrast between the two datasets, especially for Marketing.

Figure 38: Using sequential data to simulate simultaneous data
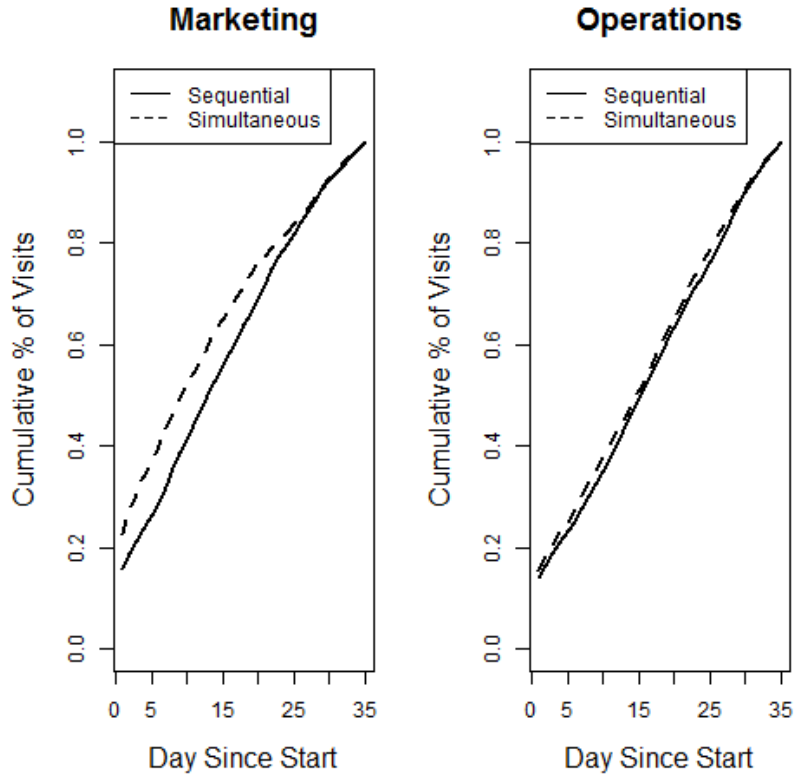
Second, we use the parameter estimates from the simultaneous data to perform out-of-sample predictions of the sequential data. In the simultaneous data, the lectures and quizzes were all available from the moment individuals started consuming content. Therefore, in the counterfactual scenario, we randomly drew a "course start time" (relative to the individual's own consumption start time) from the observed sequential data, and then had lectures and quizzes "released" each week relative to this course start time, for 4 weeks.

Figure 39 compares the simulated data using the parameter estimates from the simultaneous data. In this figure, the dotted simultaneous curve represents the in-sample fit, while the solid sequential curve represents the counterfactual out-of-sample fit. We see that using the simultaneous data to predict the sequential data fails to capture the difference between the two datasets, and moreover overestimates the concavity in consumption across the weeks.

Figure 39: Using simultaneous data to simulate sequential data
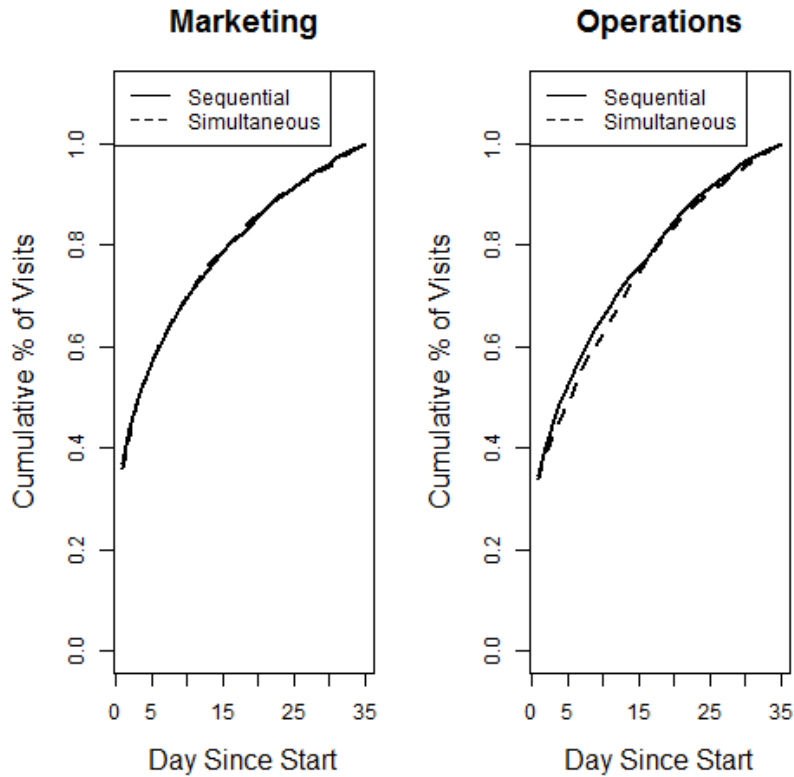
In summary, we conducted counterfactual simulations where the content release policy was changed from weekly sequential to on-demand simultaneous release, and vice versa. We then took advantage of a natural experiment policy change where the Marketing and Operations courses offered on the Coursera platform actually transitioned from sequential to simultaneous release. The predictions of our simulations were able to capture the change in course content visit patterns going from sequential to simultaneous data, but not from simultaneous to sequential data. In future work, we will explore these differences to understand why the predictions are accurate in one direction but not the other.

4.8. Predicting Downstream Behaviors

Binge consumption patterns may be predictive of more downstream behaviors that are of interest to both instructors on the Coursera platform and the firm itself. For example, in

prior work, clumpiness has been shown to be predictive of customer lifetime value (Zhang, Bradlow, and Small 2014), variety-seeking has been used to improve predictions of brand choice (McAlister 1982; Chintagunta 1998), and binging on Hulu has been related to advertising response (Schweidel and Moe 2016).

In our online learning context, a notable characteristic in the consumption patterns of learners in both our sequential and simultaneous datasets was the long runs of course content consumption without breaks, which we referred to as temporal binging, and the low rate of switching between Marketing and Operations, which we referred to as content binging. In the following subsections, we examine whether our temporal and content binging metrics are predictive of both within and cross-course behaviors. In these analyses, we focus on learners from the sequential release dataset.

4.8.1. Within-Course Behaviors

First we look at whether temporal and content binging patterns are predictive of activity within the focal courses, Marketing and Operations. Using simple linear regression analyses, we find that binging in earlier weeks of the course is predictive of binging in later weeks. For example, the coefficient from the regression of temporal binging in week 2 on week 1 is significant ($\beta = 0.44, p < 0.001$), and this relationship is robust across all 5 weeks, and for content binging as well.

We also find that more binging predicts more course completion, in terms of the percentage of the content that has been visited so far (temporal binging: $\beta = 0.03, p < 0.001$, content binging: $\beta = 2.37, p < 0.001$). Interestingly, more temporal binging is actually correlated with lower grades in the course ($\beta = -0.02, p < 0.001$ for average quiz scores in Operations, and $\beta = -0.01, p < 0.05$ for final grades in Marketing), which is consistent with the lay intuition that cramming is not good for test-taking, or indicates that the individuals who cram may also be those who are not good learners and already behind in the course. These results are also consistent with work that finds that distributed (vs. massed) learning leads

to better memory and retention (Childers and Tomasello 2002; Bloom and Shuell 1982; Cepeda et al. 2006).

4.8.2. Cross-Course Behaviors

In 2015 when we observed the sequential dataset, two other courses were being offered in addition to Marketing and Operations as part of the Wharton Online Business Foundations package: Introduction to Accounting and Introduction to Finance. Among the 467 individuals in our original sample who registered and paid for Marketing and Operations, 83.7% had also registered for Accounting, while 81.5% had registered for Finance (and had visited at least one lecture/quiz). This allows us to look at cross-course downstream behaviors.

A question of interest for the Coursera platform, as with many new product introductions, is whether the diffusion and adoption of one product can help predict the adoption of related products (Wind and Mahajan 1997; Van den Bulte 2000). For online learning environments like Coursera, "adoption" can mean engagement in course content *or* payment for different certification levels (similar to basic vs. premium subscription services, for example).

Because we are able to track the same individuals across multiple courses offered by Wharton Online on the Coursera platform, we can look at how the activity in one course predicts engagement in other courses. For example, we can look at whether the number of lecture visits in Operations is predictive of whether or not individuals later registered for Accounting or Finance.

In Figure 40, we separated individuals into quartiles based on their total visits to Operations lectures, and plotted the percentage of individuals in each quartile who later registered for Accounting and Finance. We see that individuals within higher quartiles were more likely to register for Accounting and Finance.

A binary logistic regression also reveals a significant positive effect of the number of visits to Operations lectures on Accounting registration ($\beta = 0.02, z = 4.59, p < 0.001$) and Finance

registration ($\beta = 0.04, z = 5.25, p < 0.001$). Likewise, in future work, we can look further into how engagement in Operations or Marketing predicts both registration *and* payment for other courses offered by Wharton Online.

Figure 40: Operations lecture visits and cross-course registration
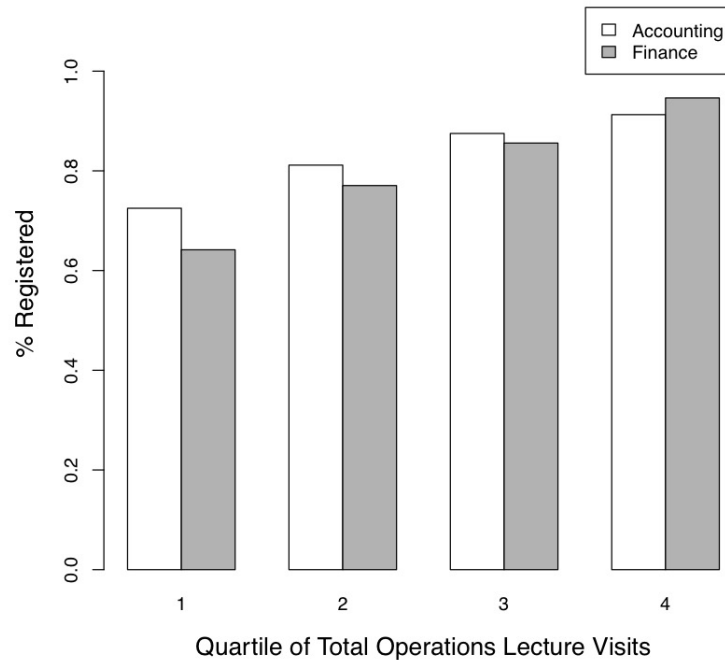
In addition, we found that content binging in Marketing and Operations was predictive of overall consumption, in terms of total number of URL visits, for both Accounting ($\beta = 39.7, p < 0.05$) and Finance ($\beta = 45.0, p < 0.01$). Content binging in Marketing and Operations was also predictive of percent completion in Accounting ($\beta = 0.53, p < 0.05$) and Finance ($\beta = 0.55, p < 0.05$). Interestingly, temporal binging within Marketing and Operations was also able to predict whether individuals ended up paying for the certificate in the Finance course ($\beta = 0.07, p < 0.05$), which suggests that perhaps individuals who are able to consumer content for longer periods of time without breaks are more likely to pursue additional courses.

In future work, we can examine whether or not these results hold for the individuals in our simultaneous dataset from 2016-2017, who had much more flexibility in when they could take each course as well as a broader set of courses available on the platform compared to the individuals in the sequential dataset.

4.9. Discussion

In this chapter, we investigated content consumption within an online learning setting where individuals can watch lecture videos and take quizzes to evaluate their accumulated knowledge. We model individual decisions to watch lectures, take quizzes, and take breaks within a two-stage decision process where they consider both the contemporaneous utility of consumption as they make progress through the course, as well as the utility from knowledge accumulation. Our model is able to capture the patterns in the data that we refer to as temporal and content binging. In addition, we use data from two different content release formats, sequential and simultaneous, and use the model estimation results from each format to make predictions about the alternative format.

One feature of our model is that we treat individual decisions to engage in content as discrete, with a finite set of break length ranges to choose from. One extension of our model would be to further assume that individuals can choose the continuous length of time they engage in a particular lecture, quiz, or break. In the context of online learning, there is reason to believe that breaks of different lengths may have different effects on utilities and the accumulation of knowledge; for example, research has shown that sleep can lead to improvements in recently acquired memories (Ellenbogen et al. 2006).

We can also extend our model to incorporate the choice of visiting new content versus repeating a lecture or a quiz. In the current model, we use Goal Progress as a measure of the utility of consumption. Goal Progress is quantified by the percentage of lectures and quizzes in a course that have been visited at least once. However, another construct that could capture the decreasing or increasing returns to consumption could be "Efficiency"

or the effect of the total number of times individuals have visited lectures and quizzes so for, which would reflect both first visits and revisits to content. Depending on the sign of the parameters, Efficiency could capture satiation or hedonic adaptation (Inman 2001; Nelson and Meyvis 2008; Nelson, Meyvis, and Galak 2009) if individuals experience decreasing returns to consumption, or fluency if individuals experience increasing returns to consumption (Chou and Ting 2003; Whittlesea and Leboe 2000; Greifeneder, Bless, and Pham 2011).

In our model, we made the assumption that individuals were making decisions using a two-stage process, where they first chose the course (vs. a break), and then choice the specific type of content within the course to consume (or a specific break length). However, there are many alternative decision trees that are plausible. For example, we might collapse decisions into a single stage consisting of 7 choices, or extend decisions to take place across three stages where individuals might first choose between consuming anything at all vs. taking a break, and then go on to choose between Marketing/Operations, and then lectures/quizzes within each course. (See Appendix A.11. for detailed examples.)

We also restricted our analyses to individuals who had registered *and* paid for the course, in order to obtain a sample with enough observations to quantify and model binge consumption at the individual level. In future work, we can extend our model to capture how individuals actually make payment decisions. For example, they might register and sample the course content before deciding to pay. We can also model more sparse consumption patterns of individuals who may quit after only a few lectures.

Finally, we focused on a short time horizon of individuals engaged in a single 5-week session simultaneously in two courses. However, individuals who have paid and did not pass the course in a particular session can return for later sessions for additional opportunities to pass the course and obtain an online certificate. In future work, we plan to examine how individuals make "repeat registration" decisions after the initial payment investment, and also how individuals complete portfolios or "bundles" of courses (e.g., all four courses in

115

the Business Foundations bundle) over multiple years, and how their binge consumption patterns relate to the long-term completion of these bundles.

In summary, this chapter provides one way of using behavioral theory constructs to model binge consumption patterns of individuals in an online learning setting, within a framework where there is tension between the utility of consumption and the utility from knowledge accumulation. By shedding light on how individuals make daily decisions to engage in content within specific courses, our work provides implications for how content providers should time content release and make predictions about future course engagement.

CHAPTER 5 : Conclusion

We examine how consumers dynamically update their preferences across three different online settings: online shopping, online gaming, and online learning. Consumers can update their preferences based on information they get from the environment and through experience with the available products. Table 12 gives an overview. In each chapter, we build a mathematical model of consumer decisions that parameterizes specific phenomena from behavioral theory. We test that our model captures the specific patterns observed in the data via in-sample simulations of consumers' sequences of choices. We also test the effects of specific managerial interventions using out-of-sample counterfactual simulations, and empirically verify our predictions using lab experiments or by taking advantage of data following natural experiment policy changes.

Table 12: Overview of chapters

	Online Shopping (Chapter 2)	Online Gaming (Chapter 3)	Online Learning (Chapter 4)
1. Mathematical model of consumer decisions	Sequential sampling model with forward-looking and threshold-updating	Dynamic variety-seeking preferences based on perceptual map distances	Course consumption with goal progress, knowledge accumulation, and forward-looking
2. Test that model captures patterns in data (in-sample)	Search length, sequence of fixations and clicks on AOIs	Sequence of map choices, switching	Temporal and content binging
3. Counterfactual simulation of managerial interventions (out-of-sample)	Preference-based product ordering	New product release	Sequential vs. simultaneous content release
4. Empirical verification of counterfactuals	Experimental manipulation in lab	Natural experiment policy change	Natural experiment policy change

In Chapter 2, we built a sequential sampling model of individual shopper eye fixations and clicks on the landing page of an online store, and used extensions of our model to test for forward-looking and threshold-updating behaviors. Using in-sample simulations from the posterior parameter samples, we were able to capture search length on the landing page, the sequence of fixations and clicks on AOIs, as well as the heterogeneity in search patterns across shoppers. We performed counterfactual simulations to test whether different preference-bass product orderings would change search length and the products that were clicked on. We empirically verified our counterfactual predictions using a follow-up lab experiment with different experimental conditions.

In Chapter 3, we built a model to capture the map choices of players of an online game across rounds of gameplay. We parameterized players' variety-seeking preferences using a coefficient on the perceptual distance between maps that could vary over time and with round outcomes. We used in-sample simulations to predict the sequence of map choices and map switching. We then used out-of-sample simulations to predict whether individuals would adopt new maps upon their release in various expansion packs, and compared this to observed behavior following the firm's real-world release of the new maps.

In Chapter 4, we modeled the course consumption decision of users on an online learning platform. To capture individual decisions about whether to engage in a course or take a break and whether to watch lectures or take quizzes, we used the parameterization of our model to test various theories of goal progress, knowledge accumulation, and boundedly rational forward-looking. We demonstrate that our model is able to improve in-sample simulations of individual sequences of consumption decisions, and in particular the various components of our model improves predictions of temporal and content binging that we find in the observed data. In addition, we test whether our model is able to accurately make out-of-sample counterfactual predictions by using separate datasets from when course content was released sequentially versus simultaneously, and examine whether the model results from one dataset is able to predict the patterns in the other dataset, and vice versa.

In summary, we examine three different online consumer settings that vary widely in temporal resolutions, from eye movements that occur within a second and comprise a 5-minute shopping trip, to online gaming and online video decisions that occur several times a day and may span weeks or even years. We demonstrate that building models that are grounded in behavioral theory can help predict the sequence of choices the consumers make, and also capture how these choices change with different managerial interventions.

A.1. Model simulation and recovery

We simulated data for 67 shoppers for the baseline and full models, with the true means of the parameters shown in Table 13, and the true variance-covariance matrix as a diagonal matrix with 0.1 along the diagonal. Using a hierarchical Bayes estimation procedure, we were able to recover the means of the parameters, with nearly all 95% Credible Intervals overlapping the true means.

Table 13: Recovery of Simulated Data for Baseline and Full Model

Parameter	Baseline			Forward-Looking Fixations			Empirical Threshold-Changing		
	True	Est	95% CI	True	Est	95% CI	True	Est	95% CI
β_1 Preference	1	1.01	[.69, 1.31]	1	1.07	[.75, 1.40]	1	1.26	[.6, 1.80]
β_2 Price	.5	.62	[.40, .83]	.5	.72	[.53, .92]	.5	.56	[.33, .78]
β_3 Out-of-Budget	-1	-1.21	[-1.66, -.80]	-1	-1.36	[-1.79, -.98]	-1	-1.00	[-1.73, -.42]
β_4 Category	3	3.44	[2.98, 3.91]	3	3.78	[3.52, 4.04]	3	3.55	[3.22, 3.92]
γ_1 Scroll	-1	-1.00	[-1.17, -.82]	-1	-1.29	[-1.54, -1.07]	-1	-1.12	[-1.32, -.95]
γ_2 Distance	-1	-1.19	[-1.42, -.95]	-1	-1.49	[-1.76, -1.22]	-1	-1.32	[-1.57, -1.10]
γ_3 Unseen	-1	-.93	[-1.21, -.66]	-1	-1.07	[-1.41, -76]	-1	-.69	[-1.01, -.39]
ρ	7	7.38	[6.89, 7.83]	7	7.82	[7.39, 8.18]	6	5.89	[5.52, 6.29]
λ_1	N/A	N/A	N/A	.05	.08	[.06, .10]	N/A	N/A	N/A
λ_3	N/A	N/A	N/A	N/A	N/A	N/A	-3	-2.89	[-3.79, -2.07]

A.2. Experiment 1 Shopping Instructions

Participants in Experiment 1 were randomly assigned to one of the following goal conditions, with the wording based on literature suggesting that consumers may be motivated by utilitarian, practical considerations or hedonic, pleasure-seeking ones (Holbrook and Hirschman 1982).

Utilitarian goal instructions: "You have decided to go shopping for a new top. Think of this as a shopping trip for practical, every day, casual clothing, something you might wear to work or class. Your goal is to find a top."

Hedonic goal instructions: "You have decided to go shopping for a new top. Think of this as a shopping trip for frivolous, indulgent, fun clothing, something you might wear to a party. Your goal is to find a top, and also to enjoy the shopping experience."

A.3. Grid Search for Preference Value Updating Rate α_s

To determine the updating rate α_s for the preference feature, we used a grid search and estimated the baseline model using five different values of α_s. Table 14 shows that in both Experiments 1 and 2, we find that the model with $\alpha_s = .1$ results in the lowest DIC. Therefore, we use this value for the remaining analyses.

Table 14: DIC values for different updating rates

α_s for Preference	Experiment 1	Experiment 2
.1	8802	17486
.3	8807	17606
.5	8834	17640
.7	8838	17646
.9	8849	17660

A.4. Alternative Distance Variable Specification

As an alternative to the binary Distance variable that we used in the model estimation, which captured "near" versus "far" AOIs, we tried using the Euclidean pixel distance between the center of the AOIs. For example, if the centers of two AOIs were 200 pixels apart vertically on the screen and 400 pixels apart horizontally, then the Distance between them would be $\sqrt{200^2 + 400^2} = 447.21$. The pixel distance between an AOI and itself if 0.

We estimated the baseline model using Euclidean Distance and found that the average simulated number of fixations was 25.93 (10.24), RMSE = 16.46. Figure 41 compares the results for Euclidean and Binary Distances, and we see that the Euclidean distance tends

to more severely under-predict the total number of fixations. One explanation is that since the relationship between the effort and the Euclidean pixel distance is linear, there is a very high probability of re-fixating on the same AOI, which has 0 distance from itself, resulting in less exploration of other products on the page and shorter search length.

Figure 41: Observed vs. simulated fixations for alternative distance variables

Comparison of Distance Variables
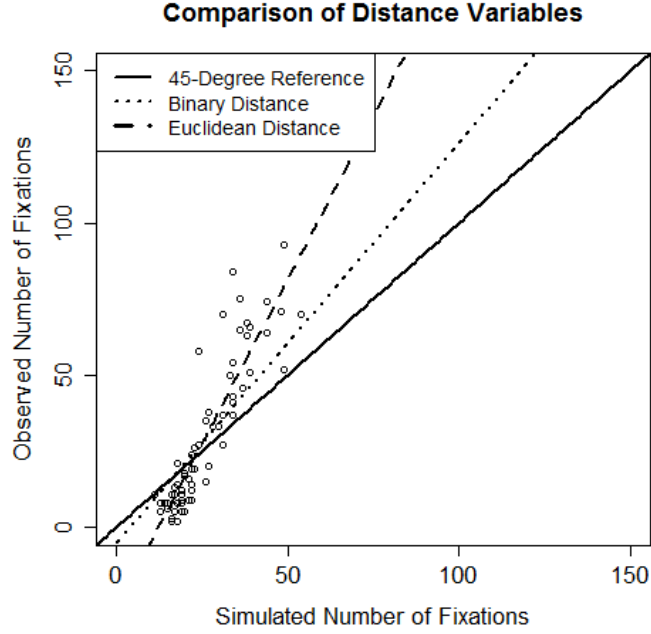
A.5. Calculation of Exchange Rate of Attractiveness and Fixations

$$V = (\beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 x^3) - \alpha(\gamma_0 + \gamma_1 x + \gamma_2 x^2 + \gamma_3 x^3)$$
$$= (\beta_0 - \alpha\gamma_0) + (\beta_1 - \alpha\gamma_1)x + (\beta_2 - \alpha\gamma_2)x^2 + (\beta_3 - \alpha\gamma_3)x^3 \qquad (A.1)$$
$$= \Delta_0 + \Delta_1 x + \Delta_2 x^2 + \Delta_3 x^3$$

To solve for the maximum, we take the derivative of V with respect to x and set to 0:

$$V' = \Delta_1 + 2\Delta_2 x + 3\Delta_3 x^2 = 0 \qquad (A.2)$$

122

One solution to this is $\alpha = \frac{\beta_1}{\gamma_1}$. To ensure this is a maximum and not a minimum, we check that $\beta_3 - \frac{\beta_1}{\gamma_1}\gamma_3 < 0$.

A.6. Estimation Results for Experiment 2 Conditions

Table 15 shows the estimation results for all participants in Experiment 2, as well as separated by ordering condition (random, worst-first, and best-first). We see that across conditions, the mean values of the parameters are very close in terms of sign and order of magnitude. There are some expected differences between conditions. For example, the preference coefficient β_1 is smaller in magnitude and the effort coefficient γ_2 is larger in the best-first condition compared to in other conditions because all the products at the top of the page had high liking ratings, so they were differentiated more by effort rather than attractiveness. Furthermore, the decision threshold intercept ρ is largest in the worst-first condition where shoppers made the most fixations on average, and smaller in the best-first condition where shoppers made the fewest fixations.

Table 15: Parameter estimation results for full model

	All (N=132)			Random (N=41)			Worst-First (N=45)			Best-First (N=46)		
	μ	95% CI	σ^2	μ	95% CI	σ^2	μ	95% CI	σ^2	μ	95% CI	σ^2
β_1 Preference	.59	[.44, .71]	.39	.65	[.42, .87]	.43	.95	[.51, 1.29]	1.41	.27	[.08, .46]	.38
β_2 Price	-.07	[-.17, .03]	.24	-.11	[-.34, .12]	.46	-.08	[-.29, .14]	.45	-.05	[-.27, .17]	.48
β_3 Out-of-Budget	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
β_4 Category	1.78	[1.36, 2.13]	1.40	1.94	[1.24, 2.68]	1.94	1.71	[.78, 2.48]	2.46	1.50	[.57, 2.19]	2.37
γ_1 Distance	-2.98	[-3.36, -2.64]	.85	-3.75	[-4.60, -2.85]	2.36	-3.52	[-4.27, -2.71]	1.52	-2.41	[-3.06, -1.93]	.84
γ_2 Effort	-.66	[-.86, -.47]	.87	-.88	[-1.45, -.40]	2.03	-.37	[-.65, -.10]	.60	-1.15	[-1.97, -.67]	1.72
γ_3 Unseen	-.50	[-.77, -.28]	1.05	-.84	[-1.44, -.32]	1.96	-.23	[-.72, 0.17]	1.24	-.75	[-1.26, -.32]	1.47
ρ	7.06	[6.69, 7.51]	.98	7.42	[6.72, 8.19]	1.10	7.79	[7.06, 8.68]	1.59	6.66	[5.63, 7.70]	1.68
λ_1	.21	[.12, .32]	.005	.12	[.08, .16]	.001	.21	[.05, .42]	.014	.13	[.03, .38]	.003
λ_3	-2.29	[-2.53, -2.06]	.25	-1.20	[-1.96, -.41]	.01	-2.22	[-2.71, -1.86]	.48	-2.67	[-3.30, -2.01]	.65

A.7. List of lecture and quiz names and video run times

Table 16: List of lectures and quizzes in Marketing

Number	Week	Type	Name	Time
1	1	Lecture	Marketing 101: Building Strong Brands Part I	15:10
2	1	Lecture	Marketing 101: Building Strong Brands Part II (4:10)	4:10
3	1	Lecture	Strategic Marketing	11:39
4	1	Lecture	Segmentation and Targeting	12:45
5	1	Lecture	Brand Positioning	12:48
6	1	Lecture	Brand Mantra: The Elevator Speech	9:41
7	1	Lecture	Experiential Branding	13:24
8	1	Quiz	Quiz #1	
9	2	Lecture	From Product-Centric to Customer-Centric Management	15:25
10	2	Lecture	Cracks in the Product-Centric Approach	9:49
11	2	Lecture	Data-Driven Business Models	4:26
12	2	Lecture	Three Cheers for Direct Marketing	3:51
13	2	Lecture	Which Firms Are Customer Centric?	12:11
14	2	Lecture	What is Customer Centricity?	11:28
15	2	Lecture	Living in a Customer-Centric World	14:48
16	2	Lecture	More Reflections on Customer Centricity	3:21
17	2	Lecture	Questions on Customer Centricity	6:00
18	2	Quiz	Quiz #2	
19	3	Lecture	Introduction and Execution	2:09
20	3	Lecture	Friction	4:39
21	3	Lecture	Online/Offline Competition	4:51
22	3	Lecture	The Long Tail Part 1	10:58
23	3	Lecture	The Long Tail Part 2	9:55
24	3	Lecture	Preference Isolation	14:37
25	3	Lecture	Customers and Digital Marketing	9:49
26	3	Lecture	Influence and How Information Spreads	11:02
27	3	Lecture	Pricing Strategies 1: Introduction	11:14
28	3	Lecture	Distribution Strategies 1: Introduction	13:38
29	3	Lecture	Distribution Strategies 2: Channel Design	13:39
30	3	Quiz	Quiz #3	
31	4	Lecture	Brand Messaging & Communication	12:08
31	4	Lecture	Brand Elements: Choosing a Brand Name	19:57
31	4	Lecture	Brand Elements: Color & Taglines	11:41
31	4	Lecture	Brand Elements: Packaging	10:09
31	4	Lecture	Brand Elements: Persuasion	13:59
31	4	Lecture	Repositioning a Brand	18:58
31	4	Quiz	Final Exam	

Table 17: List of lectures and quizzes in Operations

Number	Week	Type	Name	Time
1	1	Lecture	Intro Session 1	7:55
2	1	Lecture	Intro Session 2	7:56
3	2	Lecture	Module 2 Session 1 Video	9:29
4	2	Lecture	Module 2 Session 2 Video	11:24
5	2	Lecture	Module 2 Session 3 Video	15:21
6	2	Lecture	Module 2 Session 4 Video	6:35
7	2	Lecture	Module 2 Session 5 Video	7:25
8	2	Lecture	Module 2 Session 6 Video	10:56
9	2	Lecture	Module 2 Session 7 Video	14:14
10	2	Lecture	Module Review of Process Analysis	26:16
11	2	Quiz	Module 2 - Process Analysis	
12	3	Lecture	Module 3 Session 1 Video	7:56
13	3	Lecture	Module 3 Session 2 Video	9:39
14	3	Lecture	Module 3 Session 3 Video	6:44
15	3	Lecture	Module 3 Session 4 Video	6:22
16	3	Lecture	Module 3 Session 5 Video	12:34
17	3	Lecture	Module 3 Session 6 Video	8:35
18	3	Lecture	Module 3 Session 7 Video	9:11
19	3	Lecture	Module 3 Session 8 Video	10:00
20	3	Lecture	Review of Productivity	19:57
21	4	Lecture	Module 4 Session 1 Video	10:30
22	4	Lecture	Module 4 Session 2 Video	19:40
23	4	Lecture	Module 4 Session 3 Video	12:17
24	4	Lecture	Module 4 Session 4 Video	9:12
25	4	Lecture	Module 4 Session 5 Video	8:26
26	4	Lecture	Module 4 Session 6 Video	7:05
27	4	Lecture	Module 4 Review	19:12
28	4	Quiz	Module 4 - Quality	
29	4	Quiz	Final Exam - Module 2	
30	4	Quiz	Final Exam - Module 3	
31	4	Quiz	Final Exam - Module 4	

A.8. Determining Break Lengths

In order to assess how sensitive our analyses were to different break cutoff lengths, we tried different cutoff times, ranging from 15 minutes to 2 weeks, when constructing each individual's sequence of choices. We looked at two metrics to determine how the cutoff times influence the choice process: the total number of breaks and the average lecture/quiz run (consecutive lectures/quizzes in either course with no breaks in between), which are both measures of temporal binging. Note that the number of switches between Marketing and Operations, which is a measure of content binging, is not affected by the break cutoff times.

Figure 42 plots the total breaks and average lecture/quiz runs, averaged across individuals, for each cutoff time. We see that there is a distinct jump between 1 and 24 hours, with the curves flattening out around 48 hours. This suggests that there may be a large number of breaks that are less than 24 hours long. So we decided on 1 hour as the break cutoff length.

Figure 42: Total breaks and average runs for different cutoff times
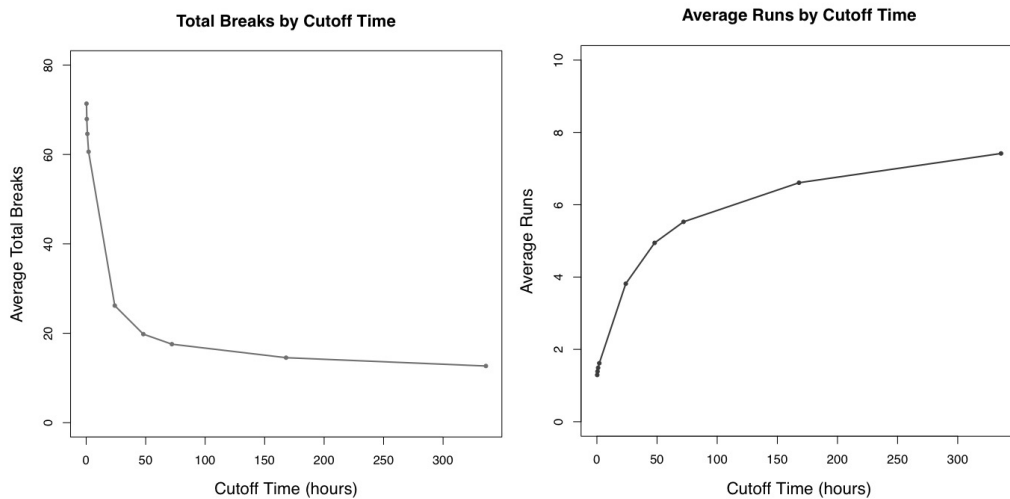
In order to determine the number of break length options for each individual, we looked

at the distribution of break lengths across all individuals in our sample. Figure 43 shows these distributions for all break lengths, as well as a closeup of break lengths that were less than a week long. We see that there is a distinct temporal pattern where the break lengths are concentrated around the minimum break length of 2 hours, as well as 24-hour intervals. Therefore, we chose break lengths that ranged between the peaks: 1-12 hours, 12-36 hours, and 36+ hours (up to 5 weeks).

Figure 43: Distribution of break lengths
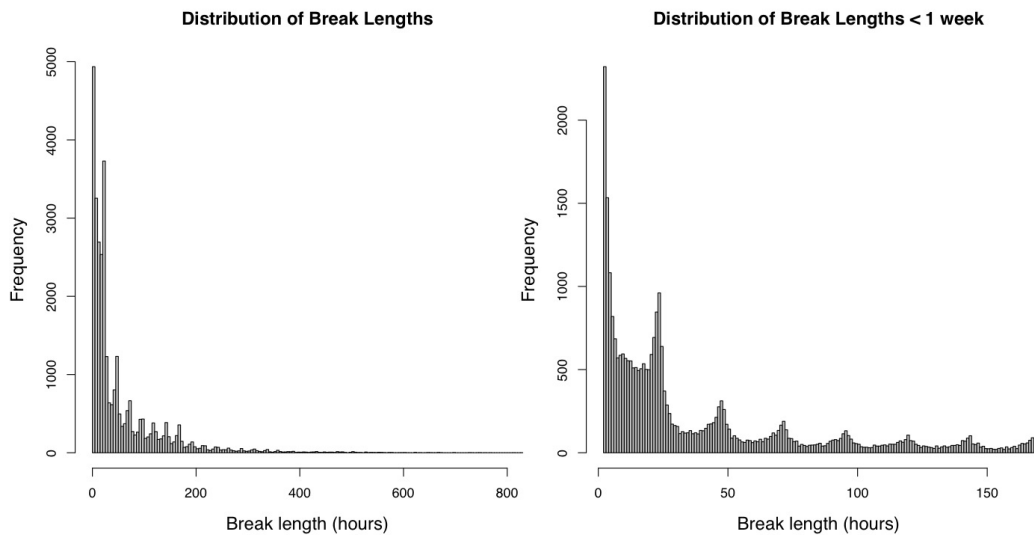
A.9. Parameter Recovery

To determine that our model was identified, we simulated data for 500 individuals, with the parameters for each individual drawn from a multivariate normal distribution with mean and covariance similar to the values of the parameters estimated using the observed data. Table 18 shows the true mean values for the full model with the following constructs: Intercepts, Goal Progress, Completion, Knowledge, and One-Stage Ahead. Note that in this simulation, instead of the Time construct, we used Weekly intercepts. We used an MCMC procedure to estimate the model, which resulted in estimated 95% Credible Intervals that

contained the true means for each parameter.

Table 18: Summary of true and estimated parameters for simulated data

Construct	Parameter	True Mean	Est. Mean	Est. 95% CI	Est. Variance
Intercepts	β_{0M}	1.50	1.40	[1.16, 1.64]	1.48
	β_{0O}	0.50	0.42	[0.17, 0.68]	1.05
	α_{0M}	-2.0	-1.95	[-2.24, -1.73]	1.35
	α_{0O}	-2.0	-2.23	[-2.52, -1.96]	1.62
	δ_1	0.05	0.26	[-0.06, 0.53]	1.64
	δ_2	-0.5	-0.57	[-0.85, -0.34]	1.88
Week	β_{1M}	0.20	0.19	[-0.04, 0.23]	0.54
	β_{2M}	0.10	0.01	[-0.15, 0.19]	0.54
	β_{3M}	0.01	-0.21	[-0.37, -0.03]	0.67
	β_{4M}	-0.50	-0.43	[-1.62, -0.27]	0.70
	β_{1O}	-1.0	-1.31	[-1.59, -0.96]	1.60
	β_{2O}	-0.50	-0.71	[-1.02, -0.39]	1.49
	β_{3O}	-0.50	-0.49	[-0.71, -0.39]	1.52
	β_{4O}	0.10	-0.06	[-0.31, 0.23]	0.99
Goal Progress	β_{5M}	0.20	0.14	[-0.04, 0.30]	0.68
	β_{6M}	-0.20	-0.03	[-0.29, 0.37]	1.74
	β_{5O}	-1.0	-0.71	[-1.11, -0.36]	1.51
	β_{6O}	-1.0	-1.13	[-1.58, -0.59]	3.75
Completion	β_{7M}	-0.50	-0.29	[-0.45, -0.10]	0.54
	β_{8M}	-0.20	-0.16	[-0.09, -0.25]	0.48
	β_{7O}	-0.50	-0.51	[-0.97, -0.18]	1.05
	β_{8O}	0.50	0.32	[0.13, 0.51]	0.93
Knowledge	α_{1M}	2.00	1.50	[1.30, 1.70]	1.10
	α_{2M}	-0.50	-0.48	[-0.45, -0.51]	0.96
	α_{1O}	2.00	1.83	[1.75, 1.86]	1.50
	α_{2O}	0.50	0.23	[0.13, 0.32]	0.50
One-Stage Ahead	γ	0.10	0.30	[0.05, 0.51]	1.15

A.10. Predicting Quiz Scores and Event Lengths

We do not explicitly model the values of quiz scores or the length of calendar time that individuals spend on each lecture, quiz, and break (with the exception of short/medium/long break range choices) as an endogenous outcome in our full model. However, when we simulate data from the posterior draws obtained via the MCMC procedure to determine whether we can recover the temporal and content binging patterns observed in the data, we need to "plug in" a quiz score so that the individual can Bayesianly update her beliefs about her quiz-taking abilities whenever she takes a quiz, and also the calendar times of events so that she can move forward through the 5 weeks of the course.

We regress the quiz scores and break lengths from the observed data on a number of explanatory variables, and then use the estimates from the regressions to "realistically" plug in quiz scores and break lengths when simulating individual choice sequences. This is consistent with the literature that uses rational expectations to generate predictive distributions of endogenous variables (i.e., price; Muth 1961).

Tables 19, 20, and 21 show the results from regressing quiz scores and the event lengths on the following variables: the total number of times the individual had visited Marketing quizzes (T_{MQ}), Marketing lectures (T_{ML}), Operations quizzes (T_{OQ}), and Operations lectures (T_{OL}) at the time the quiz was taken, the total number of times the individual had visited the three break options (T_{B1}, T_{B2}, and T_{B3}), an indicator for whether the individual had visited all available lectures/quizzes at least once (percentages represented by G_{MQ}, G_{ML}, G_{OQ}, and G_{OL}), and the log of the time spent on the quiz.

Table 19: Regression results for quiz scores

	Quiz Score (M)	Quiz Score (O)
Intercept	0.62 (0.05)***	0.14 (0.04)***
Lag $[j-1]$	0.13 (0.03)***	0.05 (0.02)***
Lag $[j-2]$	0.02 (0.03)	0.13 (0.02)***
Lag $[j-3]$	0.13 (0.03)***	0.11 (0.02)***
$\log(T_{MQ})$	-0.04 (0.02)*	-0.09 (0.02)***
$\log(T_{ML})$	0.02 (0.01)**	-0.01 (0.01)
$\log(T_{OQ})$	0.02 (0.01)**	0.09 (0.01)***
$\log(T_{OL})$	-0.00 (0.01)	-0.02 (0.01)**
$\log(T_{B1})$	-0.02 (0.01)*	0.01 (0.01)
$\log(T_{B2})$	0.02 (0.01)*	0.03 (0.01)*
$\log(T_{B3})$	-0.00 (0.02)	0.03 (0.02)
$\mathbb{1}(G_{MQ}=1)$	0.07 (0.02)***	0.01 (0.01)
$\mathbb{1}(G_{ML}=1)$	-0.01 (0.02)	-0.00 (0.02)
$\mathbb{1}(G_{OQ}=1)$	-0.01(0.02)	0.09 (0.02)***
$\mathbb{1}(G_{OL}=1)$	0.06 (0.04)	0.02 (0.03)
$\log(\text{time})$	0.00 (0.01)	0.03 (0.01)***
Adj-R^2	0.12	0.43

*** $p < 0.001$, ** $p < 0.01$, *$p < 0.05$

Table 20: Regression results for quiz/lecture event lengths

	MQ Length	ML Length	OQ Length	OL Length
Intercept	0.10 (0.02)***	0.11 (0.01)***	0.16 (0.03)***	0.05 (0.01)***
Lag $[j-1]$	0.09 (0.03)***	0.24 (0.01)***	0.14 (0.02)***	0.23 (0.01)***
Lag $[j-2]$	0.02 (0.03)	0.13 (0.02)***	0.10 (0.0)***	0.14 (0.01)***
Lag $[j-3]$	-0.04 (0.03)	0.11 (0.01)***	0.02 (0.02)	0.11 (0.01)***
$\log(T_{MQ})$	0.01 (0.01)*	-0.01 (0.00)	-0.02 (0.01)*	0.01 (0.001)
$\log(T_{ML})$	-0.00 (0.00)	-0.02 (0.00)***	0.00 (0.01)	-0.02 (0.00)***
$\log(T_{OQ})$	0.00 (0.00)	0.00 (0.00	-0.04 (0.01)***	-0.01 (0.00)**
$\log(T_{OL})$	-0.01 (0.00)	-0.01 (0.00)***	-0.002 (0.01)	-0.01 (0.00)**
$\log(T_{B1})$	0.02 (0.01)*	0.04 (0.00)***	0.02 (0.01)**	0.04 (0.00)***
$\log(T_{B2})$	0.01 (0.00)	0.00 (0.00)	0.03 (0.01)***	000 (0.00)
$\log(T_{B3})$	0.01 (0.01)	0.01 (0.00)	0.01 (0.02)	0.02 (0.01)***
$\mathbb{F}(G_{MQ}=1)$	-0.00 (0.01)	-0.04 (0.01)***	0.03 (0.01)*	0.02 (0.01)**
$\mathbb{F}(G_{ML}=1)$	0.01 (0.01)	-0.04 (0.00)***	-0.01 (0.01)	-0.02 (0.01)***
$\mathbb{F}(G_{OQ}=1)$	-0.05 (0.02)**	-0.00 (0.00)	-0.00 (0.01)	-0.01 (0.01)
$\mathbb{F}(G_{OL}=1)$	-0.00 (0.02)	0.01 (0.01)	0.00 (0.02)	-0.01 (0.01)
Adj-R^2	0.02	0.15	0.08	0.13

*** $p < 0.001$, ** $p < 0.01$, *$p < 0.05$

Table 21: Regression results for break event lengths

|	B1 Length	B2 Length	B3 Length
Intercept	0.10 (0.02)***	0.11 (0.01)***	0.16 (0.03)***
Lag $[j-1]$	0.09 (0.03)***	0.24 (0.01)***	0.14 (0.02)***
Lag $[j-2]$	0.02 (0.03)	0.13 (0.02)***	0.10 (0.0)***
Lag $[j-3]$	-0.04 (0.03)	0.11 (0.01)***	0.02 (0.02)
$\log(T_{MQ})$	0.01 (0.01)*	-0.01 (0.00)	-0.02 (0.01)*
$\log(T_{ML})$	-0.00 (0.00)	-0.02 (0.00)***	0.00 (0.01)
$\log(T_{OQ})$	0.00 (0.00)	0.00 (0.00	-0.04 (0.01)***
$\log(T_{OL})$	-0.01 (0.00)	-0.01 (0.00)***	-0.002 (0.01)
$\log(T_{B1})$	0.02 (0.01)*	0.04 (0.00)***	0.02 (0.01)**
$\log(T_{B2})$	0.01 (0.00)	0.00 (0.00)	0.03 (0.01)***
$\log(T_{B3})$	0.01 (0.01)	0.01 (0.00)	0.01 (0.02)
$\mathbb{F}(G_{MQ}=1)$	-0.00 (0.01)	-0.04 (0.01)***	0.03 (0.01)*
$\mathbb{F}(G_{ML}=1)$	0.01 (0.01)	-0.04 (0.00)***	-0.01 (0.01)
$\mathbb{F}(G_{OQ}=1)$	-0.05 (0.02)**	-0.00 (0.00)	-0.00 (0.01)
$\mathbb{F}(G_{OL}=1)$	-0.00 (0.02)	0.01 (0.01)	0.00 (0.02)
Adj-R^2	0.02	0.15	0.08

*** $p < 0.001$, ** $p < 0.01$, *$p < 0.05$

A.11. Alternative Decision Trees

In our model, we assume that individuals are making decisions in 2 stages. However, alternative decision "trees" are plausible, as illustrated in Figure 44. Different decision trees have different psychological interpretations, and may result in better or worse fit (i.e., Neslin et al. 2014 review the differences between brand and channel choice).

For example, we might collapse stage 1 of the 2-stage model and assume that individuals are actually making a 1-stage decision between all 7 choice options. The main difference between the 1 and 2-stage models would be the utilities of quizzes/lectures and short/medium/long break lengths.

Figure 44: Decision trees for models with 1 vs. 2 vs. 3 stages.
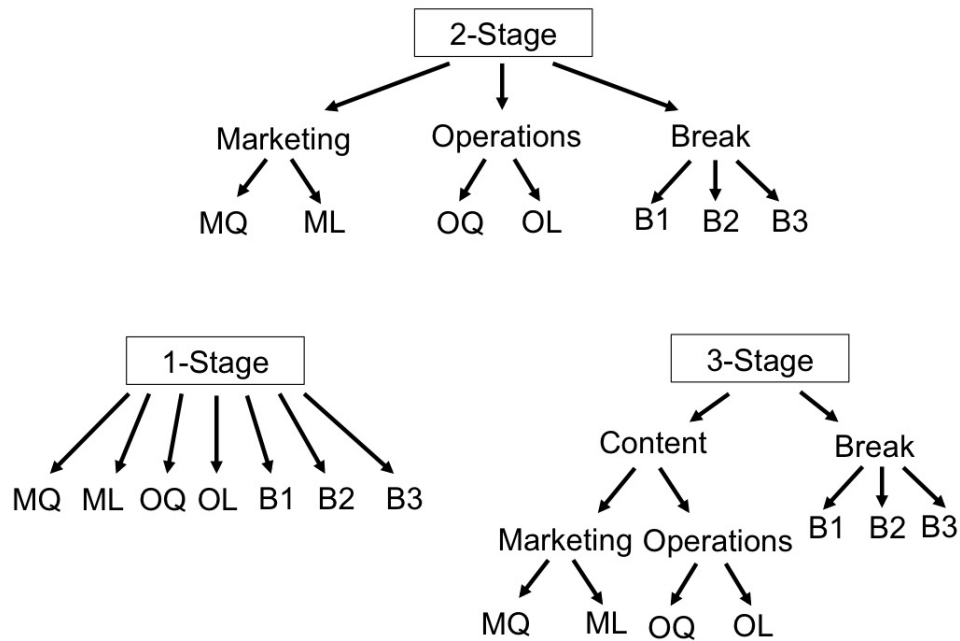
We might also assume that individuals are making decisions within a 3-stage process by first making a choice between general content vs. break, and then deciding on the specific

course, Marketing vs. Operations. Note that in the 2-stage model, temporal and content binging are both captured in stage 1. In the 3-stage model, temporal binging occurs in stage 1, as individuals decide between content and break, while content binging occurs in stage 2 of the content branch, as individuals decide between the two courses.

We note that by adding "stage-ahead" parameters, which is analogous to the inclusive value terms within a nested logit that induces correlation between options within nests (i.e., stage 2 of the 2-stage model, stage 2 and 3 of the 3-stage model), we can further assess the validity and differences between these alternative decision trees. For example, the 1-stage model is equivalent to a 2-stage model in which individuals have a discount factor of $\gamma = 1$ on the utilities in stage 2 when incorporating these utilities into their decision in stage 1.

BIBLIOGRAPHY

K. Adam. Learning while searching for the best alternative. *Journal of Economic Theory*, 101(1):252–280, 2001.

J. W. Alba and J. W. Hutchinson. Dimensions of consumer expertise. *Journal of Consumer Research*, 13(4):411–454, 1987.

F. M. Bass. A new product growth for model consumer durables. *Management Science*, 15 (5):215–227, 1969.

G. S. Becker and K. M. Murphy. A theory of rational addiction. *Journal of Political Economy*, 96(4):675–700, 1988.

M. Ben-Akiva and B. Boccara. Discrete choice models with latent choice sets. *International Journal of Research in Marketing*, 12(1):9–24, 1995.

J. Berger, A. T. Sorensen, and S. J. Rasmussen. Positive effects of negative publicity: When negative reviews increase sales. *Marketing Science*, 29(5):815–827, 2010.

K. C. Berridge. From prediction error to incentive salience: mesolimbic computation of reward motivation. *European Journal of Neuroscience*, 35(7):1124–1143, 2012.

J.-M. Blin and J. A. Dodson. The relationship between attributes, brand preference, and choice: A stochastic view. *Management Science*, 26(6):606–619, 1980.

K. C. Bloom and T. J. Shuell. Effects of massed and distributed practice on the learning and retention of second-language vocabulary. *The Journal of Educational Research*, 74 (4):245–248, 1981.

A. Bonezzi, C. M. Brendl, and M. De Angelis. Stuck in the middle: The psychophysics of goal pursuit. *Psychological Science*, 22(5):607–612, 2011.

F. Branco, M. Sun, and J. M. Villas-Boas. Optimal search for product information. *Management Science*, 58(11):2037–2056, 2012.

B. J. Bronnenberg and C. F. Mela. Market roll-out and retailer adoption for new brands. *Marketing Science*, 23(4):500–518, 2004.

B. J. Bronnenberg, J. B. Kim, and C. F. Mela. Zooming in on choice: How do consumers search for cameras online? *Marketing Science*, 35(5):693–712, 2016.

M. Brown, C. J. Flinn, and A. Schotter. Real-time search in the laboratory and the market. *American Economic Review*, 101(2):948–74, 2011.

R. E. Bucklin and C. Sismeiro. A model of web site browsing behavior estimated on clickstream data. *Journal of Marketing Research*, 40(3):249–267, 2003.

J. R. Busemeyer and J. T. Townsend. Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432459, 1993.

C. Camerer and T. Hua Ho. Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874, 1999.

C. F. Camerer, T.-H. Ho, and J.-K. Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.

N. J. Cepeda, H. Pashler, E. Vul, J. T. Wixted, and D. Rohrer. Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychological Bulletin*, 132(3): 354380, 2006.

P. Chandon, J. W. Hutchinson, E. T. Bradlow, S. Young, M. Wedel, and R. Pieters. *Visual Marketing: From Attention to Action*. Lawrence Erlbaum Assoc., Mahwah, NJ, 2007.

P. Chandon, J. W. Hutchinson, E. T. Bradlow, and S. H. Young. Does in-store marketing work? effects of the number and position of shelf facings on brand attention and evaluation at the point of purchase. *Journal of Marketing*, 73(6):1–17, 2009.

D. Chandrasekaran and G. J. Tellis. Global takeoff of new products: Culture, wealth, or vanishing differences? *Marketing Science*, 27(5):844–860, 2008.

Y. Chen and S. Yao. Sequential search with refinement: Model and application with click-stream data. *Management Science*, 63(12):4345–4365, 2016.

S. E. Chick and P. Frazier. Sequential sampling with economics of selection procedures. *Management Science*, 58(3):550–569, 2012.

J. B. Childers and M. Tomasello. Two-year-olds learn novel nouns, verbs, and conventional actions from massed or distributed exposures. *Developmental Psychology*, 38(6):967978, 2002.

P. K. Chintagunta. Inertia and variety seeking in a model of brand-purchase timing. *Marketing Science*, 17(3):253–270, 1998.

P. K. Chintagunta, R. Jiang, and G. Z. Jin. Information, learning, and drug diffusion: The case of cox-2 inhibitors. *Quarterly Marketing and Economics*, 7(4):399–443, 2009.

T.-J. Chou and C.-C. Ting. The role of flow experience in cyber-game addiction. *CyberPsychology & Behavior*, 6(6):663–675, 2003.

M. Corbetta and G. L. Shulman. Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3):201, 2002.

B. De los Santos and S. Koulayev. Optimizing click-through in online rankings with endogenous search refinement. *Marketing Science*, 36(4):542–564, 2017.

B. G. Dellaert and G. Häubl. Searching in choice mode: consumer decision processes in product search with recommendations. *Journal of Marketing Research*, 49(2):277–288, 2012.

J. M. Ellenbogen, J. C. Hulbert, R. Stickgold, D. F. Dinges, and S. L. Thompson-Schill. Interfering with theories of sleep and memory: sleep, declarative memory, and associative interference. *Current Biology*, 16(13):1290–1294, 2006.

T. Erdem. A dynamic analysis of market structure based on panel data. *Marketing Science*, 15(4):359–378, 1996.

T. Erdem and M. P. Keane. Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1):1–20, 1996.

P. S. Fader, B. G. Hardie, and K. L. Lee. Rfm and clv: Using iso-value curves for customer base analysis. *Journal of Marketing Research*, 42(4):415–430, 2005.

X. Gabaix and D. Laibson. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics*, 121(2):505–540, 2006.

A. Gelman and D. B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, 1992.

A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. *Bayesian Data Analysis*. CRC press Boca Raton, FL, 2014.

K. Gielens and J.-B. E. Steenkamp. Drivers of consumer acceptance of new packaged goods: An investigation across products and countries. *International Journal of Research in Marketing*, 24(2):97–111, 2007.

G. Gigerenzer and D. G. Goldstein. Reasoning the fast and frugal way: models of bounded rationality. *Psychological review*, 103(4):650, 1996.

T. J. Gilbride and G. M. Allenby. A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science*, 23(3):391–406, 2004.

J. C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 148–177, 1979.

M. Givon. Variety seeking through brand switching. *Marketing Science*, 3(1):1–22, 1984.

S. Gluth, J. Rieskamp, and C. Büchel. Neural evidence for adaptive strategy selection in value-based decision-making. *Cerebral Cortex*, 24(8):2009–2021, 2013.

A. Gopalakrishnan, R. Iyengar, and R. J. Meyer. Consumer dynamic usage allocation and learning under multipart tariffs. *Marketing Science*, 34(1):116–133, 2014.

B. R. Gordon and B. Sun. A dynamic model of rational addiction: Evaluating cigarette taxes. *Marketing Science*, 34(3):452–470, 2015.

R. Greifeneder, H. Bless, and M. T. Pham. When do people rely on affective and cognitive feelings in judgment? a review. *Personality and Social Psychology Review*, 15(2):107–141, 2011.

P. M. Guadagni and J. D. Little. A logit model of brand choice calibrated on scanner data. *Marketing Science*, 2(3):203–238, 1983.

G. Häubl, B. G. Dellaert, and B. Donkers. Tunnel vision: local behavioral influences on consumer decisions in product search. *Marketing Science*, 29(3):438–455, 2010.

J. R. Hauser, G. L. Urban, G. Liberali, and M. Braun. Website morphing. *Marketing Science*, 28(2):202–223, 2009.

J. D. Hey. Search for rules for search. *Journal of Economic Behavior & Organization*, 3(1): 65–81, 1982.

J. D. Hey. Still searching. *Journal of Economic Behavior & Organization*, 8(1):137–144, 1987.

M. B. Holbrook and E. C. Hirschman. The experiential aspects of consumption: Consumer fantasies, feelings, and fun. *Journal of Consumer Research*, 9(2):132–140, 1982.

E. Honka and P. Chintagunta. Simultaneous or sequential? search strategies in the us auto insurance industry. *Marketing Science*, 36(1):21–42, 2016.

Y. Huang and J. W. Hutchinson. Counting every thought: Implicit measures of cognitive responses to advertising. *Journal of Consumer Research*, 35(1):98–118, 2008.

Y. Huang and J. W. Hutchinson. The roles of planning, learning, and mental models in repeated dynamic decision making. *Organizational Behavior and Human Decision Processes*, 122(2):163–176, 2013.

S. K. Hui, E. T. Bradlow, and P. S. Fader. Testing behavioral hypotheses using an integrated model of grocery store shopping path and purchase behavior. *Journal of consumer research*, 36(3):478–493, 2009a.

S. K. Hui, P. S. Fader, and E. T. Bradlow. Path data in marketing: An integrative framework and prospectus for model building. *Marketing Science*, 28(2):320–335, 2009b.

S. K. Hui, J. J. Inman, Y. Huang, and J. Suher. The effect of in-store travel distance on unplanned spending: Applications to mobile promotion strategies. *Journal of Marketing*, 77(2):1–16, 2013.

S. K. Hui, T. Meyvis, and H. Assael. Analyzing moment-to-moment data using a bayesian

functional linear model: Application to tv show pilot testing. *Marketing Science*, 33(2): 222–240, 2014.

J. W. Hutchinson. Discrete attribute models of brand switching. *Marketing Science*, 5(4): 350–371, 1986.

J. W. Hutchinson. Netscal: A network scaling algorithm for nonsymmetric proximity data. *Psychometrika*, 54(1):25–51, 1989.

J. J. Inman. The role of sensory-specific satiety in attribute-level variety seeking. *Journal of Consumer Research*, 28(1):105–120, 2001.

R. Iyengar, C. Van den Bulte, and T. W. Valente. Opinion leadership and social contagion in new product diffusion. *Marketing Science*, 30(2):195–212, 2011.

J. Jacoby and D. B. Kyner. Brand loyalty vs. repeat purchasing behavior. *Journal of Marketing Research*, pages 1–9, 1973.

A. P. Jeuland. Brand choice inertia as one aspect of the notion of brand loyalty. *Management Science*, 25(7):671–682, 1979.

B. E. Kahn. Consumer variety-seeking among goods and services: An integrative review. *Journal of Retailing and Consumer Services*, 2(3):139–148, 1995.

B. E. Kahn and A. M. Isen. The influence of positive affect on variety seeking among safe, enjoyable products. *Journal of Consumer Research*, 20(2):257–270, 1993.

B. E. Kahn and J. S. Raju. Effects of price promotions on variety-seeking and reinforcement behavior. *Marketing Science*, 10(4):316–337, 1991.

B. E. Kahn, M. U. Kalwani, and D. G. Morrison. Measuring variety-seeking and reinforcement behaviors using panel data. *Journal of Marketing Research*, pages 89–100, 1986.

D. Kahneman. A perspective on judgment and choice: mapping bounded rationality. *American Psychologist*, 58(9):697, 2003.

H. Katahira. Perceptual mapping using ordered logit analysis. *Marketing Science*, 9(1): 1–17, 1990.

D. Keltner, P. C. Ellsworth, and K. Edwards. Beyond simple pessimism: effects of sadness and anger on social perception. *Journal of Personality and Social Psychology*, 64(5):740, 1993.

J. B. Kim, P. Albuquerque, and B. J. Bronnenberg. Online demand under limited consumer search. *Marketing Science*, 29(6):1001–1023, 2010.

R. Kivetz, O. Urminsky, and Y. Zheng. The goal-gradient hypothesis resurrected: Purchase

acceleration, illusionary goal progress, and customer retention. *Journal of Marketing Research*, 43(1):39–58, 2006.

S. Koulayev. Search for differentiated products: identification and estimation. *The RAND Journal of Economics*, 45(3):553–575, 2014.

I. Krajbich, C. Armel, and A. Rangel. Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10):12921298, 2010.

V. Kumar, R. Venkatesan, T. Bohling, and D. Beckmann. Practice prize reportthe power of clv: Managing customer lifetime value at ibm. *Marketing Science*, 27(4):585–599, 2008.

J. S. Larson, E. T. Bradlow, and P. S. Fader. An exploratory look at supermarket shopping paths. *International Journal of Research in Marketing*, 22(4):395–414, 2005.

R. S. Lazarus. Progress on a cognitive-motivational-relational theory of emotion. *American psychologist*, 46(8):819–834, 1991.

J. S. Lerner, D. A. Small, and G. Loewenstein. Heart strings and purse strings: Carryover effects of emotions on economic decisions. *Psychological Science*, 15(5):337–341, 2004.

J. Liechty, R. Pieters, and M. Wedel. Global and local covert visual attention: Evidence from a bayesian hidden markov model. *Psychometrika*, 68(4):519–541, 2003.

S. A. Lippman and K. F. McCardle. Uncertain search: a model of search among technologies of uncertain values. *Management Science*, 37(11):1474–1490, 1991.

J. G. March. Exploration and exploitation in organizational learning. *Organization Science*, 2(1):71–87, 1991.

L. McAlister. A dynamic attribute satiation model of variety-seeking behavior. *Journal of Consumer Research*, 9(2):141–150, 1982.

W. W. Moe and P. S. Fader. Dynamic conversion behavior at e-commerce sites. *Management Science*, 50(3):326–335, 2004.

A. L. Montgomery, S. Li, K. Srinivasan, and J. C. Liechty. Modeling online browsing and path analysis using clickstream data. *Marketing Science*, 23(4):579–595, 2004.

P. Morgan and R. Manning. Optimal search. *Econometrica: Journal of the Econometric Society*, pages 923–944, 1985.

J. F. Muth. Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, pages 315–335, 1961.

L. D. Nelson and T. Meyvis. Interrupted consumption: Disrupting adaptation to hedonic experiences. *Journal of Marketing Research*, 45(6):654–664, 2008.

L. D. Nelson, T. Meyvis, and J. Galak. Enhancing the television-viewing experience through commercial interruptions. *Journal of Consumer Research*, 36(2):160–172, 2009.

S. A. Neslin, K. Jerath, A. Bodapati, E. T. Bradlow, J. Deighton, S. Gensler, L. Lee, E. Montaguti, R. Telang, R. Venkatesan, et al. The interrelationships between brand and channel choice. *Marketing Letters*, 25(3):319–330, 2014.

T. Otter, J. Johnson, J. Rieskamp, G. M. Allenby, J. D. Brazell, A. Diederich, J. W. Hutchinson, S. MacEachern, S. Ruan, and J. Townsend. Sequential sampling models of choice: some recent advances. *Marketing letters*, 19(3-4):255, 2008.

M. M. Pegler. *Visual merchandising and display.* Fairchild Publications New York, 2006.

M. Poor, A. Duhachek, and S. Krishnan. The moderating role of emotional differentiation on satiation. *Journal of Consumer Psychology*, 22(4):507–519, 2012.

R. Ratcliff and P. L. Smith. A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2):333, 2004.

R. Ratcliff and F. Tuerlinckx. Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9(3):438–481, 2002.

K. Rayner. Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, 124(3):372, 1998.

E. Reutskaja, R. Nagel, C. F. Camerer, and A. Rangel. Search dynamics in consumer choice under time pressure: An eye-tracking study. *American Economic Review*, 101(2): 900–926, 2011.

H. A. Roehm Jr and M. L. Roehm. Revisiting the effect of positive mood on variety seeking. *Journal of Consumer Research*, 32(2):330–336, 2005.

R. Roy, P. K. Chintagunta, and S. Haldar. A framework for investigating habits,the hand of the past, and heterogeneity in dynamic brand choice. *Marketing Science*, 15(3):280–299, 1996.

T. Satomura, M. Wedel, and R. Pieters. Copy alert: A method and metric to detect visual copycat brands. *Journal of Marketing Research*, 51(1):1–13, 2014.

W. Schultz. Neuronal reward and decision signals: from theories to data. *Physiological Reviews*, 95(3):853–951, 2015.

D. A. Schweidel and W. W. Moe. Binge watching and advertising. *Journal of Marketing*, 80(5):1–19, 2016.

P. Seetharaman. Modeling multiple sources of state dependence in random utility models: A distributed lag approach. *Marketing Science*, 23(2):263–271, 2004.

P. Seetharaman and P. Chintagunta. A model of inertia and variety-seeking with marketing variables. *International Journal of Research in Marketing*, 15(1):1–17, 1998.

G. Shaffer and Z. J. Zhang. Pay to switch or pay to stay: Preference-based price discrimination in markets with switching costs. *Journal of Economics & Management Strategy*, 9(3):397–424, 2000.

S. W. Shi, M. Wedel, and F. Pieters. Information acquisition during online decision making: A model-based exploration using eye-tracking data. *Management Science*, 59(5):1009–1026, 2013.

H. A. Simon. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1):99–118, 1955.

K. Sridhar, R. Bezawada, and M. Trivedi. Investigating the drivers of consumer cross-category learning for new products using multiple data sets. *Marketing Science*, 31(4): 668–688, 2012.

P. Stüttgen, P. Boatwright, and R. T. Monroe. A satisficing choice model. *Marketing Science*, 31(6):878–899, 2012.

R. B. Towal, M. Mormann, and C. Koch. Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proceedings of the National Academy of Sciences*, 110(40):E3858–E3867, 2013.

J. T. Townsend and F. G. Ashby. *Stochastic modeling of elementary psychological processes*. CUP Archive, 1983.

M. Trivedi, F. M. Bass, and R. C. Rao. A model of stochastic variety-seeking. *Marketing Science*, 13(3):274–297, 1994.

R. M. Ursu, Q. Wang, and P. K. Chintagunta. Search duration. *Working Paper*, 2017.

C. Van den Bulte. New product diffusion acceleration: Measurement and analysis. *Marketing Science*, 19(4):366–380, 2000.

T. Van Zandt, H. Colonius, and R. W. Proctor. A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review*, 7(2):208–256, 2000.

M. Wadhwa, B. Shiv, and S. M. Nowlis. A bite to whet the reward appetite: The influence of sampling on reward-seeking behaviors. *Journal of Marketing Research*, 45(4):403–413, 2008.

J. P. Walsh. Managerial and organizational cognition: Notes from a trip down memory lane. *Organization science*, 6(3):280–321, 1995.

M. Wedel, R. Pieters, and J. Liechty. Attention switching during scene perception: how

goals influence the time course of eye movements across advertisements. *Journal of Experimental Psychology: Applied*, 14(2):129, 2008.

M. L. Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.

B. W. Whittlesea and J. P. Leboe. The heuristic basis of remembering and classification: Fluency, generation, and resemblance. *Journal of Experimental Psychology: General*, 129 (1):84, 2000.

J. Wind and V. Mahajan. Issues and opportunities in new product development: An introduction to the special issue. *Journal of Marketing Research*, 34(1):1–12, 1997.

L. Yang, O. Toubia, and M. G. De Jong. A bounded rationality model of information search and choice in preference measurement. *Journal of Marketing Research*, 52(2):166–183, 2015.

G. Zauberman and J. G. Lynch Jr. Resource slack and propensity to discount delayed investments of time versus money. *Journal of Experimental Psychology: General*, 134(1): 23, 2005.

Y. Zhang, E. T. Bradlow, and D. S. Small. Predicting customer value using clumpiness: From rfm to rfmc. *Marketing Science*, 34(2):195–208, 2014.