

Binge Consumption of Online Content

Tong (Joy) Lu
Doctoral Candidate
Wharton Marketing Department
University of Pennsylvania
tonglu@wharton.upenn.edu

Eric T. Bradlow
K.P. Chao Professor of Marketing
Wharton Marketing Department
University of Pennsylvania
ebradlow@wharton.upenn.edu

J. Wesley Hutchinson
Stephen J. Heyman Professor of Marketing
Wharton Marketing Department
University of Pennsylvania
jwhutch@wharton.upenn.edu

June 2017

Binge Consumption of Online Content

Abstract

Binge consumption of online content has emerged as a trending behavioral phenomenon among customers of online streaming services. Online content can range from pure entertainment to pure education, with various content providers spanning the spectrum. Here, we specifically focus on bingeing within an online education setting, using a clickstream dataset from Coursera (one of the world’s most popular online education platforms) in which we observe individual-level lecture and quiz consumption patterns across multiple courses. We extend the literature by distinguishing between “temporal bingeing,” where individuals consume multiple pieces of content in a single sitting, and “content bingeing,” where individuals consume content from the same course in succession.

First, we endogenize bingeing as an outcome of utility maximization using a mathematical model of the decision processes of individual consumers. The key feature of our model is the tension between the contemporaneous utility of consumption (i.e., enjoyment and motivation to engage in the learning experience) and the long-run accumulation of knowledge (i.e., memory and usage of acquired facts). Within an online learning setting, individuals have specific attainable goals (i.e., passing the course, learning the content) and we have a concrete measure of their knowledge (i.e., quiz scores). Our model is able to endogenously capture bingeing behavior, including which course people choose to consume, whether they consume a lecture or a quiz, and whether they take breaks of different lengths. Second, we compare the performance of our model under the assumption that individuals are either myopic or forward-looking in a boundedly-rational way. Third, we conduct counterfactual simulations to determine how the timing of content release affects binge consumption, cross-course consumption, and knowledge accumulation. We test these predictions using data obtained after a “natural experiment” policy change when the Coursera platform went from sequential to simultaneous course content release.

Keywords: *binge consumption, digital media consumption, bounded rationality*

1 Introduction

Binge consumption has become increasingly prevalent as consumers shift from traditional media outlets, where content is released on a fixed schedule, to online streaming services that offer content to consumers on-demand. Various content providers and marketing research firms have come up with different definitions of binge consumption of media, usually focused on the number of episodes of a TV show that viewers watch in a single sitting or how quickly viewers progress through episodes. For example, the Digital Democracy survey defines binge-watching as watching three or more episodes of a TV series in one sitting (Deloitte 2015), while Netflix considers whether viewers finish a 13-episode season within a week (Jurgensen 2013).

In our work, we distinguish between two types of bingeing: “temporal bingeing” and “content bingeing.” Temporal and content bingeing are both related to patterns of consumer behavior that have been modeled in other marketing settings. To illustrate, suppose an individual is taking two online courses that each consist of a series of lecture videos. Temporal bingeing would mean that the individual consumes multiple lectures in a single sitting rather than spreading them out over time. Temporal bingeing 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 bingeing 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 over time. The distinction between content bingeing 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).

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 or 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 their reduced-form approach, we endogenize bingeing as an outcome of a consumer’s decision making process within a model that assumes consumers balance the utility of consumption against the desire to accumulate the knowledge needed to pass online quizzes, potentially to obtain an online certificate.

1.1 Research Contribution

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 could be categorized as educational. Here, we specifically focus on bingeing within an online education setting, using a clickstream dataset from Coursera, in which we observe individual-level lecture video and quiz consumption across multiple courses, as well as outcome variables such as quiz scores, which learners accumulate to obtain a certificate for passing a given course.

Binge-watching 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), and has also be characterized as addiction in extreme cases (Becker and Murphy 1988; Gordon and Sun 2015). In contrast, the key feature of our model is the “mathematical tension” between the contemporaneous utility of consumption and the long-run accumulation of knowledge.

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). Consumers may also exhibit patterns of decreasing motivation because they initially

overestimate the resources or “slack” they will have to invest in the course in the future (Zauberman and Lynch 2005). Our model flexibly allows for both increasing and decreasing motivation to consume course content over time, as well as non-monotonic patterns such as the “stuck in the middle” effect where the reference point 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). In addition, our model allows for changes in consumption patterns that may occur as individuals continue to consume once they have “completed” the course content, after which the desire to consume content may drop.

On the other hand, in the context of online education, an important goal of consumers is to gain expertise or accumulate knowledge (Alba and Hutchinson 1987) to pass the course, which can be evaluated through quiz scores. Thus, we allow the utility of taking quizzes to depend on individuals’ beliefs about their quiz-taking abilities, which are Bayesianly updated when they get feedback through their quiz scores. We also allow the utility of quizzes to depend on the knowledge accumulated from watching lectures; however, for some individuals watching more lectures may actually be an indication of a *lack* of knowledge.

Importantly, binge behaviors (both temporal binging and content binging) are not explicit parameters of our model (e.g., we do not have a parameter for state-dependence), but are instead captured endogenously within a mathematical model of consumer decision processes. This is similar to the approach of Erdem and Keane (1996) who endogenize variety-seeking as a result of reducing uncertainty about brands. Rather than assuming consumers are infinitely forward-looking and solving an infinite time-horizon dynamic programming problem, we consider consumers to be boundedly rational in taking into account future utilities. We compare how our model is able to capture temporal binging and content binging under the assumption that individuals are either myopic or capable of thinking a finite number of “steps” ahead (Camerer, Ho, and Chong 2004). In the myopic version of the model, we assume that individuals are maximizing the utility of their choices at each event without consideration of how the current choice may impact the utilities of future events. Allowing individuals to be forward-looking in a boundedly rational way (i.e., thinking one step ahead)

means the model may take into account the utility of future choices when learners are making the current choice.

Using our model, we can examine how consumption in certain courses can be used to predict consumption patterns in other courses, which has implications for new product launch, cross-selling, and bundling. We also test how the timing of content release affects consumption and knowledge accumulation using counterfactual simulations. We then compare our predictions to actual data obtained following a natural experiment policy change that shifted content release from weekly “sequential” release to on-demand “simultaneous” release for the courses that we examine from Coursera. 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. Finally, we can look at how bingeing is related to payment for certificates across different courses, which has implications for monetization of subscription services and CLV (Kumar 2008; Fader, Hardie, and Lee 2005; Zhang, Bradlow, and Small 2015).

1.2 Outline of Paper

In Section 2, we describe a clickstream dataset obtained from Coursera where we observe the lecture and quiz consumption patterns of students across two courses offered by Wharton Online: Marketing and Operations. In addition, we present an overview of our model that endogenizes binge-watching by treating individuals as boundedly rational decision makers when they make their lecture and quiz consumption choices across courses, including the specific choice set and the components of the utility function based on behavioral theories. In our model, we assume that each choice is comprised of a two-stage decision process where individuals first choose a content “category” (Marketing vs. Operations vs. Break), and then select a content “type” within the category (lecture/quiz, short/medium/long break). In Section 3, we provide descriptive evidence of binge consumption within the individuals observed in our dataset; importantly, we distinguish between content bingeing and temporal bingeing.

Section 4 lays out our formal model, including the parameterization of each component of the choice utilities. We also compare and contrast our two-stage decision process model with the standard nested logit. Section 5 presents our estimation procedure and results. We use hierarchical Bayes Markov chain Monte Carlo (MCMC) sampling to obtain parameter estimates, which allows us to account for unobserved heterogeneity and obtain posterior predictive checks at the individual-level for a rigorous assessment of model adequacy. In Section 6, we conduct a series of counterfactuals to see how binge consumption patterns might change with different content release schedules. In Section 7, we look at whether the patterns observed in Marketing and Operations are predictive of behaviors in other courses, akin to assessing the value of this data for different product launches. Section 8 concludes with directions for future research in the area of online dynamic consumption.

2 Data and Model Overview

We examine the behavior of individuals engaged in two courses offered by Wharton Online through Coursera in 2015: Introduction to Marketing (“Marketing”) and Introduction to Operations Management (“Operations”). 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. We then filtered our sample by looking at individuals who had registered for both Marketing and Operations,¹ resulting in 13,690 unique individuals. Since individuals can register without actually consuming any of the content, we further filtered them by whether they had visited at least one

¹We chose to examine individuals within the overlap of these two courses because we are interested in both temporal and content binging.

quiz or lecture within *both* Marketing and Operations, resulting in 2,943 individuals. Finally, we focus on registered users who had also paid for both courses, giving us a final sample of 508 individuals. Thus, we are focusing on the most committed students who have consumed most or all of the course, which allows us to model their binge consumption patterns.²

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 URL arrivals, so 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 arrival. 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. We observed a mean of 83.8 URL arrivals per individual (SD = 51.8).

Each week, new lecture and quiz materials were released, and individuals could revisit material from previous weeks. Figure 1 shows the weekly release schedule of lecture and quiz content for both Marketing and Operations (see Appendix A 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 bingeing, 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.

Figure 2 plots the density of URL arrivals 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 when there is a second spike in activity as individuals “cram” for quizzes. Note that individuals can’t (and don’t) engage in content that has not been made available yet.

²Our model can be applied to a broader set of online courses with similar lecture/quiz structures, and although here we condition on payment, our model can also be extended to make predictions about if and when individuals pay.

³We ignore URL arrivals to 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 1: Sequential Release of Content for Marketing and Operations

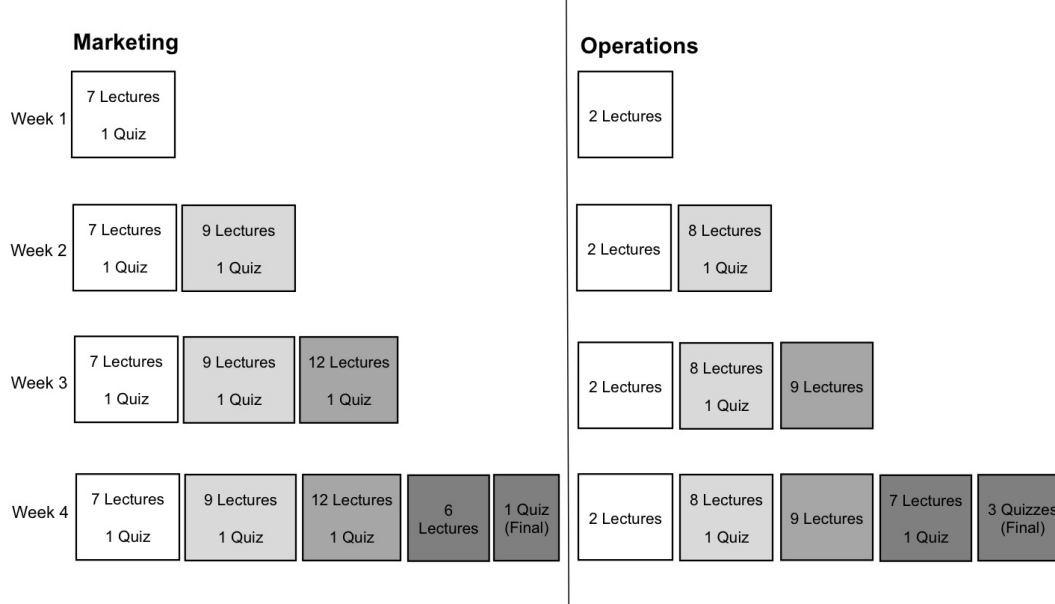


Figure 2: Density of engagement in content separated by week.

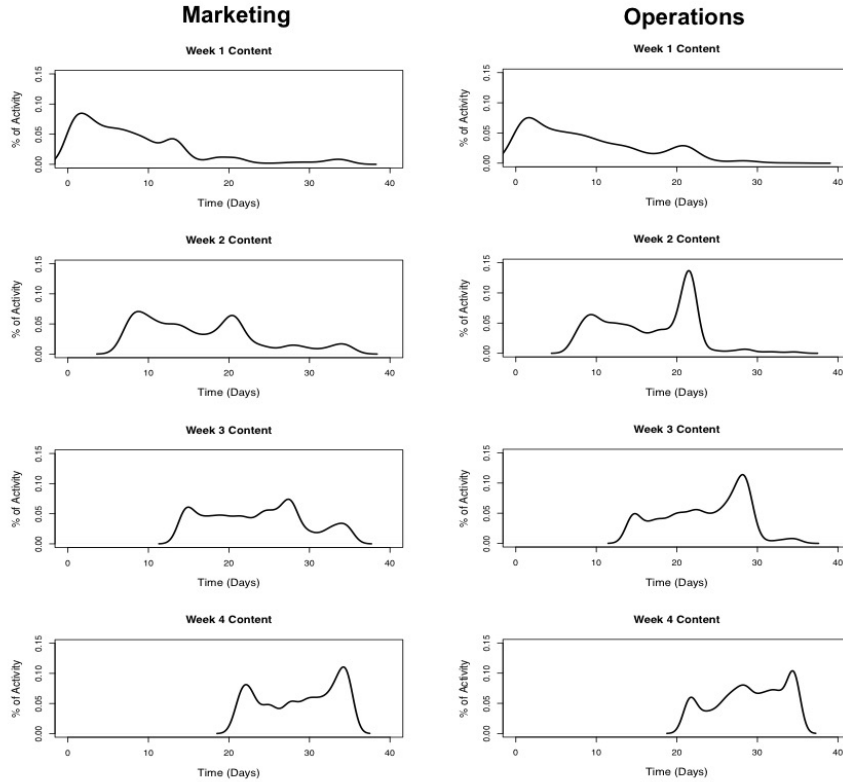


Figure 3 plots the sequence of URL arrivals for a single individual (each arrival is referred to as an “event” j), separated by Marketing and Operations. The lectures and quizzes are numbered from 1 to 38 for Marketing and from 1 to 31 for Operations, in the order of their appearance on the website. In this example, it appears that the individual’s visits are relative evenly distributed across the different lectures and quizzes. In fact, the individual mostly seems to progress systematically through the material by watching lectures and taking quizzes one-by-one in order, with few “skips” ahead or back to previous content.

Figure 3: Sequence of lecture and quiz choices at each event j for a sample individual.


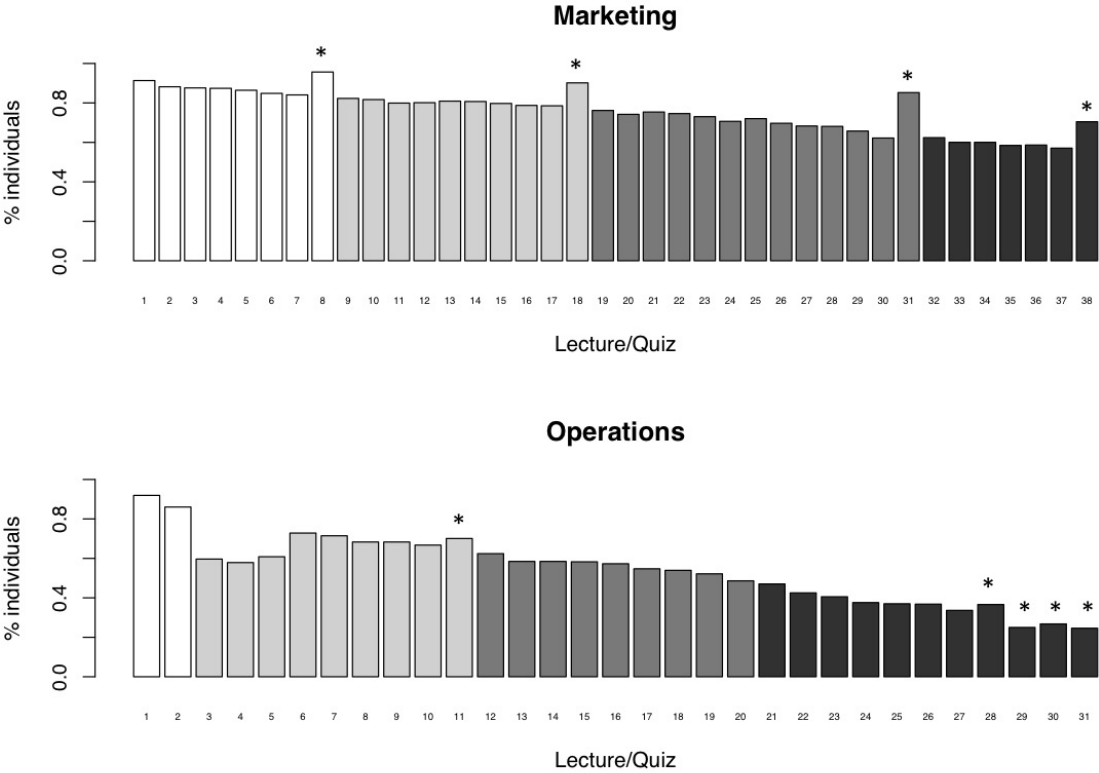


Figure 4 plots bars for each lecture and quiz representing the percentage of individuals who visited that particular lecture or quiz at least once. The shades of grey indicate the week that the content was release, while the stars indicate the bars that represent quizzes. We make three main observations. 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 4: Percentage of individuals who viewed each lecture and quiz at least once



2.1 Choice Model Overview

We assume that at each event j individuals are making a two-stage decision (see Figure 5). 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 5: Two-stage decision process. Individuals can choose between lectures and quizzes within either Marketing or Operations, as well as different ranges of break lengths.

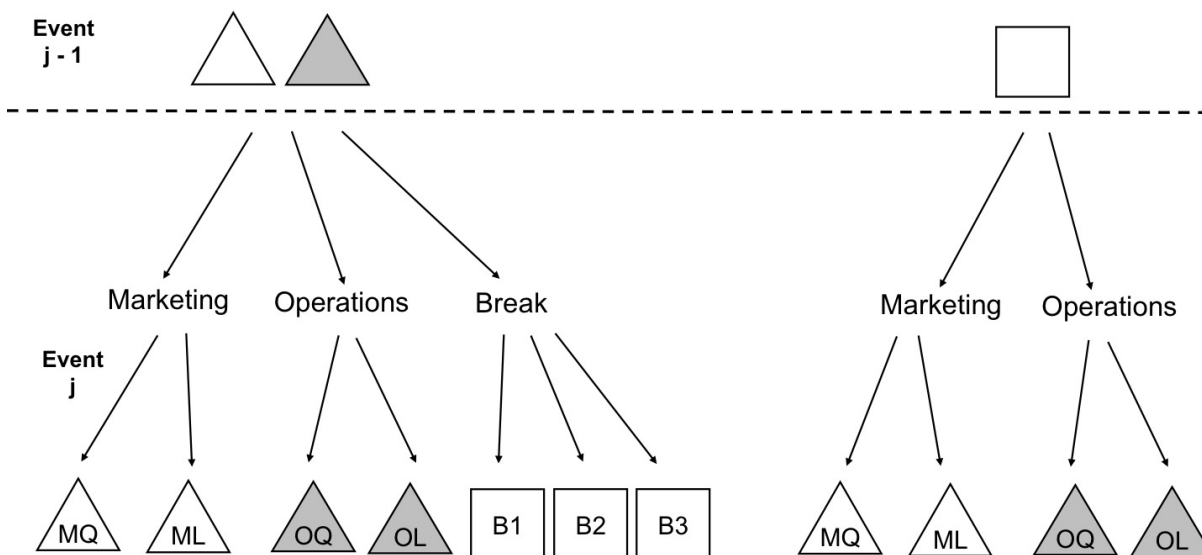
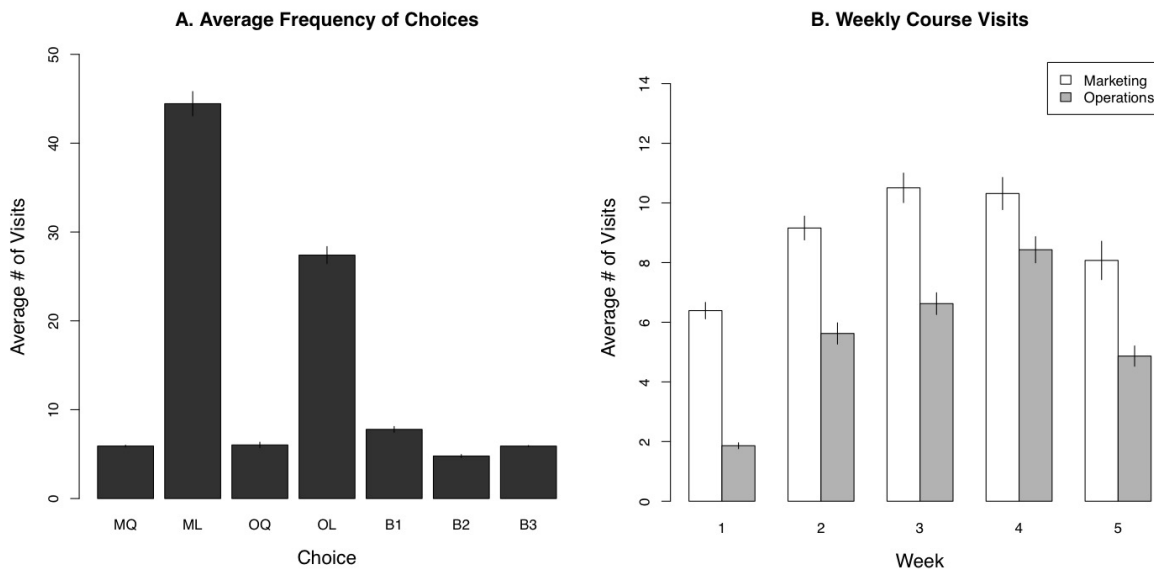


Figure 6 Panel A shows the frequency of each choice across the entire course, averaged across participants. In general, individuals chose Marketing more often than Operations, which is consistent with the fact that Marketing has more lectures and quizzes, as well as less attrition over time. Individuals 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 then the Medium Break (B2). Panel B shows the frequencies for lectures and quizzes within Marketing and

⁴The 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 B.

Operations, divided by week. We see that activity in both courses peaks during Weeks 3 and 4. There is very little activity for Operations in Week 1, possibly because there were no Operations quizzes in Week 1 and therefore little motivation to watch the Operations lectures.

Figure 6: Number of times each choices was made, averaged across individuals.



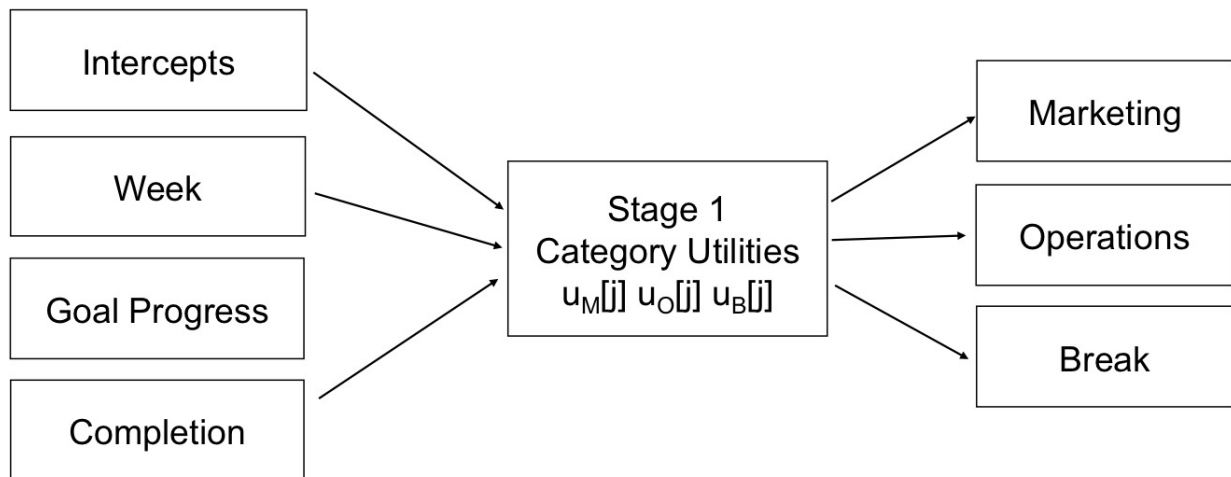
We assume three restrictions on the choice set that impact our model and the likelihood function. First, if no quizzes were available in a particular week (i.e., Operations during week 1) or all quiz attempts were “used up” (i.e., Marketing quizzes were limited to 3 attempts each), then these options were not available in stage 2 of the event.⁵ Second, observed activity of individuals is restricted to the 5 weeks encompassing the latest due date of all quizzes, 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, 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 5).

⁵In 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).

2.2 Model Overview: Stage 1 Choices

Figure 7 depicts the constructs that contribute to the utilities of the options available in stage 1 of each individual’s decision process. In stage 1 of each event j , individuals choose between Marketing, Operations, and Break.

Figure 7: Stage 1 choice at event j .



First, we include intercepts to capture the baseline propensity of individuals to engage in Marketing over Operations and Break. Second, we allow the utilities to vary weekly to capture the patterns shown in Figure 6, which might contribute to some of the content binging observed in the data (i.e., individuals may simply be switching between Marketing and Operations each week for time management purposes).

Third, 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 influenced by individuals goals to complete the course. Goal gradient theory (Kivetz, Urminsky,

and Zheng 2006) predicts that as individuals approach completion of a goal, they accelerate their “effort” towards the goal. 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, individuals may actually become less motivated as they progress through the content 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), which predicts the opposite effect of goal progress on utility.

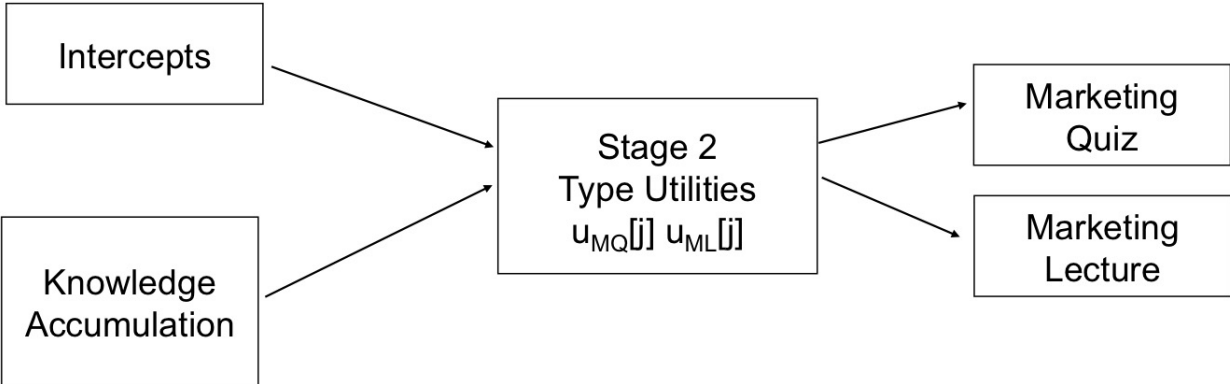
Fourth, we allow our model to capture what might happen to the utilities of the courses once individuals actually reach “completion” of their goal. In extant research that documents goal progress, the task is either “terminal,” that is the researcher stops observing the individual’s consumption patterns 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). 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 or the course ends.

2.3 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 8 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).

Figure 8: Stage 2 choice at event j , conditional on Marketing category being chosen in Stage 1.

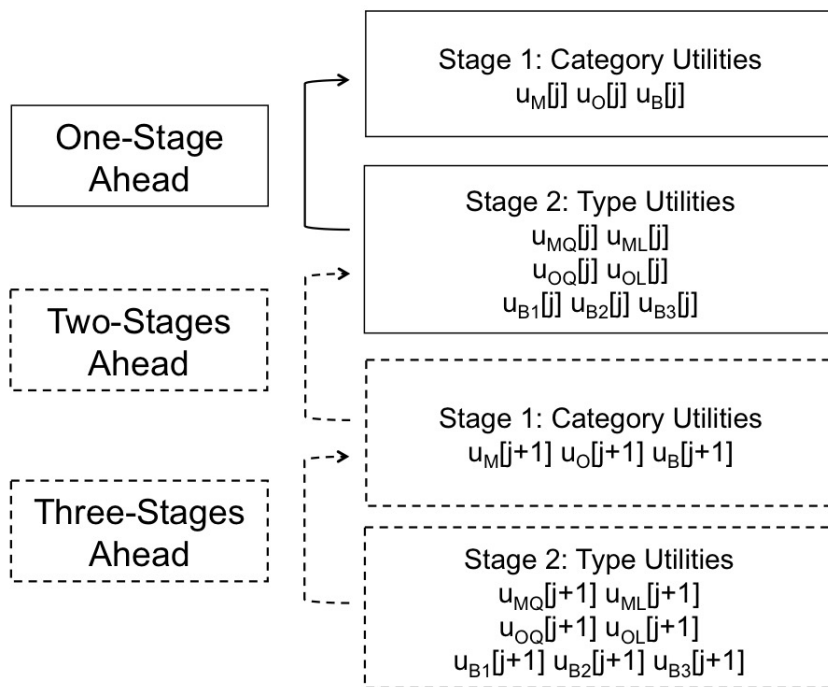


When taking online courses, individuals are motivated to accumulate knowledge as they engage in the course material. Some individuals may be purely interested in gaining knowledge by watching lectures, while others may be more interested in passing the course to earn a certificate, which may be done by passing the quizzes. So we propose that individuals' desire to take quizzes vs. watch lectures may be determined by their accumulated knowledge or beliefs about their quiz-taking abilities, which are updated in a Bayesian way as they take quizzes or watch lectures.

2.4 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 (see Figure 9). In our model, individuals make choices as if they were capable of thinking at most one stage ahead. This “forward-lookingness” can be extended farther into the future (i.e., two stages ahead, three stages ahead, etc.). The computation of utilities farther and farther into the future quickly becomes very intensive, and previous work has demonstrated that consumers may only be able to think a few steps ahead (Camerer, Ho, and Chong 2004). Limited time-horizon or “boundedly rational” models of consumer behavior may better account for the patterns of choices observed within individuals compared to infinite time-horizon models (Gabaix et al., 2006).

Figure 9: Boundedly rational forward-looking consumers during events j and $j + 1$.



Consumers may be myopic such that a decision 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. 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.

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 stage ahead may contribute to temporal binging; for example, if individuals anticipate greater utilities from stage 2 if they engage in Marketing or Operations in stage 1, then they may be less likely to choose Break in stage 1.

In the Estimation section, we will present results comparing the myopic and One-Stage Ahead models. We note that our model can be generalized to individuals thinking more stages ahead. For example, individuals may be “Two-Stages Ahead” such that during stage 1 of event j , they take into account the expected utilities from stage 2, and also assume that at stage 2 they will take into account the expected utilities from the choices in stage 1 of the next event $j + 1$ (note that these utilities then propagate up to the utilities in stage 1 of the current event j). Thinking Two-Stages Ahead may allow individuals to anticipate the effects of future goal progress or knowledge accumulation (see Appendix C for the parameterization of a Two-Stages Ahead term).

In summary, the stage 1 utilities of Marketing, Operations, and Break depend on the following constructs: Intercepts, Weekly Effects, Goal Progress, and Completion (see Figure 7). 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 and Accumulated Knowledge (see Figure 8). Moreover, individuals may be thinking One-Stage Ahead when incorporating the utilities from stage 2 to make a decision in stage 1 (see Figure 9).

3 Descriptive Analysis

Before we introduce our formal model, we present results from exploratory analyses of the data to provide descriptive evidence of binge consumption, specifically distinguishing between temporal bingeing and content bingeing. Using metrics of temporal and content bingeing, 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 bingeing patterns of the observed data, using posterior predictive checks (Gelman et al. 2014), as described in Section 5. These metrics can also be used to determine whether individuals are more likely to binge based on the counterfactual simulations described in Section 6, where we make predictions about consumption patterns with different content release schedules.

In order to assess each individual’s degree of temporal bingeing, we can look at the average length of consecutive quiz/lecture events (“runs”) with no breaks. Longer runs correspond to more temporal bingeing. To assess the degree of content bingeing, 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 bingeing. Figure 10 plots the average run lengths against the average percentage of non-switches (one dot for each person in our sample), which shows that there is a positive relationship between the two bingeing metrics ($r = 0.43$, $t(506) = 0.47$, $p < 0.001$) 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.

In order to determine whether or not individuals were statistically significant temporal bingers or content bingers, 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.

Figure 10: Distribution of average run length and percentage of non-switches

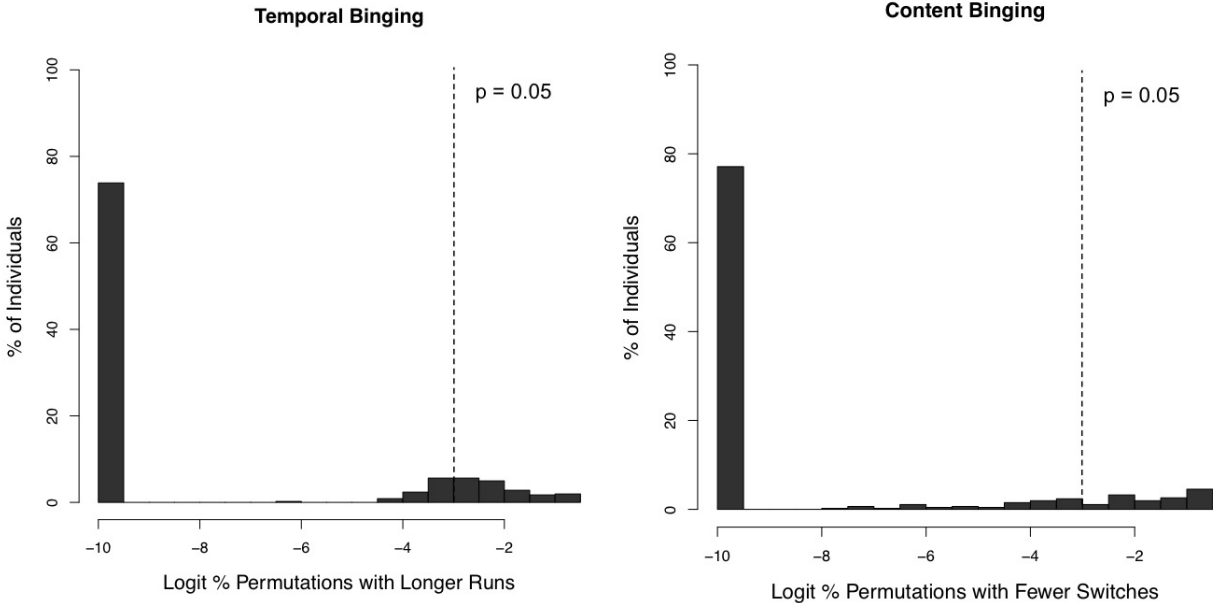


To determine whether or not individuals were temporal bingers, we calculated the percentage of random permutations with (strictly) longer average runs compared to the individual’s actual sequence. To determine whether individuals were content bingers, we calculate the percentage of random permutations 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 11 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 11. 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.

Figure 11: Distributions of logit-transformed p-values for temporal binging and content binging



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.

4.1 Two-Stage Decision Process

For each individual $i = 1, \dots, I$ we observe a sequence of events $j = 1, \dots, J_i$. At each event j , the individual makes a two-stage decision (see Figure 5) 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 among the categories of 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.

$$\mathcal{L}(\theta_i) = \prod_{j=1}^{J_i} P(S_2[j]|S_1[j]) \times P(S_1[j]) \quad (1)$$

Table 1 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.

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 7).

$$\begin{aligned} \text{Stage 1: } u_M[j] = & \text{Intercepts}_M[j] + \text{Week}_M[j] + \text{Goal Progress}_M[j] \\ & + \text{Completion}_M[j] + \epsilon_M[j] \end{aligned} \quad (2)$$

Table 1: 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)
Week	β_{1M}, β_{1O}	–	Week 1
	β_{2M}, β_{2O}	–	Week 2
	β_{3M}, β_{3O}	–	Week 3
	β_{4M}, β_{4O}	–	Week 4
Goal Progress	β_{5M}, β_{5O}	$G_M[j]$ or $G_O[j]$	Linear
	β_{6M}, β_{6O}	$G_M[j]^2$ or $G_O[j]^2$	Quadratic
Completion	β_{7M}, β_{7O}	$\mathbb{1}(G_{MQ}[j] = 1)$ or $\mathbb{1}(G_{OQ}[j] = 1)$	Quizzes
	β_{8M}, β_{8O}	$\mathbb{1}(G_{ML}[j] = 1)$ or $\mathbb{1}(G_{OL}[j] = 1)$	Lectures
Knowledge	α_{1M}, α_{1O}	$\pi_M[j]$ or $\pi_O[j]$	Quiz-Abilities
	α_{2M}, α_{2O}	$\log(T_{ML}[j])$ or $\log(T_{OL}[j])$	Lectures
	π_{0M}, π_{0O}	–	Initial beliefs (Mean)
	η_{0M}, η_{0O}	–	Initial beliefs (Precision)
Forward-Looking	γ	E	One-Stage Ahead

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 8).

$$\text{Stage 2: } u_{MQ}[j] = \text{Intercepts}_{MQ}[j] + \text{Knowledge Accumulation}_{MQ}[j] + \epsilon_{MQ}[j] \quad (3)$$

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” forward-looking, without loss of generality, as more stages can be added (as depicted in Figure 9). 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:

$$E = \log\left(e^{u_{MQ}[j]} + e^{u_{ML}[j]}\right) \quad (4)$$

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 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.

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.

4.3 Week Effect

Let $\{\beta_{1M}, \beta_{2M}, \beta_{3M}, \beta_{4M}\}$ and $\{\beta_{1O}, \beta_{2O}, \beta_{3O}, \beta_{4O}\}$ capture the weekly dummy effects for Marketing and Operations, respectively, with week 5 set to 0 as the reference.

⁶Alternatively, 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:

$$E = \frac{e^{u_{MQ}[j]}u_{MQ}[j] + e^{u_{ML}[j]}u_{ML}[j]}{e^{u_{MQ}[j]} + e^{u_{ML}[j]}} \quad (5)$$

4.4 Goal Progress

Let $G_M[j]$ and $G_O[j]$ represent the percentage of all available quizzes and lectures in Marketing and Operations, respectively, that the individual has visited at least once by event j .⁷ Under the sequential content release schedule of the courses in the observed data, the number of available lectures and quizzes changes each week in both Marketing and Operations. Later we will use the estimated model parameters to conduct counterfactual analyses of policies where all content is released on the first day of the first week.

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 $\{\beta_{5M}, \beta_{6M}\}$ and $\{\beta_{5O}, \beta_{6O}\}$ on the percentage and squared percentages of visited quizzes and lectures for each course. The following expression gives the Goal Progress effect for Marketing at event j .

$$\text{Goal Progress}_M[j] = \beta_{5M}G_M[j] + \beta_{6M}G_M[j]^2 \quad (6)$$

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 12 illustrates the values of the linear and quadratic terms within the Goal Progress construct that would result in different curves for how utility changes relative to progress.

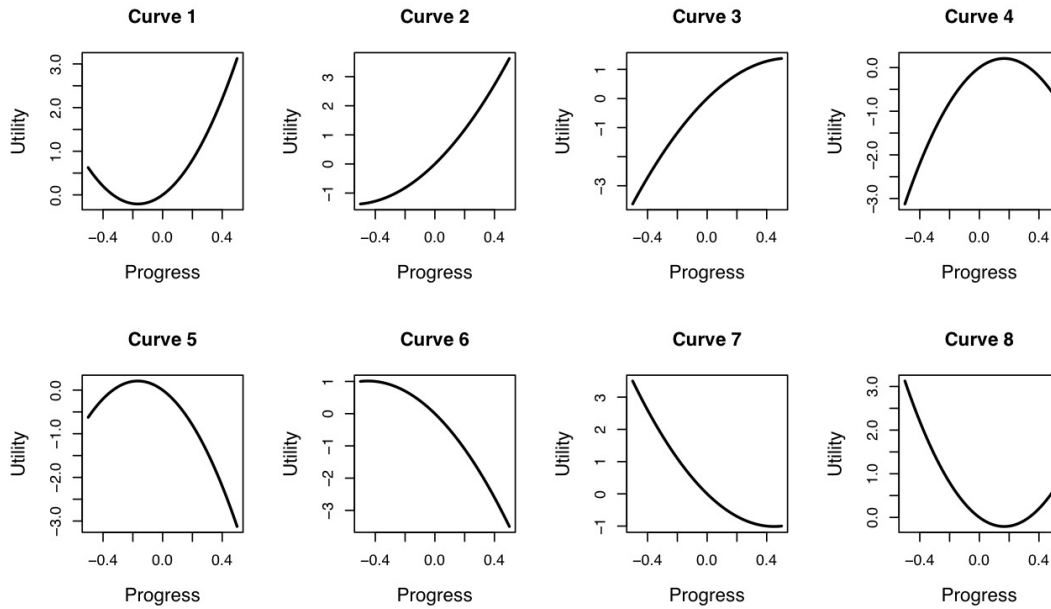
Figure 12 Panel A shows 8 possible shapes of how the utility of a course option changes with goal progress, which represents the percentage of quizzes and lectures visited at least once so far (normalized to be between -0.5 and 0.5). The curves may be monotonically increasing or decreasing, either exponentially 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 1 and 8 illustrate the “stuck

⁷Note that we center these values around 0 by subtracting 0.5.

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 6 and 7 illustrate mispredictions in future time slack (Zauberman and Lynch 2005).

Figure 12: Illustration of the Goal Progress construct

A. Goal Progress Curves



B. Corresponding Linear and Quadratic Values

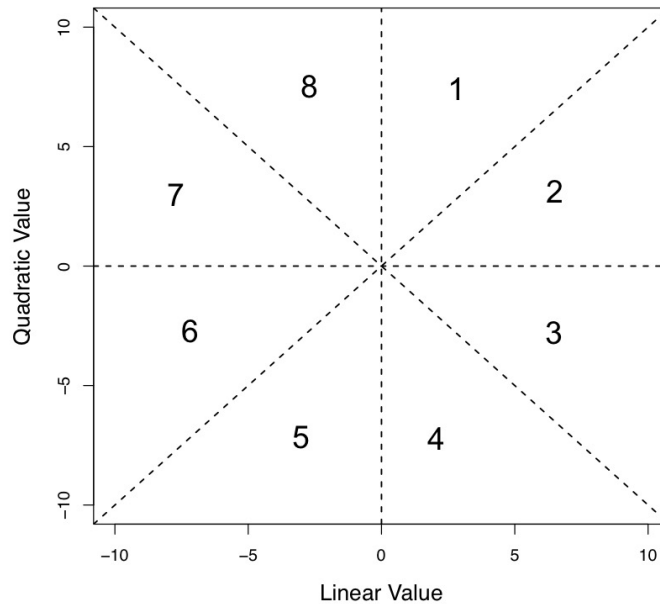


Figure 12 Panel B shows the corresponding values of the linear (β_{5M} and β_{5O}) and quadratic (β_{6M} and β_{6O}) coefficients that would result in each of the 8 Goal Progress curves. Using each individual’s parameter estimates, we can classify individuals in our sample as different types of Goal Progress learners, similar to how Gilbride and Allenby (2002) classify individuals by the type of screening rules they use during choice.

4.5 Completion

Before new content is released to “reset” the goal of visiting all lectures and quizzes at least once, we 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. Since the Goal Progress component of our model is not able to capture this “completion effect,” we use indicator variables, with coefficients $\{\beta_{7M}, \beta_{8M}\}$ for Marketing and $\{\beta_{7O}, \beta_{8O}\}$ 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 .

$$\text{Completion}_M[j] = \beta_{7M}\mathbb{1}(G_{MQ}[j] = 1) + \beta_{8M}\mathbb{1}(G_{ML}[j] = 1) \quad (7)$$

4.6 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. 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.

At event $j = 1$, the individual starts with beliefs about her own quiz-taking abilities within Marketing and Operations, with mean $\pi_M[1] = \pi_{0M}$ or $\pi_O[1] = \pi_{0O}$, and precision

$$\eta_M[1] = \eta_{0M} \text{ or } \eta_O[1] = \eta_{0O}.$$

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.

$$\pi_M[j + 1] = \frac{\eta_M[j]}{\eta_M[j] + \phi} \pi_M[j] + \frac{\phi}{\eta_M[j] + \phi} x[j] \quad (8)$$

$$\eta_M[j + 1] = \eta_M[j] + \phi \quad (9)$$

Knowledge may also accumulate when individuals watch lectures. Let $T_{ML}[j]$ and $T_{OL}[j]$ represent the total number of times that the individual visited lectures in Marketing and Operations by event j , with α_{2M} and α_{2O} as the coefficients. The following gives the expression for Knowledge Accumulation within Marketing, which is an empirically determined weighted average of knowledge beliefs based on feedback from quiz scores and watching lectures:

$$\text{Knowledge Accumulation}_{MQ}[j] = \alpha_{1M} \pi_M[j] + \alpha_{2M} T_{ML}[j] \quad (10)$$

5 Estimation

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; Brooks and Gelman 1997). After thinning the chains to reduce auto-correlation, we were left with 300 posterior samples for each parameter.

To verify that our model is identified, 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 508 individuals). We then estimated the model using this simulated data to determine that we could recover the true parameter values (see results in Appendix D). A broader set of simulation results (available upon request) demonstrate the identification of our model under a variety of parameter conditions.

In the following subsections, we describe the estimation results for different versions of our model. 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 model fit. To compare the fit across the model variations, we calculated the Deviance Information Criterion (DIC)⁸ 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.

5.1 Model Comparison

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

⁸We calculate DIC for each model using the equation $-4E_{\theta} \log(p(y|\theta)) + 2\log(p(y|\hat{\theta}))$, where y represents the observed data and $\hat{\theta}$ represented the estimated parameter values (Gelman et al. 2014).

model, as shown in Table 2. We see that adding the parameters for the Week, Goal Progress, Completion, Knowledge Accumulation, and One-Step Ahead constructs improves the model fit in terms of decreasing the DIC. Note that the fit statistics in Table 2 are for nested models, for example the “Goal Progress and Completion” model includes the constructs of Intercepts, Week, Goal Progress, *and* Completion.

Based on the DIC, the full model with all the constructs is the winning model. However, 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).

Table 2: Models fit statistics for different nested versions of the model.

	DIC	Run Length	% Switches
Observed	N/A	5.19	0.09
Intercepts	1148920	5.19	0.43
Week	1141621	5.61	0.34
Goal Progress + Completion	1134736	4.02	0.31
Knowledge Accumulation	1125537	4.63	0.31
One-Step Ahead	1124651	4.78	0.30

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 508 individuals in our sample. 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 the 5 weeks of the course.

Although we treat individual choices as discrete in the model estimation, individuals are also moving through continuous calendar time (for example, the calendar time determines the value of the dummy variables in the Week construct, as well as what lectures and quizzes are 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, 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 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 allows her to update her quiz-taking abilities. See Appendix E for further details.

After obtaining the simulated choice sequences for each individual sample, we can examine how closely the simulated patterns fit the observe data. First, we look at whether the models can capture the temporal and content binging patterns of the data, which was the central motivation for this paper. Temporal binging can be proxied by the average run length (without breaks), while content binging can be proxied by the percentage of choices where individuals switch between Marketing and Operations (disregarding breaks), as we originally described in the Descriptive Analysis section.

In addition to the DIC values, Table 2 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 ob-

served percentage, we note that there is no state dependence parameter in our model, so the patterns of content bingeing are being captured purely by the theory-driven constructs in our model. In future work, we can examine how state dependence and our constructs are related.

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 (see Figure 13).

Figure 13: Individual-level posterior predictive checks for the observed vs. simulated data, comparing the total visits to each choice, the average run lengths, and the % switches.

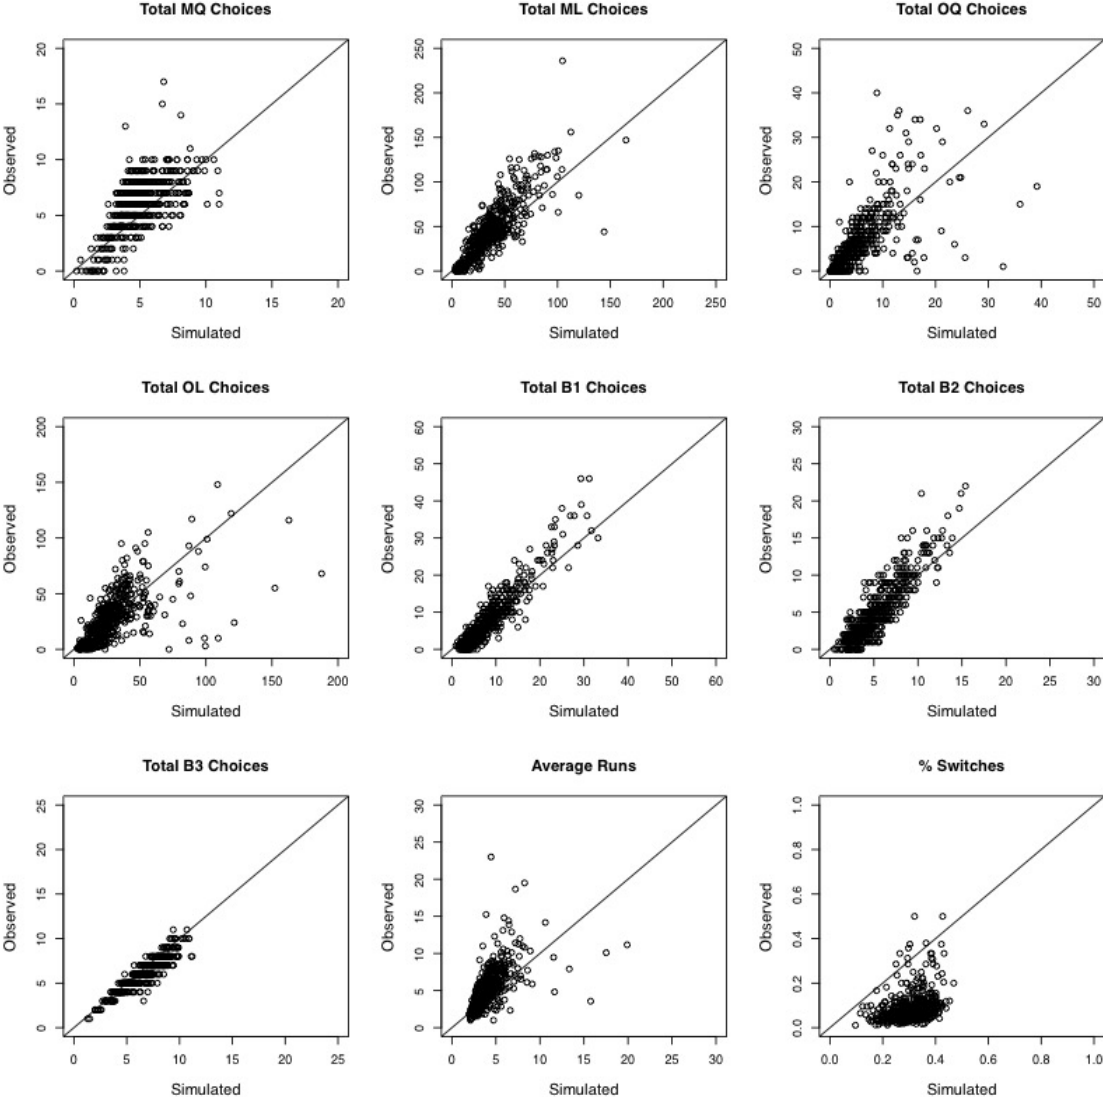


Figure 13 plots the observed vs. simulated values for the total number of times that individuals selected each of the 7 choices (MQ, ML, OQ, OL, B1, B2, B3), as well as the run length and % switches. Each point in each plot represents the mean across simulations for one individual. For most of the statistics, the simulated values roughly match the distribution of observed values, noting that the simulated total choices are sensitive to the values that are used for the calendar time that each event takes up.

5.2 Estimated Parameters: Stage 1

We present the estimated parameters for the full 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 $\theta_i \sim MVN(\mu, \Omega)$.

Table 3 shows the results of the estimation of the parameters in Stage 1, while the results for Stage 2 are presented in the next subsection. In Table 3, 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 .

Based on the values of the estimated parameters in the Intercepts construct, we see that Marketing has a higher baseline utility β_{0M} compared to Operations β_{0O} , which accounts for the higher choice shares for Marketing in stage 1. The parameters in the Week construct, $\{\beta_{1M}, \beta_{2M}, \beta_{3M}, \beta_{4M}\}$ and $\{\beta_{1O}, \beta_{2O}, \beta_{3O}, \beta_{4O}\}$, capture the pattern seen in Figure 6 Panel B where the choice shares for Marketing is higher in the earlier weeks, while the choice shares for Operations peaks around Week 4.

Looking at the parameter values of the Goal Progress construct, we can determine which of the 8 curves depicted in Figure 12 describes the change in utility with respect to goal progress for the population as a whole, as well as for each individual. In Marketing, the linear term has a mean of $\beta_{5M} = 0.67$ while the quadratic term has a mean of $\beta_{6M} = -0.36$, which corresponds to Curve 3. Thus, there appears to be a small “goal gradient” effect at the aggregate level in Marketing, so the utility of Marketing increases as individuals get closer

to visiting all quizzes and lectures at least once. On the other hand, for Operations, the linear and quadratic terms have means of $\beta_{5O} = 0.49$ and $\beta_{6M} = -4.09$, which corresponds to Curve 4. Thus, there appears to be a large non-monotonic effect such that the utility of Operations first increases and then decreases as individuals near completion.

Table 3: Summary of Estimated Parameters in Stage 1

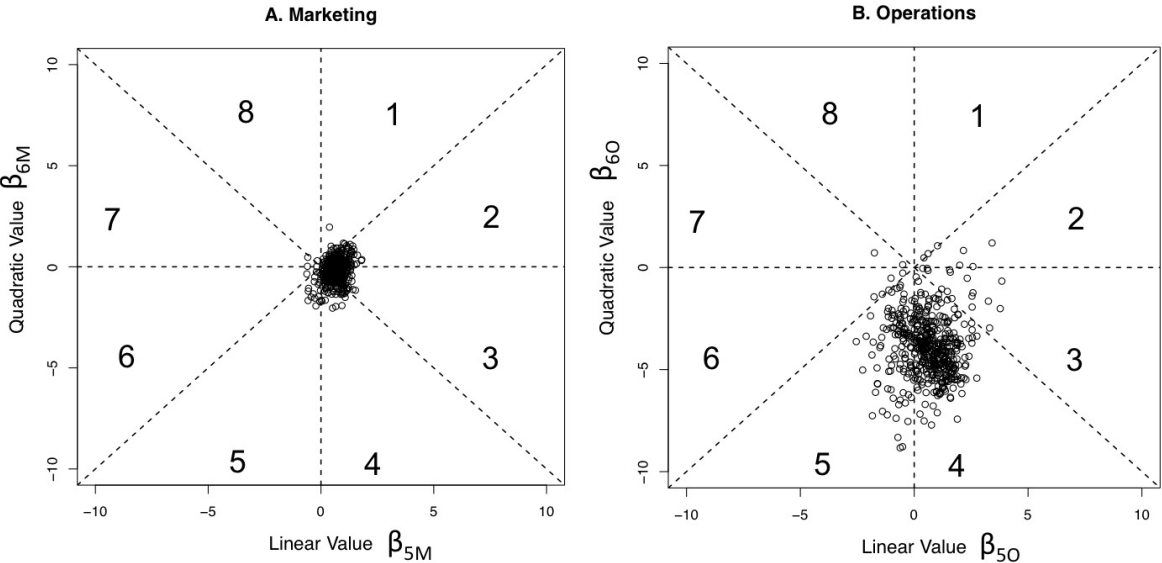
Construct	Parameter	Mean μ	95% CI	Mean σ^2
Intercepts	β_{0M}	1.14	[1.03, 1.24]	0.31
	β_{0O}	0.78	[0.68, 0.88]	0.48
Week	β_{1M}	0.22	[0.13, 0.30]	0.32
	β_{2M}	0.12	[0.02, 0.21]	0.38
	β_{3M}	-0.01	[-0.10, 0.07]	0.48
	β_{4M}	-0.36	[-0.48, -0.27]	0.69
	β_{1O}	-0.26	[-0.39, -0.12]	0.75
	β_{2O}	-0.20	[-0.34, -0.03]	1.27
	β_{3O}	-0.05	[-0.22, 0.12]	1.42
	β_{4O}	0.01	[-0.13, 0.13]	0.92
Goal Progress	β_{5M}	0.67	[0.57, 0.78]	0.58
	β_{6M}	-0.36	[-0.60, -0.09]	0.58
	β_{5O}	0.49	[0.33, 0.72]	2.18
	β_{6O}	-4.09	[-4.48, -3.69]	7.87
Completion	β_{7M}	-2.10	[-2.26, -1.93]	1.44
	β_{8M}	-0.40	[-0.50, -0.30]	0.49
	β_{7O}	-0.65	[-0.84, -0.46]	1.12
	β_{8O}	-0.76	[-0.97, -0.57]	1.10
One-Step Ahead	$\log(\gamma)$	-1.66	[-2.07, -1.42]	0.87

In Figure 14, we plot the posterior means of the linear and quadratic terms within the Goal Progress construct for each individual 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, 45% of individuals have parameters that classify them as Curve 3 and 29% are classified as Curve 2, which both indicate a goal gradient process, differing only by whether the shape of the increase is convex or concave. In Operations, 71% of individuals are classified as Curve 4, while 23% are classified as Curve 5, which both indicate a non-monotonic effect where utility increases and then decreases.

One 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. This may be interpreted as a slack effect where 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.

Figure 14: Goal Progress individual-level parameter estimates for Marketing and Operations



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 ($\beta_{7M} = -2.08$).

The log of the One-Step Ahead construct γ has a value of -1.66, which corresponds to $\gamma = 0.19$. 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.

5.3 Estimated Parameters: Stage 2

Table 4 shows the estimates for the stage 2 parameters. Note that when estimating the models, we set the initial means of the quiz-taking abilities to $\pi_{0M} = \pi_{0O} = 0$ and the precision to $\eta_{0M} = \eta_{0O} = 1$. This can be interpreted as individuals having a “weak” prior belief that they have zero 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 4: Summary of Estimated Parameters in Stage 2

Construct	Parameter	μ	95% CI	σ^2
Intercepts	α_{0M}	-2.12	[-2.31, 1.98]	0.55
	α_{0O}	-4.35	[-4.71, -4.11]	2.78
	δ_1	0.01	[-0.07, 0.10]	0.50
	δ_2	-0.40	[-0.48, -0.31]	0.40
Knowledge	α_{1M}	2.49	[2.20, 2.81]	0.56
	α_{2M}	-0.60	[-0.68, -0.51]	0.15
	α_{1O}	1.57	[1.24, 1.86]	5.30
	α_{2O}	0.51	[0.38, 0.67]	0.53

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.

Watching more lectures has a negative effect on the utility of taking quizzes in Marketing (α_{2M}), and a positive effect in Operations (α_{2O}). In Marketing, it is possible that individuals who are doing poorly on quizzes simply need to watch more lectures. In Operations, the

effect is more 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.

6 Counterfactuals

The data where we observe individuals engaged in Marketing and Operations occurred in 2015. Since the two courses were first offered in 2013, the structure of the content has undergone two main changes. First, the number of lectures and quizzes was reduced in 2014 before the period where we observed our data. Second, in 2016, the courses switched from a “sequential” release schedule where new content was made available each week for 4 weeks, to a “simultaneous” or on-demand release schedule where all content for the course was available starting on day 1.

It is of interest to the firm and the academic institution offering the content how these changes to the courses affect individuals’ overall engagement in the course, how much they progress through the material each week, and how much knowledge they accumulate. Thus, we conduct two counterfactual simulations, where we first vary the total number of available lectures and quizzes each week, and then change the lectures and quizzes from sequential to simultaneous release. In both counterfactuals, we examine how the total number of visits to content in both courses changes. We also look at whether there are any changes in binge-watching patterns, based on our proxies for temporal binging (average run length) and content binging (% switches), as well as whether individuals end up with higher or lower quiz-taking abilities.

For the second counterfactual, where content is assumed be released simultaneously rather than sequentially, we are able to empirically verify our predictions using a *new* data set of individual clickstream behavior for the Marketing and Operations courses in 2016 from *after* the switch was made to simultaneous content.

6.1 Simulation

First, we conducted a counterfactual simulation where we doubled the number of lectures and quizzes offered each week in both Marketing and Operations. Second, we conducted a counterfactual simulation where all the lectures and quizzes for each course were available on the first day of the first week, rather than being released sequentially on a weekly schedule.

In both counterfactuals, the main factor that is changing is the number of available lectures and quizzes at any given time, which in turn affects the percentage of lectures visited so far. Thus, we expect the Goal Progress construct, which takes into account the percentage of content visited, to drive the differences in the simulated statistics.

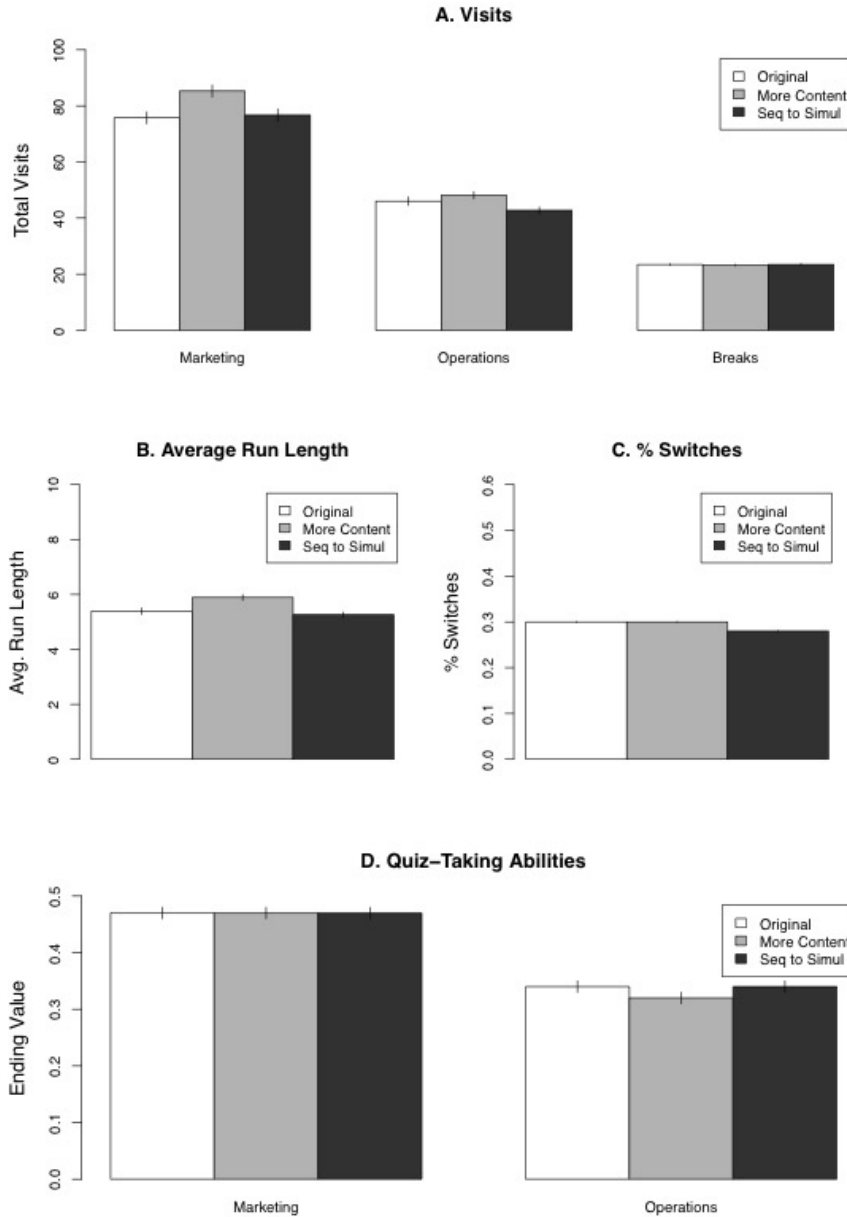
Figure 15 compares the statistics of the simulated choice pathways for the original dataset to the statistics from the counterfactual scenarios (“More Content” and “Sequential to Simultaneous”). Panel A compares the total visits to Marketing, Operations, and Breaks (i.e., stage 1 choices). Panels B and C compare the average run lengths and the percentage of switches. Panel D compares the final quiz-taking abilities for Marketing and Operations (i.e., the value of $\pi_M[j]$ and $\pi_O[j]$ at the last event $j = J_i$).

The counterfactual simulations predict that when there are more lectures and quizzes available, individuals will engage more in both Marketing and Operations, and will also exhibit more temporal binging with longer average run lengths. However, we see that although individuals visit more content, they actually seem to end up with lower quiz-taking abilities in Operations. One reason why we don’t see any differences in the quiz-taking abilities for Marketing is that there is a limit to the number of times individuals can attempt Marketing quizzes, and therefore quiz-taking abilities are not further updated with more visits to Marketing content. These changes are driven by the assumptions in our model that individuals account for both the contemporaneous utility of consuming course content (i.e., being motivated to make progress) and also the long-term accumulation of knowledge.

When content is released simultaneously, we predict there to be little change in the total number of visits, but there may be a slight decrease in the percentage of switches, which corresponds to more content binging. In the next subsection, we will examine how the dis-

tribution of content visits across weeks changes when the release schedule switches from sequential to simultaneous.

Figure 15: Comparison of choice statistics for original simulations and two counterfactuals



Although individuals taking online courses only need to make a one-time payment to access the course content, that is they don't pay in proportion to how much course content they end up consuming, how much they engage in the course could affect other marketing out-

comes such as ad revenue (Schweidel and Moe 2016) and CLV (Kumar 2008; Fader, Hardie, and Lee 2005; Zhang, Bradlow, and Small 2015). Similarly, how much knowledge they gain when taking a course could have downstream consequences on their future engagement in other courses on the platform.

6.2 Empirical Verification of Simultaneous Release Counterfactual

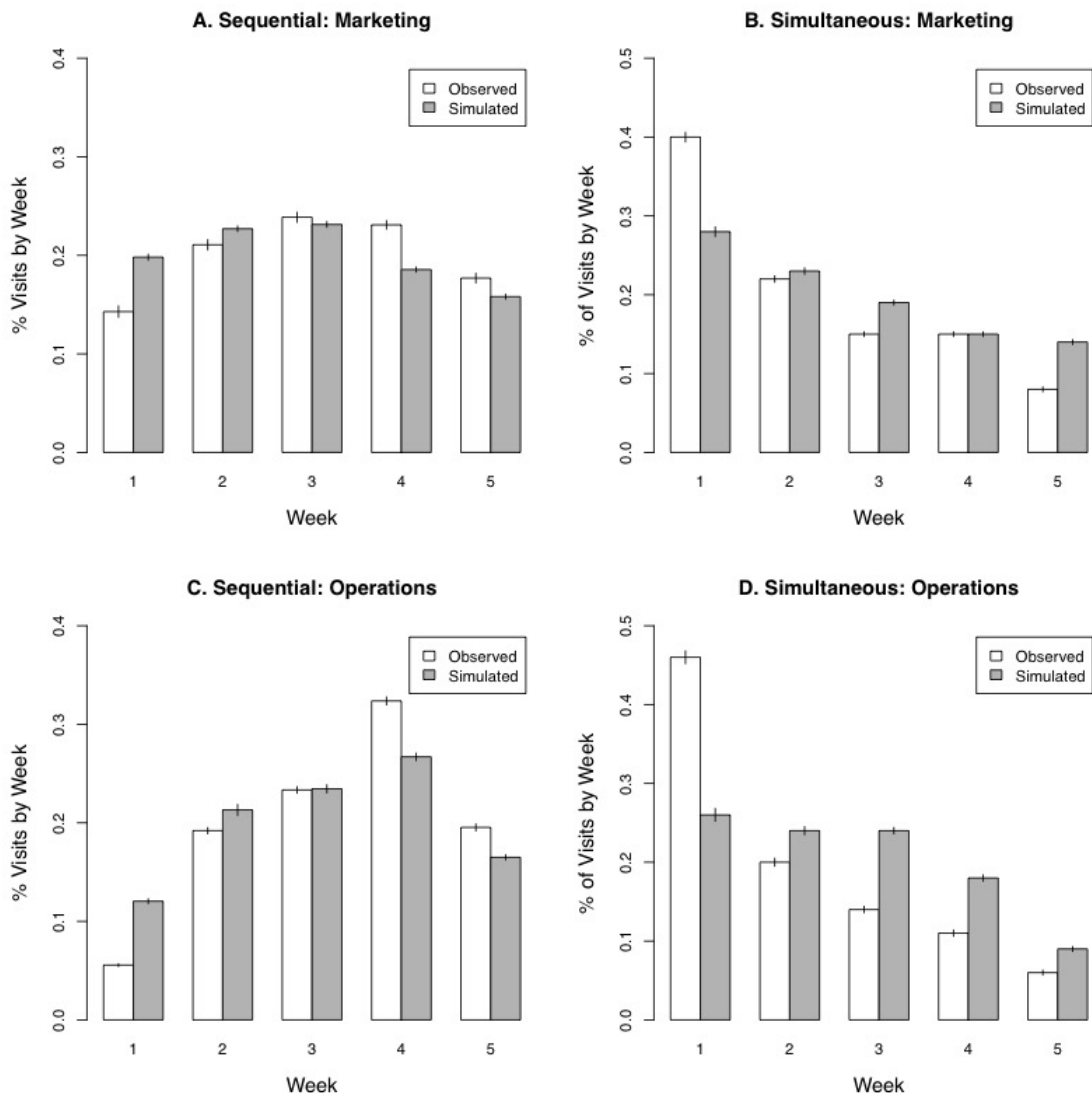
For the simultaneous release counterfactual, we are able to empirically validate our counterfactual simulations using a dataset on individual engagement in Marketing and Operations in 2016, *after* a simultaneous release or “on-demand” policy was implemented. From the day that individuals registered for the courses, they had access to *all* the content, instead of having to wait each week for new content to be released.

This new data sample consists of 1907 individuals who had paid for both Marketing and Operations. Since individuals were not restricted to any specific 5 weeks to complete the course, as with the sequential release policy, for each individual we set their “start date” to the first day that they engaged in the course, and look at their activity in the subsequent 5 weeks.

To compare the fit statistics from the counterfactual simulations to the observed patterns in the data, we examine individual’s weekly engagement in Marketing and Operations in each of the 5 weeks. Figure 16 compares the observed and simulated data for both sequential and simultaneous release, separately for Marketing and Operations.

In Panels A and C, we assess the fit of our model for the original data obtained during the sequential release policy. We compare the observed data to the data generated from simulations using the posterior MCMC samples. Panel A plots the distribution of visits for Marketing content across the 5 weeks of the course for the observed and simulated data, averaged across individuals in the sample. We see that our model is able to simulate the non-monotonic pattern of weekly visits, where consumption peaks at Week 3. Similarly, Panel C plots the distribution of weekly visits for Operations content, and again our model captures the non-monotonic pattern where consumption peaks at Week 4.

Figure 16: Weekly course visits with sequential vs. simultaneous content release, comparison of observed and simulated data



In Panels B and D, we examine how well our counterfactual simulations are able to capture the patterns in the new data obtained after the shift to the simultaneous release policy.⁹ In our counterfactual simulations, we predict the weekly distribution of activity to shift from the non-monotonic pattern observed under the sequential release policy (see Panel A for Marketing and Panel C for Operations), to a decreasing pattern (see Panel B for Marketing and Panel D for Operations). The large spike in the observed activity in Week 1 is likely

⁹Since individuals observed under the simultaneous release policy did not have a common “start date,” we use the time that they first visited content in each course as the course’s start date when plotting the distribution of content visits across weeks. We also apply this rule when plotting the corresponding counterfactual simulations to make the distributions more comparable.

due to individuals self-selecting to consume content when they have a lot of time, and since more content is available, they can make more progress in a single week.

In summary, we conducted a counterfactual simulation where the content release policy was changed from weekly sequential to on-demand simultaneous release. 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 directionally capture the consumption patterns of totally different individuals observed a year later taking the same Marketing and Operations course under a new content release schedule.

7 Predicting Engagement in Other Courses

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).

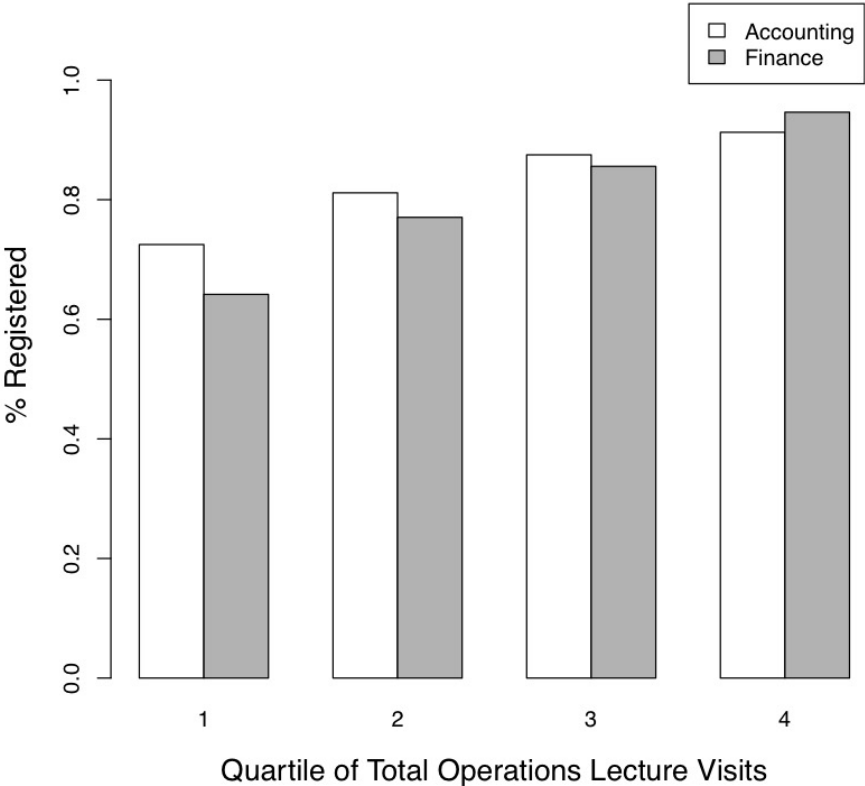
In 2015 when we observed our main data set, two other courses were being offered in addition to Marketing and Operations as part of the Wharton Online Business Foundations package: Introduction to Financial Accounting (“Accounting”) and Introduction to Finance (“Finance”). Among the 508 individuals in 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).

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 registration 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 17, 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 17: Operations Lecture Visits and Cross-Course Registration



8 Discussion

In this paper, we investigated binge consumption within an online learning setting where individuals can consume content by watching lecture videos, as well as by taking quizzes that evaluate their accumulated knowledge. We find evidence that individuals engaged in two online courses, Marketing and Operations, are both temporal binging by consuming a lot of content in succession with few breaks in between, and content binging by switching infrequently between courses.

To capture these consumption patterns, we endogenized binge-watching by modeling each individual’s observed activity as a series of 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 well as the utility from knowledge accumulation. We then conducted counterfactuals to test how consumption patterns may change with different content release formats, and tested our predictions using a new dataset obtained after a policy change regarding the content release schedule.

One feature of our model is that we treat individual decisions to engage in content as discrete choices, 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 differential 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 endogenize 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 far,

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 assumed that individuals are boundedly rational and capable of thinking a finite number of steps ahead. We compared the fit of our model under the assumption that individuals are myopic versus thinking One-Stage Ahead, where they factor in the utilities of the second stage of the decision process (i.e., choosing between lectures/quizzes or short/medium/long break lengths) into the utilities of the first stage (i.e., choosing between Marketing, Operations, and Break). We found that adding the One-Stage Ahead term resulted in a better model fit and simulated temporal bingeing patterns that were closer to observed patterns.

Our model can further be generalized to individuals thinking multiple stages ahead. In particular, a Two-Stages Ahead model would allow us to endogenize the length of the breaks between content consumption. For example, in a myopic model, the length of a break may be modeled as a set of probabilities in a discrete case (i.e., short vs. medium vs. long breaks) or as some type of hazard function in a continuous case (Seetharaman and Chintagunta 2003; Schweidel and Moe 2016). Alternatively, if individuals are thinking Two-Stages Ahead, then they may take into account how the utility of content consumption in the future changes depending on the length of the break. In our model, if breaks increase the utility of consumption, then we would expect longer breaks to lead to greater utility of future content consumption, but longer breaks may also reduce learning and expected scores on upcoming quizzes – hence the “tension” in the model between the utility of consumption and the utility from knowledge accumulation. Thus, endogenizing break lengths under the Two-Stages Ahead assumption is another way that our model could capture temporal bingeing patterns, since individuals who are taking fewer breaks in order to learn more in the future are temporal bingeing, by our definition.

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

In summary, our paper 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.

9 References

- Alba, Joseph W., and J. Wesley Hutchinson (1987), "Dimensions of Consumer Expertise," *Journal of Consumer Research*, 13(4): 411-454.
- Becker, Gary S. and Kevin M. Murphy (1988), "A Theory of Rational Addiction," *Journal of Political Economy*, 96(4): 675-700.
- Ben-Akiva, Moshe, and Bruno Boccara (1995), "Discrete Choice Models with Latent Choice Sets," *International Journal of Research in Marketing*, 12(1): 9-24.
- Bonezzi, Andrea, C. Miguel Brenda, and Matteo De Angelis (2011), "Stuck in the Middle: The Psychophysics of Goal Pursuit," *Psychological Science*, 22(5): 607-612.
- Brooks, Stephen, and Andrew Gelman (1998), "General Methods for Monitoring Convergence of Iterative Simulations," *Journal of Computational and Graphical Statistics*, 7(4): 434-455.
- Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong (2004), "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, 119(3): 861-898.
- Chou, Ting-Jui, and Chih-Chen Ting (2003), "The Role of Flow Experience in Cyber-Game Addiction," *Cyberpsychology & Behavior*, 6(6): 663-675.
- Deloitte (2016), "Deloitte: 70 percent of US consumers binge watch TV, bingers average five episodes per sitting," available at [<https://www2.deloitte.com/us/en/pages/about-deloitte/articles/press-releases/digital-democracy-survey-tenth-edition.html>]
- Ellenbogen, Jeffrey M., Justin C. Hubert, Robert Stickgold, David F. Dinges, and Sharon L. Thompson-Schill (2006), "Interfering with Theories of Sleep and Memory: Sleep, Declarative Memory, and Associative Interference," *Current Biology*, 16 (13): 1290-1294.
- Erdem, Tülin, and Michael Keane (1996), "Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15(1): 1-20.
- Fader, Peter S., Bruce G. S. Hardie, and Ka Lok Lee, "RFM and CLV: Using Iso-Value Curves for Customer Base Analysis," *Journal of Marketing Research*, 42(4): 415-430.
- Gabaix, Xavier, David Liaison, Guillermo Moloche, and Stephen Weinberg (2006), "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," *The American Economic Review*, 96(4):1043-1068.
- Gelman, Andrew, and Donald B. Rubin, "Inference From Iterative Simulation Using Multiple Sequences," *Statistical Science*, 7(4): 457-472.

Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin (2014), *Bayesian Data Analysis*, 3rd ed. Boca Raton, FL: CRC Press.

Gilbride, Timothy J., and Greg M. Allenby (2004), "A Choice Model with Conjunctive, Disjunctive, and Compensatory Screening Rules," *Marketing Science*, 23(3): 391-406.

Givon, Moshe (1984), "Variety Seeking Through Brand Switching," *Marketing Science*, 3(1): 1-22.

Gordon, Brett R. and Baohong Sun (2015), "A Dynamic Model of Rational Addiction: Evaluating Cigarette Taxes," *Marketing Science*, 34(3): 452-470.

Greifender, Rainer, Herbert Bless, and Michel Tuan Pham (2011), "When Do People Rely on Affective and Cognitive Feelings in Judgment? A Review," *Personality and Social Psychology Review*, 15(2): 107-141.

Inman, J. Jeffrey (2001), "The Role of Sensory-Specific Satiety in Attribute-Level Variety Seeking," *Journal of Consumer Research*, 28(1): 105-120.

Jurgensen, John (2013), "Netflix Says Binge Viewing is No 'House of Cards.'" *The Wall Street Journal*, available at [<https://www.wsj.com/articles/netflix-says-binge-viewing-is-no-8216house-of-cards8217-1386897939?tesla=y>]

Kahn, Barbara E., Manohar U. Kalwani, and Donald G. Morrison, "Measuring Variety-Seeking and Reinforcement Behaviors Using Panel Data," *Journal of Marketing Research*, 23(2): 89-100.

Kivetz, Ran, Oleg Urminsky, and Yuhuang Zheng (2006), "The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusionary Goal Progress, and Customer Retention," *Journal of Marketing Research*, 43(1): 38-58.

Kumar, V., Rajkumar Venkatesan, Tim Bohlin, and Denise Beckmann, "The Power of CLV: Managing Customer Lifetime Value at IBM," *Marketing Science*, 27(4): 585-599.

Loewenstein, George F., and Drazen Prelec (1993), "Preferences for Sequences of Outcomes," *Psychological Review*, 100 (1): 91-108.

McAlister, Leigh (1982), "A Dynamic Attribute Satiation Model of Variety-Seeking Behavior," *Journal of Consumer Research*, 9(2): 141-150.

Muth, John F (1961), "Rational Expectations and the Theory of Price Movements," *Econometrica: Journal of the Econometric Society*, 29(3): 315-355.

Nelson, Leif D., and Tom Meyvis (2008), "Interrupted Consumption: Disrupting Adaptation to Hedonic Experiences," *Journal of Marketing Research*, 45 (6): 654-664.

- Nelson, Leif D., Tom Meyvis, and Jeff Galak (2009), "Enhancing the Television-Viewing Experience Through Commercial Interruptions," *Journal of Consumer Research*, 36(2): 160-172.
- Schweidel, David A., and Wendy W. Moe (2016), "Binge Watching and Advertising," *Journal of Marketing*, 80(5): 1-19.
- Seetheraman, P.B., and Pradeep Chintagunta (2003), "The Proportional Hazard Model for Purchase Timing: A Comparison of Alternative Specifications," *Journal of Business & Economic Statistics*, 21(3): 368-382.
- Van den Bulte, Christophe (2000), "New Product Diffusion Acceleration: Measurement and Analysis," *Marketing Science*, 19 (4): 366-380.
- Whittlesea, Bruce W., and Jason P. Leboe (2000), "The Heuristic Basis of Remembering and Classification: Fluency, Generation, and Resemblance," *Journal of Experimental Psychology: General*, 129(1): 84-106.
- Wind, Jerry, and Vijay Mahajan, "Issues and Opportunities in New Product Development: An Introduction to the Special Issue," *Journal of Marketing Research*, 34 (1): 1-12.
- Zauberman, Gal, and John G. Lynch Jr., (2005), "Resource Slack and Propensity to Discount Delayed Investments of Time Versus Money," *Journal of Experimental Psychology: General*, 134(1): 23-37.
- Zhang, Yao, Eric T. Bradlow, and Dylan S. Small (2015), "Predicting Customer Value Using Clumpiness: From RFM to RFMC," *Marketing Science*, 34(2): 195-208.

10 Appendices

Appendix A: List of lecture and quiz names and video run times

Table 5: 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 6: List of lectures and quizzes in Operations

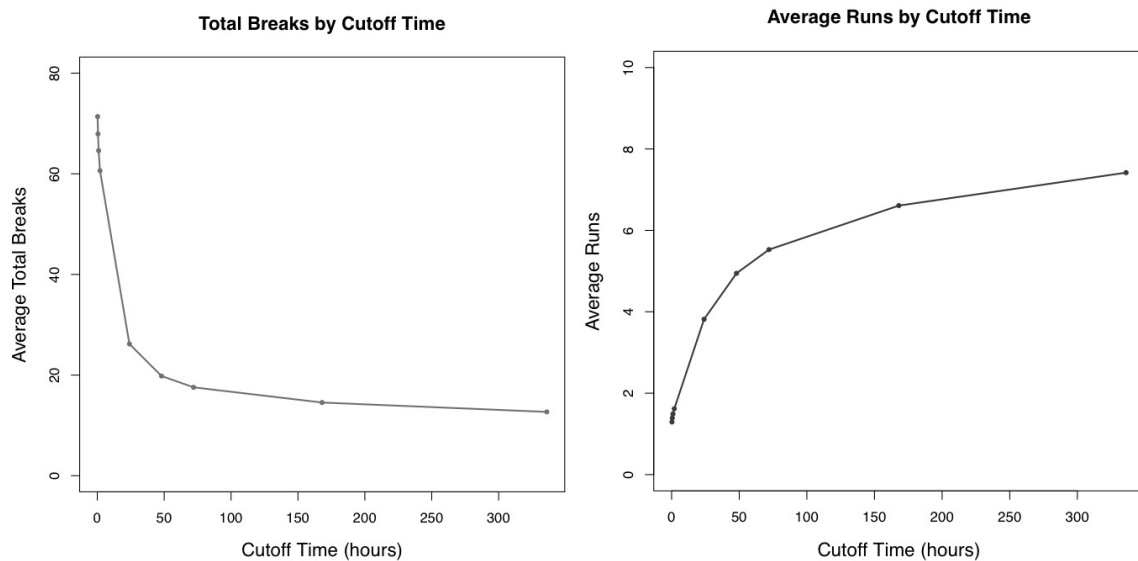
Number	Week	Type	Name	Time
1	1	Lecture	Intro Session 1	7:55
2	1	Lecture	Intro Session 2	7:56
3	2	Lecture	Module 2 Session 1 Video	9:29
4	2	Lecture	Module 2 Session 2 Video	11:24
5	2	Lecture	Module 2 Session 3 Video	15:21
6	2	Lecture	Module 2 Session 4 Video	6:35
7	2	Lecture	Module 2 Session 5 Video	7:25
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	

Appendix B: 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.

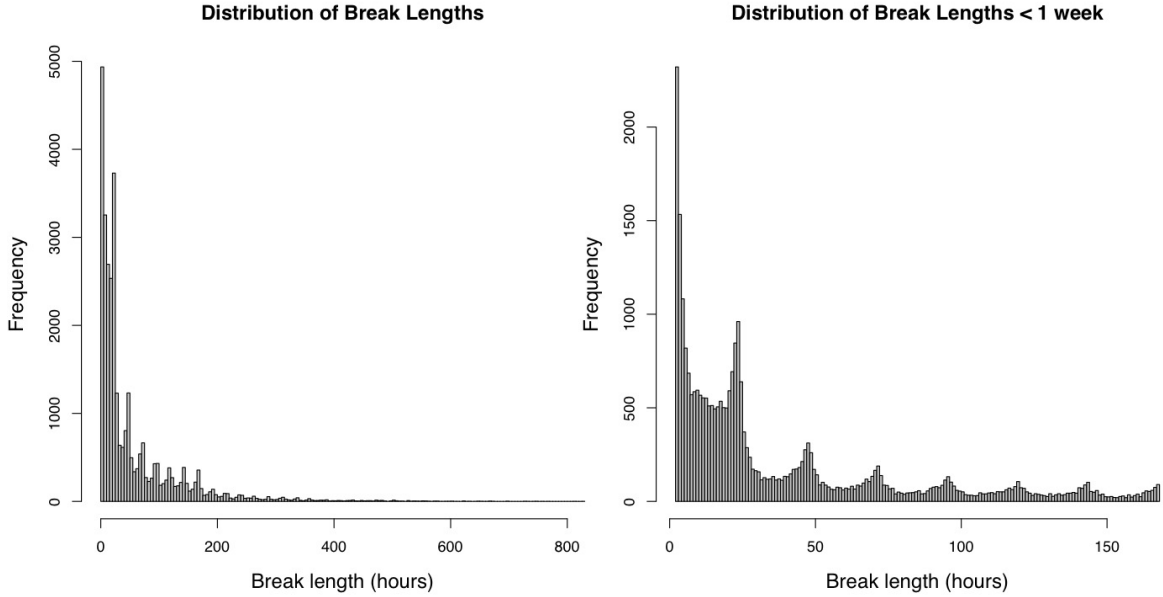
Figure 18 plots the total breaks and average lecture/quiz runs, averaged across individuals, for each cutoff time. We see that there is a distinct jump between 1 and 24 hours, with the curves flattening out around 48 hours. This suggests that there may be a large number of breaks that are less than 24 hours long. So we decided on 1 hour as the break cutoff length.

Figure 18: 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 19 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 19: Distribution of break lengths



Appendix C: Two-Stages Ahead Parameterization

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 $j + 1$:

$$F = \log\left(e^{u_M[j+1]} + e^{u_O[j+1]} + e^{u_B[j+1]}\right) \quad (11)$$

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 .

Appendix D: 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 7 shows the true mean values for the full model with the following constructs: Intercepts, Week, Goal Progress, Completion, Knowledge, and One-Stage Ahead. 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 7: 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	0.06	[-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

Appendix E: 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 8, 9, and 10 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 8: 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 9: 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)***	0.00 (0.00)
$\log(T_{B3})$	0.01 (0.01)	0.01 (0.00)	0.01 (0.02)	0.02 (0.01)***
$\mathbb{1}(G_{MQ} = 1)$	-0.00 (0.01)	-0.04 (0.01)***	0.03 (0.01)*	0.02 (0.01)**
$\mathbb{1}(G_{ML} = 1)$	0.01 (0.01)	-0.04 (0.00)***	-0.01 (0.01)	-0.02 (0.01)***
$\mathbb{1}(G_{OQ} = 1)$	-0.05 (0.02)**	-0.00 (0.00)	-0.00 (0.01)	-0.01 (0.01)
$\mathbb{1}(G_{OL} = 1)$	-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 10: 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)
$\mathbb{1}(G_{MQ} = 1)$	-0.00 (0.01)	-0.04 (0.01)***	0.03 (0.01)*
$\mathbb{1}(G_{ML} = 1)$	0.01 (0.01)	-0.04 (0.00)***	-0.01 (0.01)
$\mathbb{1}(G_{OQ} = 1)$	-0.05 (0.02)**	-0.00 (0.00)	-0.00 (0.01)
$\mathbb{1}(G_{OL} = 1)$	-0.00 (0.02)	0.01 (0.01)	0.00 (0.02)
Adj- R^2	0.02	0.15	0.08

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